Interpretable Data-Driven Modeling Reveals Complexity of Battery Aging

Bruis van Vlijmen[†1,2], Vivek Lam[†1,2], Patrick A. Asinger[†3], Xiao Cui[†1,2], Devi Ganapathi[1,2], Shijing Sun[4], Patrick K. Herring[4], Chirranjeevi Balaji Gopal[4], Natalie Geise[2], Haitao D. Deng[1,2], Henry L. Thaman[1,2], Stephen Dongmin Kang[1], Amalie Trewartha[4], Abraham Anapolsky[4], Brian D. Storey[4], William E. Gent[1,2], Richard D. Braatz[⋆3] and William C. Chueh[⋆1,2]

[1]Department of Materials Science and Engineering, Stanford University, Stanford, CA, USA.

[2]Applied Energy Division, SLAC National Accelerator Laboratory, Menlo Park, CA, USA.

[3]Department of Chemical Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA.

[4]Toyota Research Institute, Los Altos, CA, USA.

Abstract

To reliably deploy lithium-ion batteries, a fundamental understanding of cycling and aging behavior is critical. Battery aging, however, consists of complex and highly coupled phenomena, making it challenging to develop a holistic interpretation. In this work, we generate a diverse battery cycling dataset with a broad range of degradation trajectories, consisting of 363 high energy density commercial $Li(Ni,Co,Al)O_2$/Graphite + SiO_x cylindrical 21700 cells cycled across 218 unique cycling protocols. We consolidate aging via 16 mechanistic state-of-health (SOH) metrics, including cell-level performance metrics, electrode-specific capacities/state-of-charges (SOCs), and aging trajectory descriptors. Through the use of interpretable machine learning and explainable features, we deconvolute the high-dimensional correlations that contribute to battery degradation. This generalizable data-driven mechanistic framework reveals the complex interplay between cycling conditions, degradation modes, and SOH, representing a holistic approach towards understanding battery aging.

Keywords: lithium-ion batteries, machine learning, data analytics

1 Introduction

Lithium-ion batteries are a key enabler for electrifying transportation and decarbonizing the electricity grid [1–8]. Optimizing new battery designs is challenging due to the need to simultaneously meet many

---

[†]These authors contributed equally to this work

[⋆]Corresponding Author. e-mail: wchueh@stanford.edu; braatz@mit.edu

performance targets while satisfying design constraints. Improving battery lifetime is especially difficult due to the slow, nonlinear, and coupled physics of the aging process [9–18]. It is time and resource consuming to observe the impact that design choices have on battery life and understand *why* one battery degrades more rapidly than another.

Characterization at the materials and cell level generates a mechanistic understanding of battery aging [19–22]; however, the throughput is relatively low [23]. In recent years, machine learning (ML) techniques have been developed to analyze battery aging through a data-driven lens [24–39]. While ML techniques are high in throughput, a purely data-driven approach overlooks key scientific and engineering insights. Despite the predictive power of complex black-box ML models (e.g., deep learning), the relationships between cycling conditions and battery aging mechanisms are unclear. On the other hand, physics-based electrochemical simulations, such as the Doyle-Fuller-Newman model [40–42], are physically interpretable. Nonetheless, predicting battery lifetime under unseen conditions remains challenging due to the complexity of interconnected aging phenomena [14] and model parameter identifiability [43]. Yet another approach are mechanistic models which involve estimation of electrode capacities and lithium inventory [44–46]. These models capture aggregate physical mechanisms with fewer model parameters than physics-based simulations [47–51]. Tracking electrode capacities independently provides a clear picture of what types of degradation occur under various operating conditions [23, 52].

A challenge with developing and benchmarking battery aging models is that publicly available datasets do not contain a wide range of operating conditions. Existing datasets are typically collected with specific applications in mind [53–57]. For example, Attia, Severson, and colleagues focused on optimizing electric vehicle fast charging protocols [58, 59]. Diao et al. examined different temperatures to understand how temperatures accelerate battery aging [60]. Paulson, Ward, and colleagues tested various cell chemistries to understand the differences in their aging and build transferable ML models [34, 61]. As a final example, Wildfeuer et al. examined different state-of-charge (SOC) ranges and temperatures in both cycling and calendar aging tests to investigate different experimental factors, but did not apply ML techniques to analyze the large dataset [62]. There remains a significant gap in interpretable data-driven models that can be comprehensively applied to large datasets. The lack of available data spanning many use cases, including a wide range of SOC, charging, and discharging protocols, further compounds this challenge.

In this work, we develop a physically interpretable, data-driven understanding of lithium-ion battery aging. We generate a large dataset consisting of 363 cells under 218 unique cycling conditions spanning diverse use cases and aging trajectories. We apply interpretable ML with explainable features to track

16 mechanistic SOH metrics. With this framework, we begin to answer three principal questions: 1) how do cells degrade? 2) when will cells degrade? and, most critically, 3) what factors influence degradation? We demonstrate that physically meaningful features must be used in combination with methods that robustly extract feature importance [63–67]. Our approach of using interpretable features also reveals which and how mechanistic SOH metrics can be predicted from early cycle data, addressing the challenge that features used for early prediction tasks are difficult to meaningfully interpret, such as the features employed in Severson et al. [58]. With our explainable data-driven model, we analyze and understand battery aging further than would be possible with either a data-driven or physics-based approach alone. More generally, constraining ML models to use features that have clear physical meaning dramatically enhances interpretability and explainability, complementing purely data-driven featurization approaches.

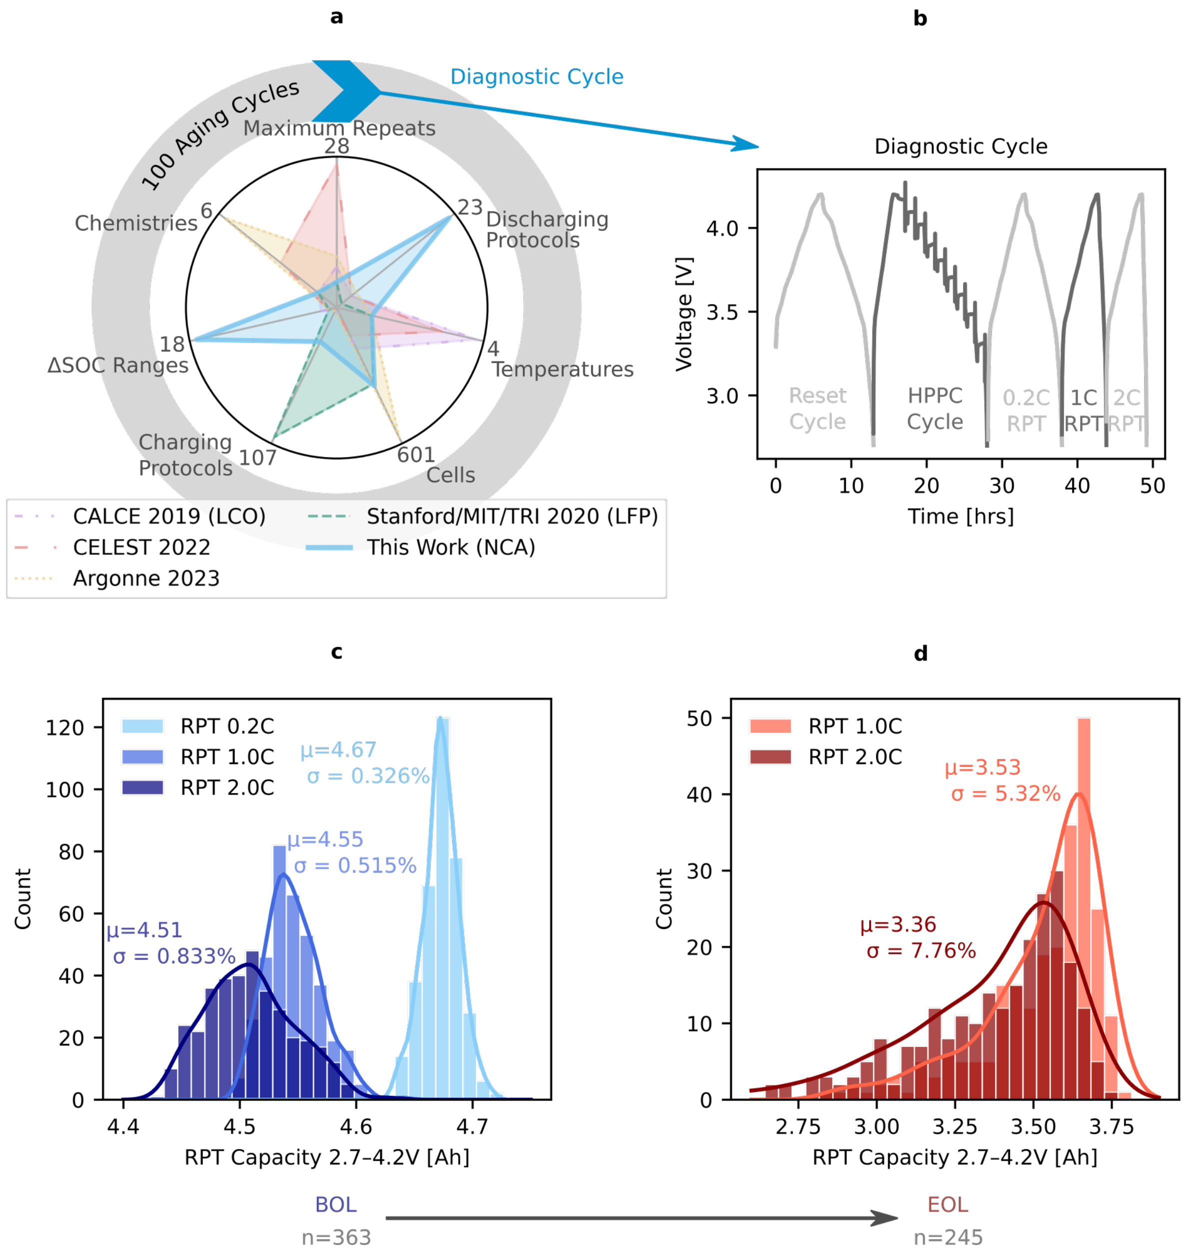

Fig. 1: Overview of dataset. a) The scope of our dataset across various cycling conditions is highlighted in the inscribed spider plot in blue compared to other large, publicly available battery cycling datasets [34, 53, 58–62, 68]. All batteries in this dataset are cycled at 25°C. The cycling experiment structure is shown schematically with the loop surrounding the spider plot. Individual cells go through a diagnostic "checkup" cycle, followed by 100 aging cycles repeating until end of life (EOL). **b**) The diagnostic cycle consisting of a reset cycle, a hybrid pulse power characterization (HPPC) [69], and three rate performance tests (RPTs) at 0.2C, 1C, and 2C discharge currents (see SI Table S2 for full conditions). Mechanistic SOH metrics are extracted from various parts of this diagnostic cycle data (see SI Section S.3 for further details). **c**) The distribution of rate-dependent capacities at beginning of life (BOL). Means and coefficients of variation are included in the plot showcasing the tight distribution at BOL. **d**) The distribution of rate-dependent capacities at end of life (EOL, defined by 0.2C RPT capacity reaching 80% of the nominal capacity, 4.84Ah). The broadened distribution showcases diverse aging and highlights the limitations of using a single mechanistic SOH metric such as the low-rate capacity (for further information on BOL to EOL variability see SI Section S.5) [70].

2 Comprehensive Aging Characterization

Our dataset contains electrochemical data from 363 Li(Ni,Co,Al)O_2/Graphite + SiO_x cylindrical 21700 cells cycled for over 2 years (Fig. 1a). To induce diverse aging trajectories, we explore a broad range of cycling voltage windows and charging and discharging rates (see Section S.2.2 for details). To cleanly compare the effects of different cycling conditions, we apply a standardized, periodic diagnostic cycle to comprehensively probe SOH over cell lifetime (typically every 100 aging cycles – Fig. 1b). Given the variation of voltage windows and charge and discharge rates throughout the dataset, we compare cell lifetimes using capacity throughput-based equivalent full cycles (EFCs) [71]. In total, we examine 218 unique aging protocols, with EFCs at end of life (EOL) ranging from 44 to 994 (or 63 to 4,641 cycles). EOL is defined as when the 0.2C rate-specific capacity, $\mathbf{Q_{RPT,0.2C}}$, reaches 80% of the nominal capacity (where 1C is 4.84A). Fig. 1a compares the diversity of our cycling conditions to other public datasets. Critically, we realize that a single health metric, such as low-rate capacity, does not capture all facets of degradation (Fig. 1c,d, and SI Section S.5). To address this gap, we automatically calculate and track 16 mechanistic SOH metrics (see SI Section S.1.1 for summary of abbreviations).

We first quantify six cell-level performance metrics: 1) total EFCs at EOL, 2) 1C rate-specific capacity: $\mathbf{Q_{RPT,1C}}$, 3) 2C capacity: $\mathbf{Q_{RPT,2C}}$, 4) ohmic resistance: $\mathbf{R_{ohm}}$, 5) charge transfer resistance: $\mathbf{R_{ct}}$, and 6) polarization resistance: $\mathbf{R_p}$. We calculate resistances through pulse measurements performed during the hybrid pulse power characterization (HPPC) sequence of the diagnostic cycle at various SOCs and timescales (see SI Section S.3.2 for definitions and calculation details for resistance metrics). Unless otherwise specified, the resistances reported are at 50% SOC.

Second, to determine electrode-specific capacities/SOCs, we implement a mechanistic model-fitting algorithm to extract seven interpretable quantities (see Methods for details): 1) negative electrode capacity: $\mathbf{Q_{NE}}$, 2) positive electrode capacity: $\mathbf{Q_{PE}}$, 3) lithium capacity: $\mathbf{Q_{Li}}$, 4) State of charge of the negative electrode near the full cell charged state: $\mathbf{SOC_{NE,4.0V}}$, 5) state of charge of the negative electrode in the discharged state: $\mathbf{SOC_{NE,2.7V}}$, 6) state of charge of the positive electrode near the charged state: $\mathbf{SOC_{PE,4.0V}}$, and 7) state of charge of the positive electrode in the discharged state: $\mathbf{SOC_{PE,2.7V}}$. We select the latter four quantities because the electrode-specific SOC near the fully discharged and fully charged states can dominate aging.

These cell-level and electrode-specific metrics are calculated at every diagnostic cycle for each cell and tracked from beginning of life (BOL) to EOL. As would be expected for commercial cells, these metrics have low variability at BOL (Fig. 1c). Importantly, this low variability also confirms that the staggered start of cycling (resulting in different calendar aging in the discharged state) does not contribute significantly to the initial conditions of the cells (SI Section S.5). However, by EOL there is high variation in the rate capability, resistance, and electrode-specific capacities/SOCs (Fig. 1d, and SI Fig. S8). This observation underscores the importance of using a comprehensive set of SOH metrics, and confirms that the cycling conditions in this work induce a wide range of degradation trajectories.

In addition to probing cell-level and electrode-specific metrics with each diagnostic cycle, we also quantify the aging trajectory over the entire battery lifetime [9, 72]. We define three trajectory descriptors: 1) knee indicator: **Knee**, 2) resistance growth factor: **R"**, and 3) negative/positive capacity (N/P) ratio: **NP Ratio**. The knee indicator describes a sudden and accelerated capacity-based degradation (i.e., a knee in the capacity vs. cycle number curve) with knee indicator > 0 if a knee exists at any point in the cell lifetime. The resistance growth factor captures the curvature of resistance with respect to EFCs, indicating whether resistance grows at an accelerating or decelerating rate during cycling. Finally, the NP Ratio captures the ratio of the estimated $\mathrm{Q_{NE}}$ and $\mathrm{Q_{PE}}$. SI Section S.7 details the calculations of these trajectory descriptors.

We combine these 16 cell-level performance metrics, electrode-specific capacities/SOCs, and trajectory descriptors (collectively called mechanistic SOH metrics) and comprehensively quantify battery aging. By concurrently assessing these metrics, we reveal their relationships to 218 cycling conditions to develop a holistic understanding of aging. Fig. 2 visualizes selected metrics calculated on all cells in the dataset.

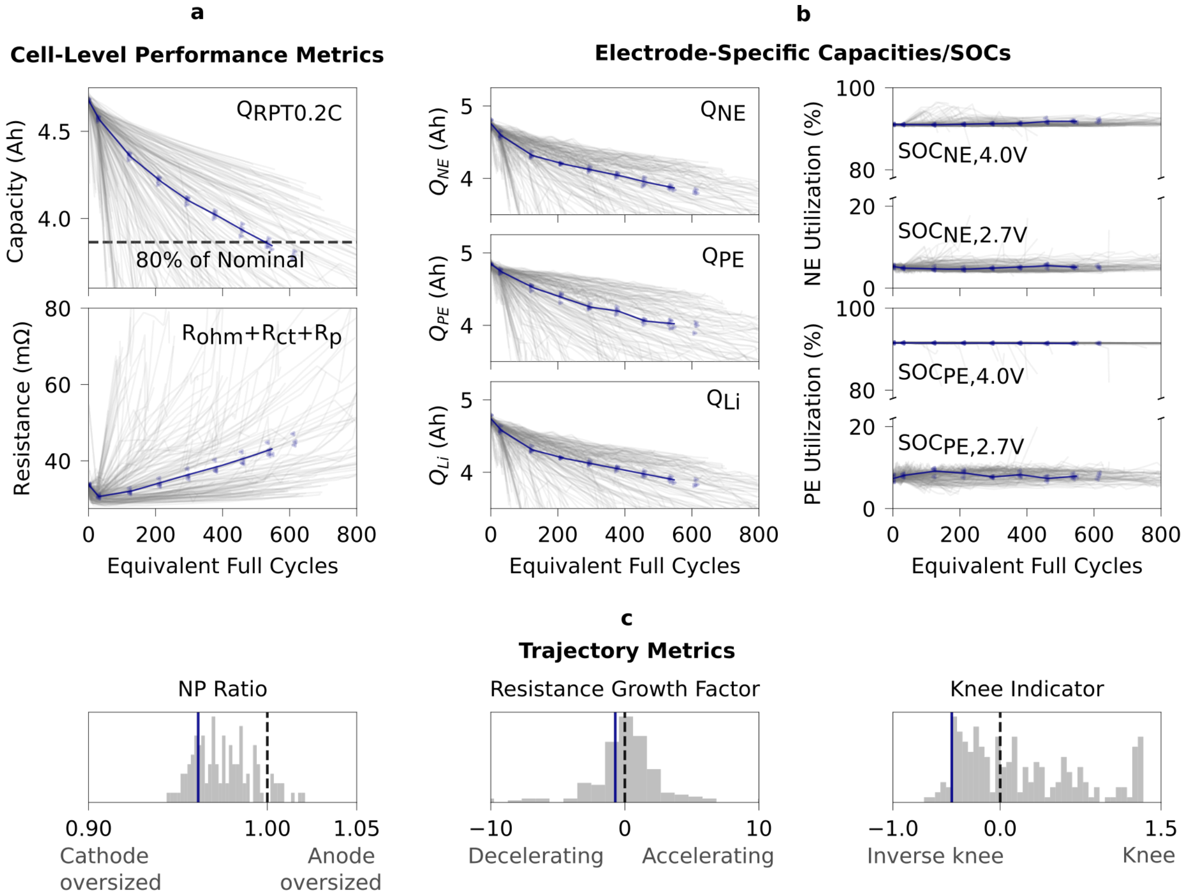

Fig. 2: Mechanistic SOH metric trajectories. a) The cell-level performance metrics column show the trajectories of selected performance metrics: the 0.2C RPT discharge capacity ($Q_{RPT,0.2C}$) (top) and the combination of R_{ohm}, R_{ct}, R_p (R_{tot}) at 50% SOC (bottom). **b**) The electrode-specific capacities/ SOCs column depict the trajectories of electrode-specific capacities, Q_{NE}, Q_{PE}, and Q_{Li}, on the left. A utilization plot showing electrode-specific SOCs at the charged and discharged state is shown on the right. **c**) The trajectory metrics row shows histograms of the values for the NP ratio, resistance growth factor, and knee indicator. The highlighted protocol (in dark blue) represents $CC_{discharge} = 0.2C$, $CC_1 = CC_2 = 0.2C$, $V_{charge} = 4.2V$ and $V_{discharge} = 2.7V$ aging conditions. This protocol has four experimental repeats shown by the scatter markers with the solid line representing the mean trajectory. The gray lines in the background showcase the mean trajectory of all other unique protocols. This protocol appears as a blue vertical bar in the trajectory metric histograms.

3 Impact of Cycling Conditions on Mechanistic SOH Metrics

By varying six cycling parameters across this dataset (SI Table S4), we induce a diverse range of EOL states and trajectory descriptors (Fig. 2a-c). To understand the impact of cycling conditions on mechanistic SOH metrics, we construct nonlinear random forest ML models, and then employ Shapley additive explanations (SHAP) analysis [73] to interpret these models.

We first develop descriptive models using cycling protocol parameters alone as inputs. These "protocol models" predict a single cell-level performance metric, such as an electrode-specific capacity or a trajectory descriptor at EOL conditions. Fig. 3a schematically depicts the structure of these models, with the cycling protocol conditions as the input features and the EOL mechanistic SOH metrics as the target outputs. With only cycling conditions as input features, the EFC model attains good performance on the training/test set (SI Section S.9.5 for performance, and SI Section S.9.1 on train/test split). Well-performing models are critical in order to extract the correct feature importance. To understand the impact of the various cycling conditions, we investigate feature importance using SHAP analysis [74]. Fig. 3b shows an example of SHAP feature importance for predicting EFC. In Fig. 3c, we devise a 96 element aging matrix representation that comprehensively visualizes how cycling conditions affect each mechanistic SOH metric, where the color indicates the magnitude of feature importance for all cycling conditions.

While no single cycling parameter dominates all mechanistic SOH metrics, it was surprising that many of the metrics are either primarily determined by a single cycling parameter, or a combination of features (e.g., the NP Ratio, a combination of $\mathrm{Q_{PE}}$ and $\mathrm{Q_{NE}}$) that are primarily influenced by a single cycling parameter . For example, the cell-level performance metrics $\mathrm{Q_{RPT,1C}}$ and $\mathrm{Q_{RPT,2C}}$ are dominated by $\mathrm{CC_2}$, the resistances $\mathrm{R_{ohm}}$ and $\mathrm{R_{ct}}$ are dominated by $\mathrm{V_{charge}}$, while $\mathrm{R_p}$ is dominated $\mathrm{CC_1}$. Some more convoluted metrics, such as the EFC, depend on multiple parameters; both $\mathrm{CC_1}$ and $\mathrm{V_{discharge}}$ are about equally important. Surprisingly, $\mathrm{V_{discharge}}$ and $\mathrm{t_{CV}}$ do not dominate aging (within the bounds of this dataset) for most of the mechanistic SOH metrics (except for EFC), despite previous reports stating their importance [75, 76].

For the electrode-specific capacities, both the positive and negative electrode are strongly affected by the magnitude of the current in the direction of lithiation. This current is $\mathrm{CC_{discharge}}$ for $\mathrm{Q_{PE}}$, and $\mathrm{CC_1}$ for $\mathrm{Q_{NE}}$. Since $\mathrm{Q_{Li}}$ also depends most strongly on $\mathrm{CC_1}$, it is possible that $\mathrm{CC_1}$ triggers mechanisms that age both $\mathrm{Q_{NE}}$ and $\mathrm{Q_{Li}}$, such as solid-electrolyte interface (SEI) growth. The electrode-specific SOCs, calculated from electrode-specific capacities, depend most strongly on $\mathrm{CC_1}$ and $\mathrm{CC_{discharge}}$ (the most important features of the electrode-specific capacities), approximately equally. Finally, for the trajectory metrics, the knee indicator depends most strongly on $\mathrm{CC_1}$ and $\mathrm{CC_{discharge}}$, the resistance growth factor (R") on $\mathrm{CC_1}$ and $\mathrm{CC_2}$, and NP ratio on the values it was constructed from, in this case, both the dominant feature from $\mathrm{Q_{NE}}$ ($\mathrm{CC_1}$), and from $\mathrm{Q_{PE}}$ ($\mathrm{CC_{discharge}}$). For detailed information on influence of cycling conditions on mechanistic SOH metrics, as well as model performance see SI Section S.9.5

With our aging matrix representation generated by interpretable ML, a battery cell designer could more intelligently identify aging mechanisms and design cycling limits. For example, if it is important to prevent capacity knees, from this analysis, we see that modifying CC_1 and $CC_{discharge}$ will have the greatest impact, whereas modifying the $V_{discharge}$ would not be effective.

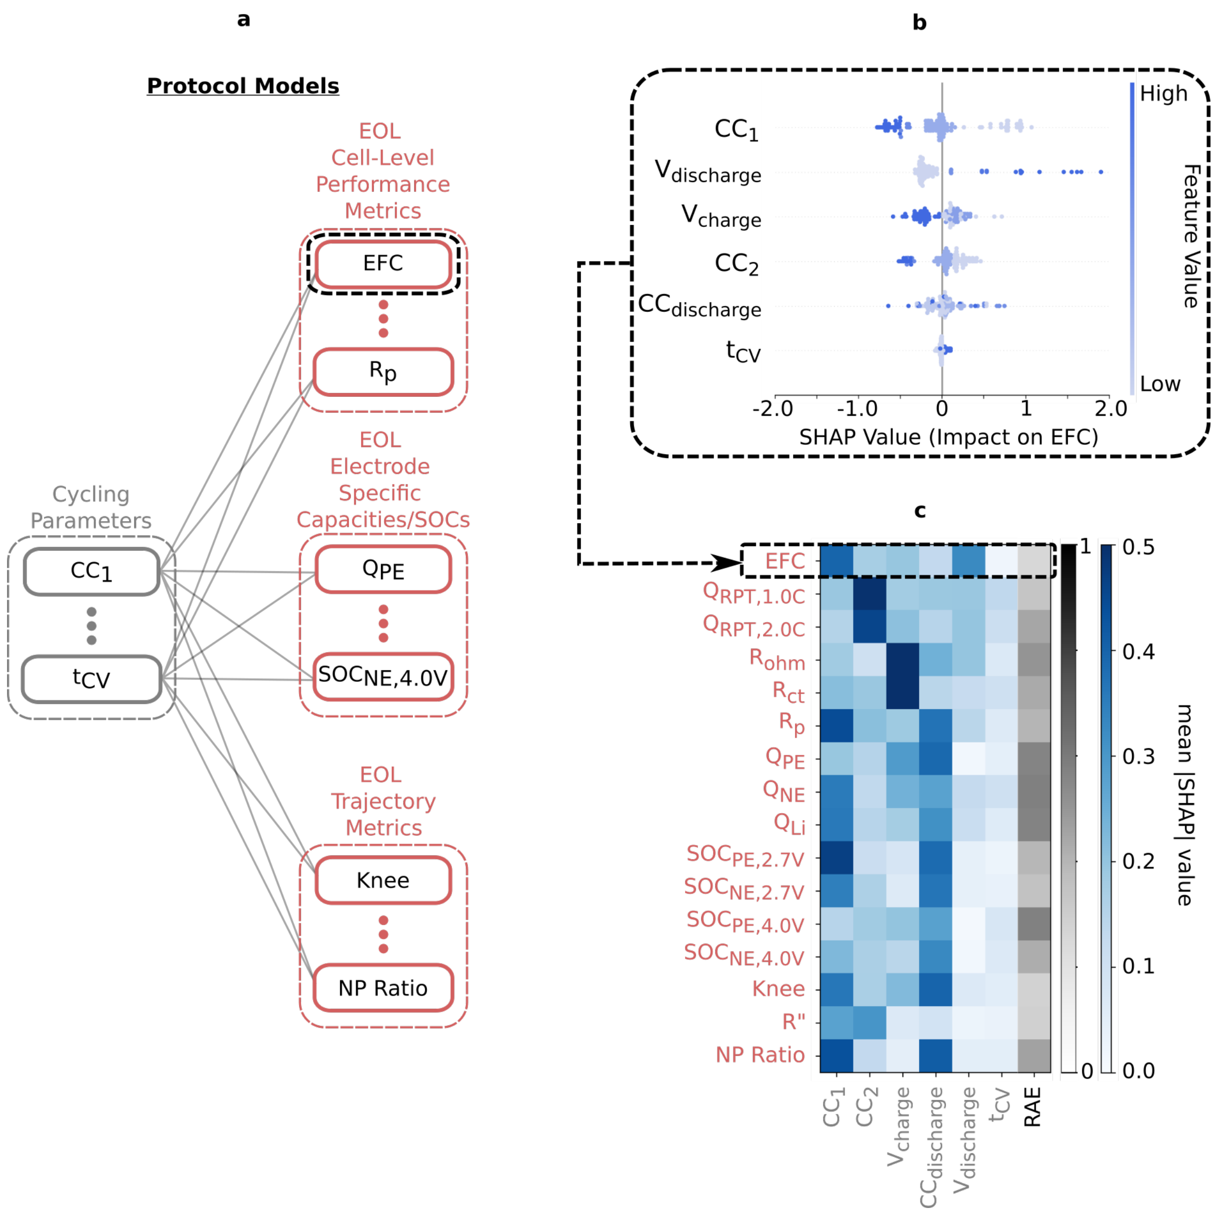

Fig. 3: Impact of cycling conditions. a) Schematic of inputs and outputs of the protocol models. Gray rectangles indicate cycling parameters, and red rectangles indicate the EOL mechanistic SOH metrics. **b**) SHAP feature importances for the protocol model predicting EFCs, marked by a dashed box in **a**. The color indicates the feature value, and horizontal location indicates the SHAP value impact on EFCs. Features are listed in descending order of importance. **c**) Replicating this approach for each mechanistic SOH metric, the matrix shows the mean absolute SHAP value of each cycling condition for each degradation metric. Darker hue indicates stronger dependence. Additionally, the RAE (relative absolute error) column indicates the test error of the models trained to predict a particular mechanistic SOH metric. SI Section S.9.5 shows the parity plot and SHAP beeswarm plot for each mechanistic SOH metric. This degradation matrix representation visualizes the impact of cycling conditions on degradation in a high-dimensional space.

4 Fundamental Investigation of Performance Metric Degradation

Having revealed the relationship between 16 mechanistic SOH metrics and cycling conditions using an aging matrix, we now demonstrate the explanatory nature of our framework by answering one important, exemplar question: “how does degradation at specific electrodes contribute to resistance growth in a battery?” Resistance growth during aging can limit the discharge capacity and energy of a battery. However, it is challenging to understand where inside a battery resistance growth originates using only full cell measurements because of the convolution of multiple effects from both electrodes. The resistances of individual electrodes are highly dependent on their respective electrode’s lithiation state and degradation. In addition, as cells age under diverse usage conditions, individual electrodes can go through various degradation pathways such as cathode structural changes [77] and anode solid electrolyte interface (SEI) formation [78]. These changes lead to varying degrees of electrode slippage or SOC shifts, adjusting the relative lithium composition of the cathode and anode at a given full cell SOC (SI Fig. S22).

To understand the complex relationship between electrode degradation and resistance growth in a full cell, we expand on the “protocol model” discussed in the previous section to include EOL electrode-specific capacities/ SOCs metrics as input features. This “explanatory model” aims to learn the relationship between the physically meaningful electrode-level features and the mechanistic SOH metric of interest (Fig. 4a). In this section, we investigate the changes in the electrode-specific capacities/SOCs and resistances with cycling. As such, the model inputs are the changes in SOH metrics from BOL to EOL (represented by Δ). The model output here is the low SOC (30%) total resistance (summation of R_{ohm}, R_{ct}, and R_{p}, SI Section S.3.2). We choose this health metric as the example target of our explanatory model because resistance at low SOCs are typically the largest and limit the discharge capacity.

Fig. 4b lists the most dominant features contributing to the observed total resistance growth. From the SHAP analysis, we observe that two electrode-specific features, $\Delta\mathrm{SOC}_{\mathrm{PE},2.7\mathrm{V}}$ and $\Delta\mathrm{SOC}_{\mathrm{NE},2.7\mathrm{V}}$, are dominant features impacting the total resistance but show opposite relationships with resistance growth (Fig. 4b,c). Surprisingly, negative electrode over-discharging ($\Delta\mathrm{SOC}_{\mathrm{NE},2.7\mathrm{V}} < 0$) leads to lower resistance increase. This is unexpected because electrode kinetics are typically most sluggish at the SOC extremes; therefore, at low SOC, we expect that resistance should increase in the direction of deeper discharge for an electrode [76].

To understand the origin of this effect, we recall how $\Delta\mathrm{SOC}_{\mathrm{PE},2.7\mathrm{V}}$ and $\Delta\mathrm{SOC}_{\mathrm{NE},2.7\mathrm{V}}$ are calculated. These quantities are calculated at a specified full cell voltage (2.7V for this example) and, as a result, are highly correlated (Fig. 4d, SI Section S.8.1). This correlation arises because when one electrode's SOC shifts, regardless of the aging mechanism, the other electrode's SOC must shift in the opposite direction to produce the same measured full cell voltage (SI Fig. S22 explores this in further detail). In general, SHAP is unable to differentiate between highly correlated features, and repeating the SHAP analysis multiple times reveals that either $\Delta\mathrm{SOC}_{\mathrm{PE,2.7V}}$ or $\Delta\mathrm{SOC}_{\mathrm{NE,2.7V}}$ can emerge as the most dominant feature (SI Fig. S20). However, if $\Delta\mathrm{SOC}_{\mathrm{NE,2.7V}}$ is removed from this explanatory model, for example, $\Delta\mathrm{SOC}_{\mathrm{PE,2.7V}}$ appears as the dominant feature (SI Fig. S21). From this analysis, we understand that, while negative electrode over-discharging ($\Delta\mathrm{SOC}_{\mathrm{NE,2.7V}} < 0$) leads to lower resistance increase, the correlated metric positive electrode over-discharging ($\Delta\mathrm{SOC}_{\mathrm{PE,2.7V}} < 0$) leads to higher resistance increase, in line with the understanding that electrode kinetics are most sluggish at SOC extremes. Combining statistical analysis with scientific understanding of battery materials, we rationalize that low SOC resistance rise is dominated by the over-discharging of the positive electrode.

Our framework exemplifies the value of SHAP as a tool for identifying correlations between input features and the target mechanistic SOH metrics. While the ML method alone does not differentiate between the contributions from two highly correlated electrodes, the explainable features together with scientific knowledge helps to hypothesize causation. Although we choose in this section to highlight and analyze low SOC resistance as one example, we emphasize that the approach generalizes to any mechanistic aging feature of interest (SI Section S.8.3).

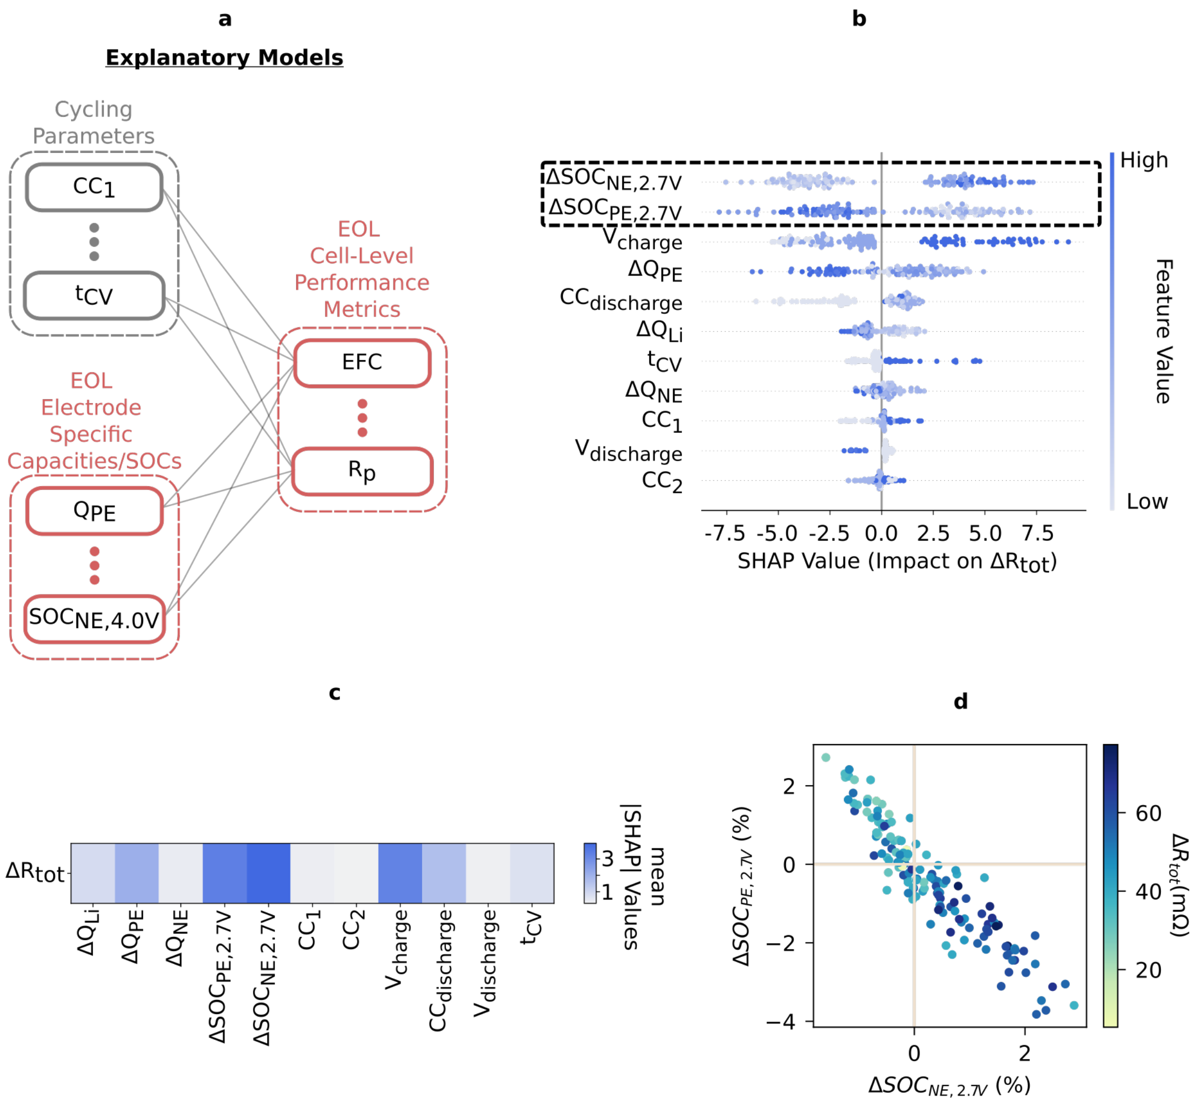

Fig. 4: Analyzing EOL cell-level performance metrics through electrode capacities/SOCs. a) Schematic representation of the explanatory models to understand the degradation of cell-level performance metrics. Gray rectangles indicate cycling parameters, and red rectangles indicate mechanistic SOH metrics that are obtained at EOL. **b**) SHAP feature importance ranking from the random forest model fit on 30% SOC total resistance in descending order. The ΔSOCs are the most important features, but show an opposite relationship with resistance increase. **c**) One example row of a matrix plot summarizing the information in the SHAP analysis. **d**) $\Delta SOC_{PE,2.7V}$ plotted against $\Delta SOC_{NE,2.7V}$ at EOL for 146 cells. Color bar indicates full cell total resistance growth at 30% SOC. The high correlation indicates that feature importance can be convoluted. With knowledge that resistance values of electrodes increase at extremes of the SOC range, we determine that resistance increase is driven by the positive electrode.

5 Early Prediction Using Explainable Features

Finally, we quantify and rationalize the predictive power of explainable features in early cycles, and demonstrate the value of features extracted from the early diagnostic cycles for early prediction of the 16 EOL mechanistic SOH metrics. Building upon our protocol model in which random forest regression models were employed to correlate EOL mechanistic aging features to cycling parameters, we construct a "diagnostic-aided model" that uses both cycling parameters and early values of the mechanistic SOH metrics (specifically, the evolution between the 1st and 3rd diagnostic cycle, Fig. 5a) as inputs to our interpretable ML model. The inclusion of features from early diagnostic cycles differentiates between cells with the same cycling parameters, giving insight into cell-to-cell variability in fixed aging conditions.

We perform similar SHAP analysis as demonstrated in the previous sections on our diagnostic-aided model and present the results in an aging matrix plot in Fig. 5b (see SI Section S.9.5 for parity plots and full shap analysis). For the mechanistic SOH metrics, the diagonal entries of the degradation matrix correspond to self prediction (i.e., predicting the EOL value of a given metric using its early value). Interestingly, while the features on this diagonal might be expected to consistently be the most predictive, this is not always the case. For example, the early prediction of R_{ct} is dominated by V_{charge}. Additionally, the early prediction of EFC is dominated by R_p, rather than by $Q_{RPT0.2C}$; the latter is the metric used to define the EOL cutoff, and thus EFC at EOL. The result highlights the importance of a detailed tracking of battery SOH. While a given degradation mode might dominate the EOL values of certain mechanistic SOH metrics, the best early indicators for the onset of that mode may be a different metric or set of metrics.

Since SHAP analysis cannot differentiate between correlated input features, in order to draw robust conclusions about the importance of early cycle features, it is necessary to also consider a "diagnostic-only" model, excluding cycling parameters as input features (SI Section S.9.4). In principle, this may affect the relative feature importance of the early cycle features which correlate with specific cycling parameters. In addition, this type of model may be preferred in cases where you either do not directly have access to cycling conditions, cycling conditions are kept constant, or the relationship to cycling conditions is not the focus [79]. In this case, the exclusion of cycling conditions does not meaningfully affect the ranking of the feature importances (see SI Section S.9.4 for further details).

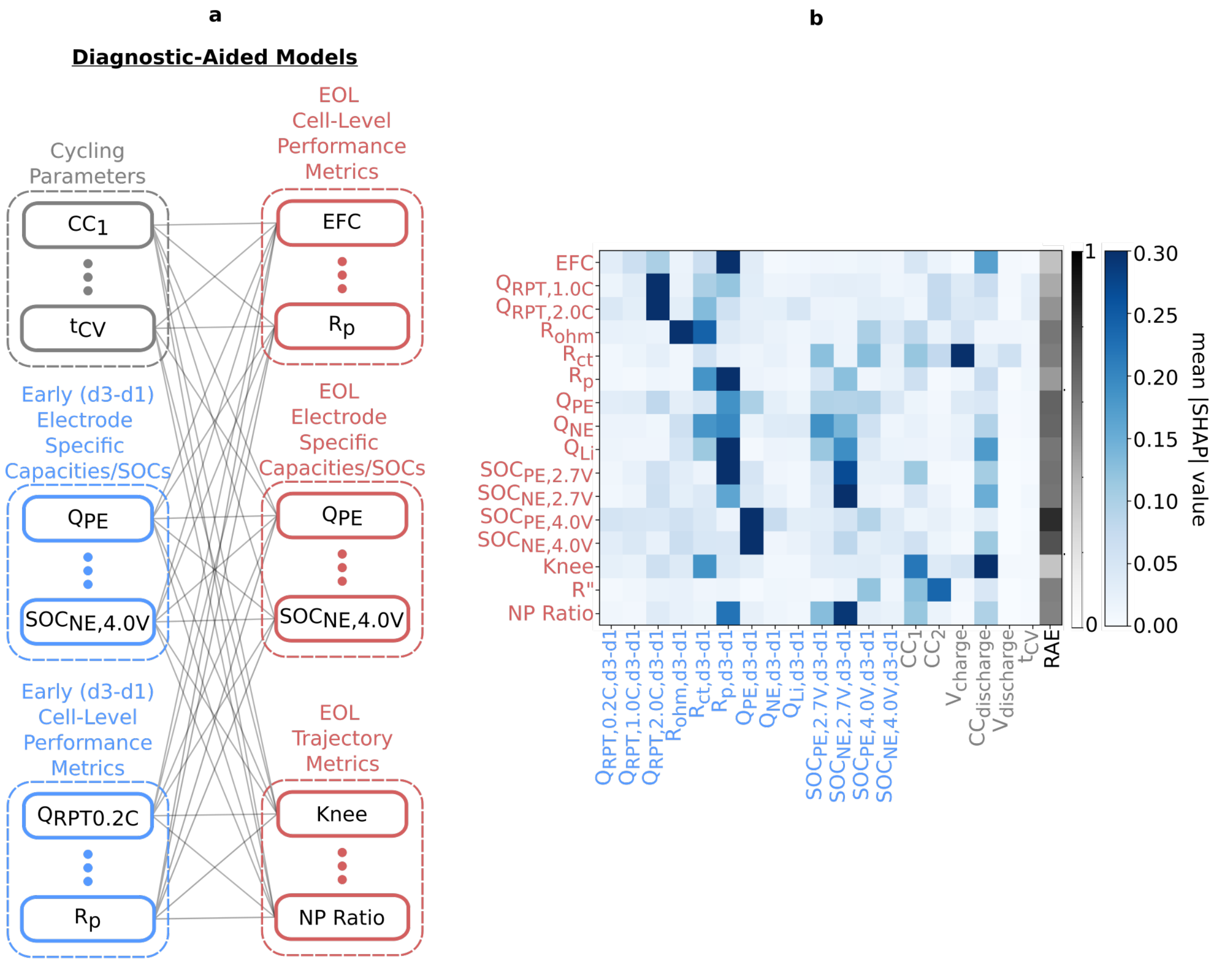

Fig. 5: Early prediction of mechanistic SOH metrics. a) Architecture of the diagnostic-aided model where the blue rectangles indicate values that are extracted early in the cycling and red rectangles are extracted at EOL. **b**) SHAP analysis degradation matrix plot showcasing the importance of cycling protocols and early prediction features on predicting mechanistic SOH metrics. All early prediction features are extracted as the difference of mechanistic SOH metrics from 1st to 3rd diagnostic cycle (annotated as d3-d1).

6 Conclusions

In this study, we develop a holistic framework for revealing and explaining coupled battery aging pathways by combining interpretable ML, physically-derived mechanistic SOH metrics, and a diverse dataset spanning over 200 distinct cycling conditions. By tracking a comprehensive set of 16 mechanistic aging features, we fully describe the battery SOH through an aging matrix, provide insight into battery degradation mechanisms and also identify mechanistic features from early cycles that enable early predictions.

Through our interpretable ML framework, we deepen our physical intuition on battery degradation with a diverse dataset. While interpretable ML tools can be used to generate hypotheses and summaries of the dataset, the findings must be further validated with physical characterization to gain confidence. We urge the field to use the dataset presented here to expand upon this work while keeping interpretability in mind as to enrich our understanding of battery degradation.

7 Methods

7.1 Data Cycling and Generation

All cells in this study were harvested from a newly purchased 2019 Tesla Model 3. These 21700 cylindrical cells were manufactured by Panasonic and tested to have a low-rate capacity of 4.84Ah. The positive electrode is NCA (approximately 90-5-5 composition) and the negative electrode is a graphite-SiO_x blend. Cells were cycled in CSZ ZP-16-2-H/AC environmental chambers set to 25°C, and fitted with 4-point contact cylindrical cell fixtures from Korea Thermo-Tech Co. Ltd. assembled by SpectraPower. The cells were cycled using two 96 channel Maccor Series 4000 battery cyclers.

The cells are subject to two types of cycling: aging cycles and diagnostic cycles. The aging cycle consists of a multi-step CC-CV charge and a CC discharge. Information on cycling protocol, parameters varied and their distribution see SI Section S.2.2. The diagnostic cycle consists of three main portions: a reset cycle, a hybrid pulse power characterization (HPPC) cycle [69], and a rate performance test (RPT) sequence. The reset cycle, resets the transient kinetics due to the aging cycles, HPPC probes resistance at different SOC increments, and the RPT extracts rate-dependent capabilities (Fig. 1b). For information on diagnostic cycle protocol see SI Table S2. This cycling data is automatically backed up to an S3 bucket and subsequently processed through the BEEP processing pipeline for use in analysis [80].

7.2 Differential Voltage Fitting

We implement differential voltage fitting (DVF) to estimate properties of the battery at the electrode level. Similar methodologies have been implemented by other groups [23, 48, 50, 81, 82]. This method extracts electrode capacities and lithium inventory: $\mathrm{Q_{PE}}$, $\mathrm{Q_{NE}}$, and $\mathrm{Q_{Li}}$. Additional information, such as the SOC of either electrode at a full cell specified voltage is further calculated: $\mathrm{SOC_{PE,2.7V}}$, $\mathrm{SOC_{NE,2.7V}}$, $\mathrm{SOC_{PE,4.0V}}$ and $\mathrm{SOC_{NE,4.0V}}$. The DVF routine employed non-invasively probes degradation by fitting the measured 0.2C RPT full cell differential voltage profile with an emulated full cell profile by stretching and translating the voltage profiles of the cathode and anode. Details of the fitting methodology and feature extraction are provided in SI Section S.3.3. Additionally, a comparison to DVF performed at C/40 is provided in S.6.1.

Reference voltage profiles for the cathode and anode are acquired through destructive tear down of the full cell to extract cathode and anode sheets. Portions of the sheet are then cycled in a pouch cells with a lithium counter electrode at various low rates. Details of the experimental electrode extraction and measurement procedure are in SI Section S.4.

7.3 Machine Learning Models

Random forest regression was chosen as the machine learning model of choice for all models in this work due to its ability in capturing non-linear relations with input features. We first generate a train/test split of the data, and cross validation folds on the training split. The cells that go in to the different splits are chosen randomly for the explanatory model, but an inside-of-domain testing scenario for the protocol only, diagnostic-only, and diagnostic-aided models (see SI Section S.9.1 for details). Random forest hyper parameters are optimized via grid search cross validation. From the subsequent trained model we report the RAE metric to accurately compare the prediction performance on different mechanistic SOH metrics of different scales and distributions. To determine feature importances we then use the SHAP python library on the fitted model to extract SHAP values for all features and datapoints (see SI Section S.9.5 for full parity plots and SHAP analysis). To summarize this information, we then take the absolute mean feature importance and report this value in the matrix plots.

8 Data Availability

Raw data and data structured via the BEEP pipeline [80] will be available at the time of publication

9 Code Availability

Code for figure generation, and random forest model building will be available at the time of publication.

References

[1] Abhishek Jaiswal. Lithium-ion battery based renewable energy solution for off-grid electricity: A techno-economic analysis. *Renewable and Sustainable Energy Reviews*, 72:922–934, May 2017.

[2] Ghassan Zubi, Rodolfo Dufo-López, Monica Carvalho, and Guzay Pasaoglu. The lithium-ion battery: State of the art and future perspectives. *Renewable and Sustainable Energy Reviews*, 89:292–308, June 2018.

[3] Charles Lorenzo, Romain Tabusse, David Bouquain, Samuel Hibon, and Daniel Hissel. Study of lithium-ion battery ageing cycled with current profiles from railway applications. In *IEEE Vehicle Power and Propulsion Conference*. IEEE, October 2021.

[4] Tianmei Chen, Yi Jin, Hanyu Lv, Antao Yang, Meiyi Liu, Bing Chen, Ying Xie, and Qiang Chen. Applications of lithium-ion batteries in grid-scale energy storage systems. *Transactions of Tianjin University*, 26(3):208–217, February 2020.

[5] Weidong Chen, Jun Liang, Zhaohua Yang, and Gen Li. A review of lithium-ion battery for electric vehicle applications and beyond. *Energy Procedia*, 158:4363–4368, February 2019.

[6] Boucar Diouf and Ramchandra Pode. Potential of lithium-ion batteries in renewable energy. *Renewable Energy*, 76:375–380, April 2015.

[7] J. B. Dunn, L. Gaines, J. C. Kelly, C. James, and K. G. Gallagher. The significance of Li-ion batteries in electric vehicle life-cycle energy and emissions and recycling's role in its reduction. *Energy & Environmental Science*, 8(1):158–168, 2015.

[8] Boya Zhou, Ye Wu, Bin Zhou, Renjie Wang, Wenwei Ke, Shaojun Zhang, and Jiming Hao. Real-world performance of battery electric buses and their life-cycle benefits with respect to energy consumption and carbon dioxide emissions. *Energy*, 96:603–613, February 2016.

[9] Peter M. Attia, Alexander Bills, Ferran Brosa Planella, Philipp Dechent, Gonçalo dos Reis, Matthieu Dubarry, Paul Gasper, Richard Gilchrist, Samuel Greenbank, David Howey, Ouyang Liu, Edwin Khoo, Yuliya Preger, Abhishek Soni, Shashank Sripad, Anna G. Stefanopoulou, and Valentin Sulzer. Review—"Knees" in lithium-ion battery aging trajectories. *Journal of The Electrochemical Society*, 169(6):060517, June 2022.

[10] Yi Li, Kailong Liu, Aoife M. Foley, Alana Zülke, Maitane Berecibar, Elise Nanini-Maury, Joeri Van Mierlo, and Harry E. Hoster. Data-driven health estimation and lifetime prediction of lithium-ion batteries: A review. *Renewable and Sustainable Energy Reviews*, 113:109254, October 2019.

[11] Juan Rivera-Barrera, Nicolás Muñoz-Galeano, and Henry Sarmiento-Maldonado. SoC estimation for lithium-ion batteries: Review and future challenges. *Electronics*, 6(4):102, November 2017.

[12] Wenlong Xie, Xinhua Liu, Rong He, Yalun Li, Xinlei Gao, Xinghu Li, Zhaoxia Peng, Suwei Feng, Xuning Feng, and Shichun Yang. Challenges and opportunities toward fast-charging of lithium-ion batteries. *Journal of Energy Storage*, 32:101837, December 2020.

[13] Jacqueline S. Edge, Simon O'Kane, Ryan Prosser, Niall D. Kirkaldy, Anisha N. Patel, Alastair Hales, Abir Ghosh, Weilong Ai, Jingyi Chen, Jiang Yang, Shen Li, Mei-Chin Pang, Laura Bravo Diaz, Anna Tomaszewska, M. Waseem Marzook, Karthik N. Radhakrishnan, Huizhi Wang, Yatish Patel, Billy Wu, and Gregory J. Offer. Lithium ion battery degradation: What you need to know. *Physical Chemistry Chemical Physics*, 23(14):8200–8221, 2021.

[14] K. B. Hatzell, A. Sharma, and H. K. Fathy. A survey of long-term health modeling, estimation, and control of lithium-ion batteries: Challenges and opportunities. In *American Control Conference*, pages 584–591. IEEE, June 2012.

[15] Jacqueline S. Edge, Simon O'Kane, Ryan Prosser, Niall D. Kirkaldy, Anisha N. Patel, Alastair Hales, Abir Ghosh, Weilong Ai, Jingyi Chen, Jiang Yang, Shen Li, Mei-Chin Pang, Laura Bravo Diaz, Anna Tomaszewska, M. Waseem Marzook, Karthik N. Radhakrishnan, Huizhi Wang, Yatish Patel, Billy Wu, and Gregory J. Offer. Lithium ion battery degradation: what you need to know. *Physical Chemistry Chemical Physics*, 23(14):8200–8221, 2021.

[16] Andrew Weng, Peyman Mohtat, Peter M. Attia, Valentin Sulzer, Suhak Lee, Greg Less, and Anna Stefanopoulou. Predicting the impact of formation protocols on battery lifetime immediately after manufacturing. *Joule*, 5(11):2971–2992, November 2021.

[17] Stavros X. Drakopoulos, Azarmidokht Gholamipour-Shirazi, Paul MacDonald, Robert C. Parini, Carl D. Reynolds, David L. Burnett, Ben Pye, Kieran B. O'Regan, Guanmei Wang, Thomas M. Whitehead, Gareth J. Conduit, Alexandru Cazacu, and Emma Kendrick. Formulation and manufacturing optimization of lithium-ion graphite-based electrodes via machine learning. *Cell Reports Physical Science*, 2(12):100683, December 2021.

[18] A. Eldesoky, M. Bauer, T. Bond, Nicholas Kowalski, J. Corsten, D. Rathore, R. Dressler, and J. R. Dahn. Long-term study on the impact of depth of discharge, c-rate, voltage, and temperature on the lifetime of single-crystal NMC811/artificial graphite pouch cells. *Journal of The Electrochemical*

Society, 169(10):100531, October 2022.

[19] Jun Lu, Tianpin Wu, and Khalil Amine. State-of-the-art characterization techniques for advanced lithium-ion batteries. *Nature Energy*, 2(3):17011, March 2017.

[20] Partha P. Paul, Eric J. McShane, Andrew M. Colclasure, Nitash Balsara, David E. Brown, Chuntian Cao, Bor-Rong Chen, Parameswara R. Chinnam, Yi Cui, Eric J. Dufek, Donal P. Finegan, Samuel Gillard, Wenxiao Huang, Zachary M. Konz, Robert Kostecki, Fang Liu, Sean Lubner, Ravi Prasher, Molleigh B. Preefer, Ji Qian, Marco-Tulio Fonseca Rodrigues, Manuel Schnabel, Seoung-Bum Son, Venkat Srinivasan, Hans-Georg Steinrück, Tanvir R. Tanim, Michael F. Toney, Wei Tong, Francois Usseglio-Viretta, Jiayu Wan, Maha Yusuf, Bryan D. McCloskey, and Johanna Nelson Weker. A review of existing and emerging methods for lithium detection and characterization in Li-ion and Li-metal batteries. *Advanced Energy Materials*, 11(17):2100372, March 2021.

[21] Daniel Juarez-Robles, Judith A. Jeevarajan, and Partha P. Mukherjee. Degradation-safety analytics in lithium-ion cells: Part i. aging under charge/discharge cycling. *Journal of The Electrochemical Society*, 167(16):160510, November 2020.

[22] Daniel Juarez-Robles, Saad Azam, Judith A. Jeevarajan, and Partha P. Mukherjee. Degradation-safety analytics in lithium-ion cells and modules: Part III. aging and safety of pouch format cells. *Journal of The Electrochemical Society*, 168(11):110501, November 2021.

[23] Christoph R. Birkl, Matthew R. Roberts, Euan McTurk, Peter G. Bruce, and David A. Howey. Degradation diagnostics for lithium ion cells. *Journal of Power Sources*, 341:373–386, February 2017.

[24] Jiale Mao, Jiazhi Miao, Yingying Lu, and Zheming Tong. Machine learning of materials design and state prediction for lithium ion batteries. *Chinese Journal of Chemical Engineering*, 37:1–11, September 2021.

[25] Xing Shu, Shiquan Shen, Jiangwei Shen, Yuanjian Zhang, Guang Li, Zheng Chen, and Yonggang Liu. State of health prediction of lithium-ion batteries based on machine learning: Advances and perspectives. *iScience*, 24(11):103265, November 2021.

[26] Siyu Jin, Xin Sui, Xinrong Huang, Shunli Wang, Remus Teodorescu, and Daniel-Ioan Stroe. Overview of machine learning methods for lithium-ion battery remaining useful lifetime prediction. *Electronics*, 10(24):3126, December 2021.

[27] Maitane Berecibar. Machine-learning techniques used to accurately predict battery life. *Nature*, 568(7752):325–326, April 2019.

[28] Xin Sui, Shan He, Søren B. Vilsen, Jinhao Meng, Remus Teodorescu, and Daniel-Ioan Stroe. A review of non-probabilistic machine learning-based state of health estimation techniques for lithium-ion battery. *Applied Energy*, 300:117346, October 2021.

[29] Sangwook Kim, Zonggen Yi, Bor-Rong Chen, Tanvir R. Tanim, and Eric J. Dufek. Rapid failure mode classification and quantification in batteries: A deep learning modeling framework. *Energy Storage Materials*, 45:1002–1011, March 2022.

[30] Weihan Li, Neil Sengupta, Philipp Dechent, David Howey, Anuradha Annaswamy, and Dirk Uwe Sauer. Online capacity estimation of lithium-ion batteries with deep long short-term memory networks. *Journal of Power Sources*, 482:228863, January 2021.

[31] Darius Roman, Saurabh Saxena, Valentin Robu, Michael Pecht, and David Flynn. Machine learning pipeline for battery state-of-health estimation. *Nature Machine Intelligence*, 3(5):447–456, April 2021.

[32] Teo Lombardo, Marc Duquesnoy, Hassna El-Bouysidy, Fabian Årén, Alfonso Gallo-Bueno, Peter Bjørn Jørgensen, Arghya Bhowmik, Arnaud Demortière, Elixabete Ayerbe, Francisco Alcaide, Marine Reynaud, Javier Carrasco, Alexis Grimaud, Chao Zhang, Tejs Vegge, Patrik Johansson, and Alejandro A. Franco. Artificial intelligence applied to battery research: Hype or reality? *Chemical Reviews*, 122(12):10899–10969, September 2021.

[33] Paul Gasper, Nils Collath, Holger C. Hesse, Andreas Jossen, and Kandler Smith. Machine-learning assisted identification of accurate battery lifetime models with uncertainty. *Journal of The Electrochemical Society*, 169(8):080518, August 2022.

[34] Noah H. Paulson, Joseph Kubal, Logan Ward, Saurabh Saxena, Wenquan Lu, and Susan J. Babinec. Feature engineering for machine learning enabled early prediction of battery lifetime. *Journal of Power Sources*, 527:231127, April 2022.

[35] Robert R. Richardson, Michael A. Osborne, and David A. Howey. Gaussian process regression for forecasting battery state of health. *Journal of Power Sources*, 357:209–219, July 2017.

[36] Benben Jiang, William E. Gent, Fabian Mohr, Supratim Das, Marc D. Berliner, Michael Forsuelo, Hongbo Zhao, Peter M. Attia, Aditya Grover, Patrick K. Herring, Martin Z. Bazant, Stephen J. Harris, Stefano Ermon, William C. Chueh, and Richard D. Braatz. Bayesian learning for rapid prediction of lithium-ion battery-cycling protocols. *Joule*, 5(12):3187–3203, December 2021.

[37] Weihan Li, Haotian Zhang, Bruis van Vlijmen, Philipp Dechent, and Dirk Uwe Sauer. Forecasting battery capacity and power degradation with multi-task learning. *Energy Storage Materials*, 53:453–466, December 2022.

[38] Penelope K. Jones, Ulrich Stimming, and Alpha A. Lee. Impedance-based forecasting of lithium-ion battery performance amid uneven usage. *Nature Communications*, 13(1), August 2022.

[39] Jiangong Zhu, Yixiu Wang, Yuan Huang, R. Bhushan Gopaluni, Yankai Cao, Michael Heere, Martin J. Mühlbauer, Liuda Mereacre, Haifeng Dai, Xinhua Liu, Anatoliy Senyshyn, Xuezhe Wei, Michael Knapp, and Helmut Ehrenberg. Data-driven capacity estimation of commercial lithium-ion batteries from voltage relaxation. *Nature Communications*, 13(1), April 2022.

[40] Thomas F. Fuller, Marc Doyle, and John Newman. Simulation and optimization of the dual lithium ion insertion cell. *Journal of The Electrochemical Society*, 141(1):1–10, January 1994.

[41] Selcuk Atalay, Muhammad Sheikh, Alessandro Mariani, Yu Merla, Ed Bower, and W. Dhammika Widanage. Theory of battery ageing in a lithium-ion battery: Capacity fade, nonlinear ageing and lifetime prediction. *Journal of Power Sources*, 478:229026, December 2020.

[42] Marc Doyle, Thomas F. Fuller, and John Newman. Modeling of galvanostatic charge and discharge of the lithium/polymer/insertion cell. *Journal of The Electrochemical Society*, 140(6):1526–1533, June 1993.

[43] Marc D. Berliner, Hongbo Zhao, Supratim Das, Michael Forsuelo, Benben Jiang, William H. Chueh, Martin Z. Bazant, and Richard D. Braatz. Nonlinear identifiability analysis of the porous electrode theory model of lithium-ion batteries. *Journal of The Electrochemical Society*, 168(9):090546, September 2021.

[44] Dong Zhang, Luis D. Couto, and Scott J. Moura. Electrode-level state estimation in lithium-ion batteries via kalman decomposition. *IEEE Control Systems Letters*, 5(5):1657–1662, November 2021.

[45] Simon E. J. O'Kane, Weilong Ai, Ganesh Madabattula, Diego Alonso Alvarez, Robert Timms, Valentin Sulzer, Jacqueline Sophie Edge, Billy Wu, Gregory J. Offer, and Monica Marinescu. Lithium-ion battery degradation: How to model it. *Phys. Chem. Chem. Phys.*, 24:7909–7922, 2022.

[46] Bor-Rong Chen, Cody M. Walker, Sangwook Kim, M. Ross Kunz, Tanvir R. Tanim, and Eric J. Dufek. Battery aging mode identification across NMC compositions and designs using machine learning. *Joule*, 6(12):2776–2793, December 2022.

[47] Chade Lv, Xin Zhou, Lixiang Zhong, Chunshuang Yan, Madhavi Srinivasan, Zhi Wei Seh, Chuntai Liu, Hongge Pan, Shuzhou Li, Yonggang Wen, and Qingyu Yan. Machine learning: An advanced platform for materials development and state prediction in lithium-ion batteries. *Advanced Materials*, 34(25):2101474, September 2021.

[48] Matthieu Dubarry, Cyril Truchot, and Bor Yann Liaw. Synthesize battery degradation modes via a diagnostic and prognostic model. *Journal of Power Sources*, 219:204–216, December 2012.

[49] Matthieu Dubarry and David Beck. Perspective on mechanistic modeling of Li-ion batteries. *Accounts of Materials Research*, 3(8):843–853, June 2022.

[50] Matthieu Dubarry and David Beck. Big data training data for artificial intelligence-based Li-ion diagnosis and prognosis. *Journal of Power Sources*, 479:228806, December 2020.

[51] Stefan Schindler, George Baure, Michael A. Danzer, and Matthieu Dubarry. Kinetics accommodation in Li-ion mechanistic modeling. *Journal of Power Sources*, 440:227117, November 2019.

[52] Matthieu Dubarry, M. Berecibar, A. Devie, D. Anseán, N. Omar, and I. Villarreal. State of health battery estimator enabling degradation diagnosis: Model and algorithm description. *Journal of Power Sources*, 360:59–69, August 2017.

[53] Gonçalo dos Reis, Calum Strange, Mohit Yadav, and Shawn Li. Lithium-ion battery data and where to find it. *Energy and AI*, 5:100081, September 2021.

[54] Wei He, Nicholas Williard, Michael Osterman, and Michael Pecht. Prognostics of lithium-ion batteries based on Dempster–Shafer theory and the Bayesian Monte Carlo method. *Journal of Power Sources*, 196(23):10314–10321, December 2011.

[55] Yinjiao Xing, Eden W.M. Ma, Kwok-Leung Tsui, and Michael Pecht. An ensemble model for predicting the remaining useful performance of lithium-ion batteries. *Microelectronics Reliability*, 53(6):811–820, June 2013.

[56] Shuzhi Zhang, Xu Guo, Xiaoxin Dou, and Xiongwen Zhang. A data-driven coulomb counting method for state of charge calibration and estimation of lithium-ion battery. *Sustainable Energy Technologies and Assessments*, 40:100752, August 2020.

[57] Damian Burzyński and Leszek Kasprzyk. A novel method for the modeling of the state of health of lithium-ion cells using machine learning for practical applications. *Knowledge-Based Systems*, 219:106900, May 2021.

[58] Kristen A. Severson, Peter M. Attia, Norman Jin, Nicholas Perkins, Benben Jiang, Zi Yang, Michael H. Chen, Muratahan Aykol, Patrick K. Herring, Dimitrios Fraggedakis, Martin Z. Bazant, Stephen J. Harris, William C. Chueh, and Richard D. Braatz. Data-driven prediction of battery cycle life before capacity degradation. *Nature Energy*, 4(5):383–391, March 2019.

[59] Peter M. Attia, Aditya Grover, Norman Jin, Kristen A. Severson, Todor M. Markov, Yang-Hung Liao, Michael H. Chen, Bryan Cheong, Nicholas Perkins, Zi Yang, Patrick K. Herring, Muratahan Aykol, Stephen J. Harris, Richard D. Braatz, Stefano Ermon, and William C. Chueh. Closed-loop optimization of fast-charging protocols for batteries with machine learning. *Nature*, 578(7795):397–402, February 2020.

[60] Weiping Diao, Saurabh Saxena, and Michael Pecht. Accelerated cycle life testing and capacity degradation modeling of $LiCoO_2$-graphite cells. *Journal of Power Sources*, 435:226830, September 2019.

[61] Logan Ward, Joseph Kubal, Susan J. Babinec, Wenquan Lu, Allison Dunlop, Steve Trask, Bryant Polzin, Andrew Jansen, and Noah H. Paulson. Dataset of NMC battery tests from CAMP, 2023 release. Technical report, Materials Data Facility, Argonne National Laboratory, Illinois, 2023.

[62] Leo Wildfeuer, Alexander Karger, Deniz Aygül, Nikolaos Wassiliadis, Andreas Jossen, and Markus Lienkamp. Experimental degradation study of a commercial lithium-ion battery. *Journal of Power Sources*, 560:232498, March 2023.

[63] Muratahan Aykol, Chirranjeevi Balaji Gopal, Abraham Anapolsky, Patrick K. Herring, Bruis van Vlijmen, Marc D. Berliner, Martin Z. Bazant, Richard D. Braatz, William C. Chueh, and Brian D. Storey. Perspective—Combining physics and machine learning to predict battery lifetime. *Journal of The Electrochemical Society*, 168(3):030525, March 2021.

[64] Alan G. Li, Alan C. West, and Matthias Preindl. Towards unified machine learning characterization of lithium-ion battery degradation across multiple levels: A critical review. *Applied Energy*, 316:119030, June 2022.

[65] Donal P. Finegan, Juner Zhu, Xuning Feng, Matt Keyser, Marcus Ulmefors, Wei Li, Martin Z. Bazant, and Samuel J. Cooper. The application of data-driven methods and physics-based learning for improving battery safety. *Joule*, 5(2):316–329, February 2021.

[66] Xiaosong Hu, Le Xu, Xianke Lin, and Michael Pecht. Battery lifetime prognostics. *Joule*, 4(2):310–346, February 2020.

[67] Valentin Sulzer, Peyman Mohtat, Antti Aitio, Suhak Lee, Yen T. Yeh, Frank Steinbacher, Muhammad Umer Khan, Jang Woo Lee, Jason B. Siegel, Anna G. Stefanopoulou, and David A. Howey. The challenge and opportunity of battery lifetime prediction from field data. *Joule*, 5(8):1934–1955, August 2021.

[68] Jiangong Zhu, Yixiu Wang, Yuan Huang, R. Bhushan Gopaluni, Yankai Cao, Michael Heere, Martin J. Mühlbauer, Liuda Mereacre, Haifeng Dai, Xinhua Liu, Anatoliy Senyshyn, Xuezhe Wei, Michael Knapp, and Helmut Ehrenberg. Data-driven capacity estimation of commercial lithium-ion batteries from voltage relaxation. *Nature Communications*, 13(1), April 2022.

[69] Jon P. Christopherson. Battery test manual for electric vehicles. Technical report, Jun 2015.

[70] Thorsten Baumhöfer, Manuel Brühl, Susanne Rothgang, and Dirk Uwe Sauer. Production caused variation in capacity aging trend and correlation to initial cell performance. *Journal of Power Sources*, 247:332–338, February 2014.

[71] Yuliya Preger, Heather M. Barkholtz, Armando Fresquez, Daniel L. Campbell, Benjamin W. Juba, Jessica Romàn-Kustas, Summer R. Ferreira, and Babu Chalamala. Degradation of commercial lithium-ion cells as a function of chemistry and cycling conditions. *Journal of The Electrochemical Society*, 167(12):120532, January 2020.

[72] Trishna Raj, Andrew A. Wang, Charles W. Monroe, and David A. Howey. Investigation of path-dependent degradation in lithium-ion batteries. *Batteries & Supercaps*, 3(12):1377–1385, September 2020.

[73] Scott Lundberg. SHAP: A game theoretic approach to explain the output of any machine learning model. https://github.com/slundberg/shap/.

[74] Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, page 4768–4777. Curran Associates, Inc., 2017.

[75] Jing Li, Jessie Harlow, Nikolai Stakheiko, Ning Zhang, Jens Paulsen, and Jeff Dahn. Dependence of cell failure on cut-off voltage ranges and observation of kinetic hindrance in $LiNi_{0.8}Co_{0.15}Al_{0.05}O_2$. *Journal of The Electrochemical Society*, 165(11):A2682–A2695, 2018.

[76] K. J. Nelson, G. L. d'Eon, A. T. B. Wright, L Ma, J. Xia, and J. R. Dahn. Studies of the effect of high voltage on the impedance and cycling performance of $Li[Ni_{0.4}Mn_{0.4}Co_{0.2}]O_2$/Graphite lithium-ion pouch cells. *Journal of The Electrochemical Society*, 162(6):A1046–A1054, 2015.

[77] Seung-Taek Myung, Filippo Maglia, Kang-Joon Park, Chong Seung Yoon, Peter Lamp, Sung-Jin Kim, and Yang-Kook Sun. Nickel-rich layered cathode materials for automotive lithium-ion batteries: Achievements and perspectives. *ACS Energy Letters*, 2(1):196–223, January 2017.

[78] M. Uitz, M. Sternad, S. Breuer, C. Täubert, T. Traußnig, V. Hennige, I. Hanzu, and M. Wilkening. Aging of Tesla's 18650 lithium-ion cells: Correlating Solid-Electrolyte-Interphase evolution with fading in capacity and power. *Journal of The Electrochemical Society*, 164(14):A3503–A3510, November 2017.

[79] Alexis Geslin, Bruis van Vlijmen, Xiao Cui, Arjun Bhargava, Patrick Asinger, Richard Braatz, and William Chueh. Battery lifetime predictions: information leakage from unblinded training. March 2023.

[80] Patrick Herring, Chirranjeevi Balaji Gopal, Muratahan Aykol, Joseph H. Montoya, Abraham Anapolsky, Peter M. Attia, William Gent, Jens S. Hummelshøj, Linda Hung, Ha-Kyung Kwon, Patrick Moore, Daniel Schweigert, Kristen A. Severson, Santosh Suram, Zi Yang, Richard D. Braatz,

and Brian D. Storey. BEEP: A python library for battery evaluation and early prediction. *SoftwareX*, 11:100506, January 2020.

[81] C. R. Birkl, E. McTurk, M. R. Roberts, P. G. Bruce, and D. A. Howey. A parametric open circuit voltage model for lithium ion batteries. *Journal of The Electrochemical Society*, 162(12):A2271–A2280, 2015.

[82] Julius Schmitt, Markus Schindler, Andreas Oberbauer, and Andreas Jossen. Determination of degradation modes of lithium-ion batteries considering aging-induced changes in the half-cell open-circuit potential curve of silicon–graphite. *Journal of Power Sources*, 532:231296, 2022.

10 Acknowledgments

This work was supported by the Toyota Research Institute through the Accelerated Materials Design and Discovery program. V.L. and D.G. acknowledge support by the National Science Foundation Graduate Research Fellowship. X.C. acknowledges support by the Stanford Data Science Scholars Program. N.R.G. acknowledges support from the U.S. Department of Defense through the National Defense Science & Engineering Graduate Fellowship Program. We thank Alexis Geslin for editing the manuscript. We also thank D'Arcy Seamon for contributions on experiments related to half-cell voltage extraction. Additionally, we thank Alex Dunn for contributions in developing the BEEP python package.

11 Author contributions

V.L., B.V.V., X.C., W.E.G., W.C.C., conceived and conducted the full-cell cycling experiments. V.L. compiled the manuscript, figures and edits from authors. P.A.A. and B.V.V. conceptualized, implemented, and constructed visualizations for the protocol-only, diagnostic-aided, and diagnostic-only machine learning models and their SHAP analysis. P.A.A., B.V.V., V.L., and X.C. conceptualized the mechanistic SOH metrics used. X.C. conceptualized and implemented the explanatory model and its SHAP analysis. P.A.A. conceptualized and implemented the DVF algorithm in strong collaboration with B.V.V.. P.K.H., C.B.G., S.S, A.T. performed data management and data processing pipeline support. N.G., V.L., and H.L.T. designed the methodology for full cell disassembly, pouch cell assembly, and half cell voltage extraction. N.G. and V.L. conducted the pouch cell experiments. D.G., V.L., B.V.V., X.C. and P.A.A. contributed to conceptualizing and implementing the visual representations of the work. All authors edited, reviewed and discussed the work. W.C.C. and R.D.B. supervised the work.

12 Competing interests

W.C.C., B.V.V., W.E.G., V.L., P.K.H., C.B.G., P.A.A., R.D.B., X.C., have filed a patent related to this work: US Application No. 20220137149A1, dated 15 May 2022.

13 Additional Information

Supplementary information will be available for this paper at publication. **Correspondence and requests for materials** should be addressed to W.C.C. or R.D.B.