

Early-Cycle Internal Impedance Enables ML-Based Battery Cycle Life Predictions Across Manufacturers

Tyler Sours^{*a}, Shivang Agarwal^{*a}, Marc Cormier^{*b}, Jordan Crivelli-Decker^a, Steffen Ridderbusch^a, Stephen L. Glazier^b, Aayush R. Singh^a, Ang Xiao^{†a}, Omar Allam^{†a}

^a*SandboxAQ, Palo Alto, California 94301, United States*

^b*NOVONIX Battery Technology Solutions, Bedford, Nova Scotia*

Abstract

Predicting the end-of-life (EOL) of lithium-ion batteries across different manufacturers presents significant challenges due to variations in electrode materials, manufacturing processes, cell formats, and a lack of generally available data. Methods that construct features solely on voltage-capacity profile data typically fail to generalize across cell chemistries. This study introduces a methodology that combines traditional voltage-capacity features with Direct Current Internal Resistance (DCIR) measurements, enabling more accurate and generalizable EOL predictions. The use of early-cycle DCIR data captures critical degradation mechanisms related to internal resistance growth, enhancing model robustness. Models are shown to successfully predict the number of cycles to EOL for unseen manufacturers of varied electrode composition with a mean absolute error (MAE) of 150 cycles. This cross-manufacturer generalizability reduces the need for extensive new data collection and retraining, enabling manufacturers to optimize new battery designs using existing datasets. Additionally, a novel DCIR-compatible dataset is released as part of ongoing efforts to enrich the growing ecosystem of cycling data and accelerate battery materials development.

1. Introduction

Lithium-ion batteries are indispensable across a wide range of applications due to their high energy densities, decreasing costs, and extended lifetimes. However, predicting the end of life (EOL) of these batteries, which is critical to ensuring their reliability and safety, remains a significant challenge due to the complex and nonlinear degradation mechanisms involved. Early and accurate prediction of battery EOL is essential for several reasons. It enables accelerated testing and validation of new battery formulations, allowing manufacturers to optimize chemistries and processes more efficiently.

Severson et al. (2019) made significant contributions by utilizing early-cycle features from standard discharge capacity curves to predict battery life, introducing an approach that provided high accuracy for end-of-life (EOL) predictions. Since then, a wide array of studies has focused on improving EOL predictions Elmahallawy et al. (2022); Ling (2022); Paulson et al. (2022); Saxena et al. (2022); Ma et al. (2020); Fei et al. (2022); Geslin et al. (2023); Saxena et al. (2021); Chen et al. (2022); Li et al. (2019); Rangel-Martinez et al. (2021); Saxena et al. (2015); Zhang et al. (2018); Hosen et al. (2021); Attia et al. (2021); Ng et al. (2020). However, many of these works have relied on capacity-based features and models that perform well only within the specific datasets on which they were trained.

* Authors contributed equally to this work.

† Correspondence to: ang.xiao@sandboxquantum.com, omar.allam@sandboxquantum.com.

When applied to different chemistries or operating conditions, their predictive power diminishes significantly [Severson et al. \(2019\)](#); [Sulzer et al. \(2021\)](#); [Finegan et al. \(2021\)](#); [Attia et al. \(2020\)](#); [Paulson et al. \(2024\)](#); [Finegan and Cooper \(2019\)](#); [Walker et al. \(2019\)](#); [Finegan et al. \(2019\)](#); [Wang et al. \(2024\)](#). [Paulson et al. \(2022\)](#) created a machine learning model that generalized across several cathode chemistries, utilizing a dataset drawn from a substantial internal archive of batteries produced in-house. Recently, [Rahmanian et al. \(2024\)](#) proposed an attention-based recurrent neural network capable of handling diverse chemistries and cycling protocols. While their model demonstrates adaptability, it still requires fine-tuning for new chemistries and is reliant on large proprietary datasets, limiting its practical application in environments with limited data availability. Therefore, predictive models for battery cycle life that generalize across manufacturers without relying on data-hungry models are still lacking. Battery performance is influenced not only by electrode chemistry but also by a variety of factors related to cell design and manufacturing processes [Baumhöfer et al. \(2014\)](#). Even cells with nominally similar chemistries can exhibit vastly different behaviors due to differences in manufacturing parameters, including electrode coating thickness, particle size distribution, slurry preparation, and precision in assembly. Such variations can affect the electrochemical performance, aging mechanisms, and overall reliability of a cell, making cross-manufacturer generalization a significantly more complex challenge than cross-chemistry generalization alone. These subtle changes profoundly impact degradation pathways, leading to different rates and patterns of failure. For instance, as the battery cycles, degradation processes such as the formation and growth of the solid electrolyte interphase (SEI) on the anode surface, lithium plating, and the accumulation of inactive lithium deposits lead to an increase in interfacial resistance. Concurrently, mechanical degradation of electrode materials, loss of active material, and electrolyte decomposition contribute to sluggish ion and electron mobility, resulting in an overall increase in internal impedance [Diao et al. \(2022\)](#); [Fermín-Cueto et al. \(2020\)](#).

Each manufacturer employs distinct cell designs, electrode materials, and processing techniques. These differences manifest in a variety of ways, such as the thickness of the electrodes, the specific anode and cathode formulations, the quality and consistency of the solid-electrolyte interface (SEI) layer, and thermal management strategies during operation. Furthermore, slight deviations in particle size distribution or the homogeneity of active material coatings can lead to differences in impedance growth and capacity fade. Thus, despite sharing broad similarities in lithium-ion technology, the cells in this study display unique degradation profiles over time, emphasizing the complexity of generalizing across manufacturers. Though works in recent years have made notable progress in cross-chemistry generalization, the challenge of cross-manufacturer generalization remains underexplored.

To address this challenge, the present work uses Direct Current Internal Resistance (DCIR) measurements from early cycles as a key feature in predictive modeling, capturing these underlying degradation processes to enable accurate EOL predictions across chemistries and operating conditions. Using the ensemble of DCIR and traditional capacity-based features from early cycles, models are developed that maintain high predictive accuracy even with simple linear regressors, achieving intrinsic generalizability without the need for complex algorithms or extensive retraining. Incorporating DCIR into routine cycling protocols offers a straightforward and cost-effective way to enhance model robustness and generalizability across varying manufacturers.

2. Cycling Protocol and Dataset Generation

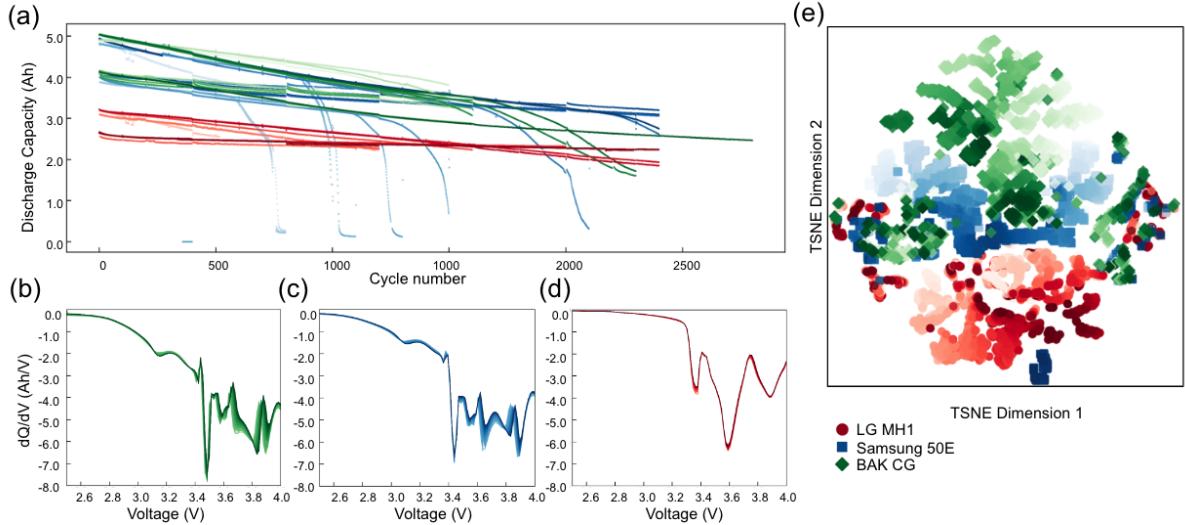


Figure 1: **Distinct cycling behavior across battery manufacturers complicates cross-manufacturer prediction when relying on capacity-based features alone.** (a) Discharge capacity as a function of cycle number for batteries from three manufacturers (LG MH1, Samsung 50E, and BAK CG). Clear differences in degradation rates and capacity loss profiles can be observed across manufacturers. (b-d) Differential capacity (dQ/dV) curves at cycle 100 for (b) BAK CG, (c) Samsung 50E, and (d) LG MH1 cells, highlighting the distinct electrochemical signatures for each manufacturer. (e) t-SNE plot of features derived from cycling data, showing clustering of cells by manufacturer, further illustrating the challenge of cross-manufacturer generalization in battery life prediction.

To address the scarcity of standardized cross-chemistry datasets for battery end-of-life (EOL) prediction, a comprehensive benchmark dataset acquired from 57 commercial lithium-ion cells from 3 different manufacturers with different cathode-anode combinations is introduced (Fig. 1). The cells were cycled with a diagnostic protocol that incorporated periodic DCIR measurements to track the internal resistance of the battery cells at incremental states-of-charge (SOC) to provide a detailed view of cell degradation and performance, a method widely utilized by OEMs and cell makers.

Battery cycling data were recorded at intervals ranging from a few seconds to a few minutes, depending on the cycling current, generating time-series data for voltage, current, and time. From these fundamental quantities, additional metrics such as capacity and energy were computed. A key metric, the discharge capacity, was tracked across cycles to monitor battery health, with the remaining percentage of initial discharge capacity serving as the state-of-health (SOH).

The discharge capacity versus cycle number for all cells is illustrated in Fig. 1a, where cells are colored by manufacturer. While most cells exhibit linear decay in discharge capacity, some demonstrate non-linear profiles, emphasizing the need for models that can accurately capture these variations. The focus of this work is on predicting the cycle number at which a fixed SOH (e.g., 85%) is reached, recognizing that different applications may require different SOH thresholds for a cell to be considered at its "end of life". With this in mind, it should be noted that while "EOL" may not be the most appropriate terminology for this predictive task, it is used throughout the manuscript for consistency with other works in the field.

2.1. Battery cell manufacturers and chemistries

The dataset presented in this work is comprised of 57 Li-ion battery cells from 3 different manufacturers, each with a distinct cathode-anode combination. Table 1 describes the difference between each cell type. The cathode in all cells are considered high-Ni; the Samsung cells contain a Nickel-Cobalt-Aluminum (NCA) cathode [Popp et al. \(2020\)](#); [Schmitt et al. \(2023\)](#) while the BAK and LG cells contain a Nickel-Manganese-Cobalt (NMC) cathode. The anode in all cells is graphite-based; the Samsung and BAK cells contain a Graphite-Silicon (Gr-Si) blend anode, while the LG cells contain a pure Graphite (Gr) anode.

Cell	Cathode	Anode	Format	Nominal capacity
Samsung INR21700-50E	NCA	Gr-Si	21700 cylinder	5 Ah
BAK N21700CG	NMC	Gr-Si	21700 cylinder	5 Ah
LG INR18650-MH1	NMC	Gr	18650 cylinder	3 Ah

Table 1: Manufacturer differences between the three cell types used in this study.

These cells differ in their electrode compositions and form factors, which directly influence their degradation behaviors. Fig.1a shows that the discharge capacity versus cycle number curves have varying characteristics depending on the manufacturer and cycling conditions; linear, sub-linear, and super-linear decay profiles are observed [Attia et al. \(2022\)](#). Fig.1b-d demonstrates the differences in electrochemistry between the three cell types considered through the evolution of the differential capacity profiles over 50 cycles (data collected on a high-resolution NOVONIX UHPC system). To visualize the differences in aging characteristics between each cell type, the differential capacities are used in a t-distributed stochastic neighbor embedding, shown in Fig.1e.

By incorporating this diversity, the presented dataset is designed to validate the generalizability of machine learning models across varying commercial cell chemistries. This approach ensures that findings presented in this work are not limited to a single cell type but are applicable across a broad spectrum of commercially relevant lithium-ion cells.

2.2. Operating conditions

Table 2 summarizes the various cycling conditions comprising this dataset. The cells were tested across three voltage ranges and three temperatures to capture a broad spectrum of degradation behaviors. Notably, only the Sam50E and BakCG cells were cycled within a 3.0 V - 4.2 V range to avoid Li-Si alloying reactions, which are specific to Gr-Si anodes. Cycling at 60°C was conducted exclusively with the 2.5 V - 4.2 V range because at high-temperatures degradation is expected to be dominated by time-dependent reactions rather than cycle-dependant reactions.

Each set of operational parameters (temperature, voltage range, and manufacturer) included three identical cells to account for cell-to-cell variability and the influence of different cycling conditions. This comprehensive approach supports the development of predictive models that are robust across a wide range of real-world scenarios.

All cells were cycled at a constant-current rate of C/3, with a constant-voltage hold at the top-of-charge. To monitor and assess cell health, a diagnostic sequence was employed every 100 cycles, starting at the 2nd cycle, that included two low-rate cycles at C/10 followed by a 10-step DCIR measurement. The DCIR measurement was performed using 10 current pulses at equally-spaced SOCs, decrementing in 10% intervals from 100% to 10%.

	V	2.5 - 4.06	2.5 - 4.2	3.0 - 4.2
T				
25	Sam50E, BakCG, LGMH1	Sam50E, BakCG, LGMH1	Sam50E, BakCG	
40	Sam50E, BakCG, LGMH1	Sam50E, BakCG, LGMH1	Sam50E, BakCG	
60	-	Sam50E, BakCG, LGMH1	-	

Table 2: Battery cell cycling conditions matrix, where T refers to the temperature (in Celsius) and V refers to the voltage range between which the cells were cycled. Three voltage ranges and three temperatures were employed to cover various degradation modes across three different cell manufacturers.

Each pulse had a 10-second duration at a current equivalent to a 1C full discharge and was followed by a 15 minute relaxation period, a design that minimizes disruption to the cell’s normal cycling and ensures accurate resistance measurements without compromising cell performance.

This protocol efficiently captures dynamic performance, internal resistance, and overall battery health under realistic operating conditions, providing a rich dataset for developing predictive models that can generalize across various chemistries and conditions.

3. Machine Learning Approach

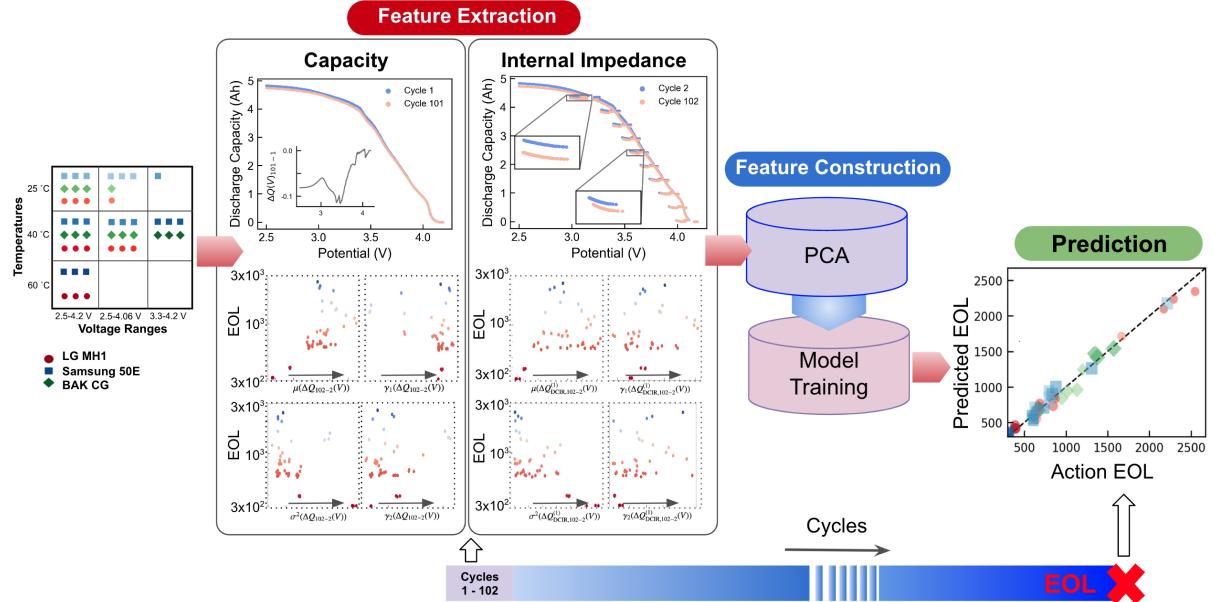


Figure 2: **Schematic representation of feature extraction for battery life prediction.** Starting with raw discharge capacity curves, features such as $\Delta Q(V)$, linear capacity fade, and DCIR-based $\Delta Q(V)$ features are extracted for state-of-health prediction.

All features were extracted from time-series data of the first 102 cycles (Fig.2). Emphasis was given to features derived from electrochemically meaningful data fields such as internal resistance, calculated from Direct Current Internal Resistance (DCIR) pulses:

1. **$\Delta Q(V)$ Features** — Following the work of [Severson et al. \(2019\)](#), univariate time-series features were extracted from the difference between discharge capacity vs. potential curves, $Q(V)$, for two specific cycles (in this case, cycles 102 and 2). The

extracted features include skewness (γ_1), kurtosis (γ_2), variance (σ^2), mean (μ), minimum (min), and maximum (max) of the function:

$$\Delta Q_{102-2}(V) = Q_{102}(V) - Q_2(V)$$

where $Q_n(V)$ represents the $Q(V)$ curve at cycle n . These statistical features capture the distributional and shape changes in $Q(V)$ curves, which provide insights into the battery's state of health over time. Fig. 2 depicts sample $Q_2(V)$ and $Q_{102}(V)$ curves, with the inset showing the corresponding $\Delta Q_{102-2}(V)$.

2. **R_{DCIR} Features** — Each DCIR measurement involves applying a current pulse at a specific state of charge (SOC) during a discharge step. The response to a current pulse is a voltage drop that allows for the calculation of an effective internal resistance: $R = V/I$. To build features for machine learning models, resistance values were extracted for each of the 10 current pulses during cycles 2 and 102 (Fig. 2).
3. **$\Delta Q_{\text{DCIR}}(V)$ Features** — Recognizing that each current pulse and its associated voltage drop response from DCIR measurements can be viewed as forming a partial $Q(V)$ curve, the methodology from [Severson et al. \(2019\)](#) is extended to calculate the same set of univariate time-series statistical transformations (mean, minimum, maximum, variance, skewness, and kurtosis) on these partial curves across the 10 SOC levels, between cycles 2 and 102. This process generated 60 additional features, which is referred to as $\Delta Q_{\text{DCIR}}^{(k)}(V)$ features, where k represents the pulse number. This approach enables us to capture the nuances in the capacity changes at various SOC levels, providing a more detailed understanding of the degradation process. For instance, the kurtosis of the changes in these partial capacity curves ($\gamma_2(\Delta Q_{\text{DCIR},102-2}^{(k)}(V))$) can reveal sharpness or outlier behavior in capacity changes, which might correlate with specific degradation mechanisms.
4. **Linear Fit Features** — To capture the overall trend in discharge capacity over time, a linear fit is applied to the discharge capacity vs. cycle number curve between two early cycles (e.g., cycles 30 and 99). The slope (p_1) and intercept (p_2) of this linear fit were used as features. These features provide insights into the rate of capacity fade and the initial state of the battery, which are crucial for predicting future performance. Although more complex fits were explored, they did not significantly improve model performance and were therefore not included in the final feature set.

These features were then combined and used to train various machine learning models. For regression-based modeling, the best performance was achieved using a simple linear elastic net model. While more complex models such as Random Forest and Multi-layer Perceptron were tested (as detailed in the Supplementary Information), Elastic Net was ultimately selected due to its regularization properties, which help mitigate overfitting—especially when using a relatively small dataset or when the features exhibit multicollinearity. This is particularly important given the nature of our feature sets, which include correlated variables from both traditional capacity-based features and those derived from DCIR measurements. Additionally, as a preprocessing step, Sequential Feature Selection (SFS) was applied for high dimensional feature sets (e.g., the DCIR features) to refine the input feature space for subsequent Principal Component Analysis (PCA). This method demonstrates the robustness of the selected features which achieve high prediction accuracy without the need for complex or resource-intensive models. Although the

models trained in this work show strong predictive accuracy, it is noted that a larger and more diverse dataset is essential to further validate these findings and ensure broader applicability. While cross-validation provides a robust estimate of model performance, the addition of an independent holdout dataset would offer a more definitive evaluation of the model’s capacity to generalize across different chemistries, manufacturers, and operating conditions.

Nevertheless, to evaluate the generalization performance of the trained models, two distinct cross-validation strategies were employed, each focusing on different aspects of the dataset variability. A *leave one triplicate out* cross-validation strategy was adopted to assess the model’s ability to generalize across different operating conditions. The dataset is composed of cells grouped into triplicates, where each triplicate is a set of three cells cycled under identical conditions (i.e., a specific combination of manufacturer, temperature, and voltage range). In this approach, the model was trained on all but one triplicate, and the remaining triplicate was used for testing. This process was repeated until all triplicates were used as test sets, thus evaluating the model across a diverse range of operating conditions. By excluding entire triplicates during each fold, this method tests the model’s ability to predict cell behavior under previously unseen conditions, effectively simulating the real-world scenario where a battery management system must generalize to new operational environments. This strategy maximizes the utility of the dataset and provides a robust assessment of the model’s generalization capabilities across temperature, voltage, and manufacturer combinations.

In addition to operating conditions, cross-manufacturer validation was performed to assess the model’s ability to generalize across different cell manufacturers. This was achieved by training the model on cells from two manufacturers and testing it on cells from the third. This approach evaluates how well the model, trained on cells from specific manufacturers, can predict the performance of cells from a different manufacturer. Since manufacturing processes introduce variability in electrode design, assembly, and material quality, this validation is critical for ensuring the model’s robustness across different production sources.

For completeness, a standard randomized 5-fold cross-validation was also performed, and the results are presented in the Supplementary Information (SI). However, the results presented in this manuscript focus on the validation strategies discussed above which are tailored to the specific challenges of generalization across operating conditions and manufacturers.

4. Results and Discussion

4.1. Cross voltage and temperature performance

To establish a performance baseline for the model across different operational conditions, an established and widely applied feature set is first examined: the original six-feature $\Delta Q(V)$ set from [Severson et al. \(2019\)](#). This feature set has been extensively used in prior work for cycle life prediction and serves as a benchmark for comparison with other feature sets. Fig. 3a shows the parity plot comparing predicted and actual EOL_{85} values using the *leave one triplicate out* cross-validation strategy, which was discussed in the previous section. In all regression tasks, error is evaluated using the mean MAE of all cross-validation (CV) validation set predictions.

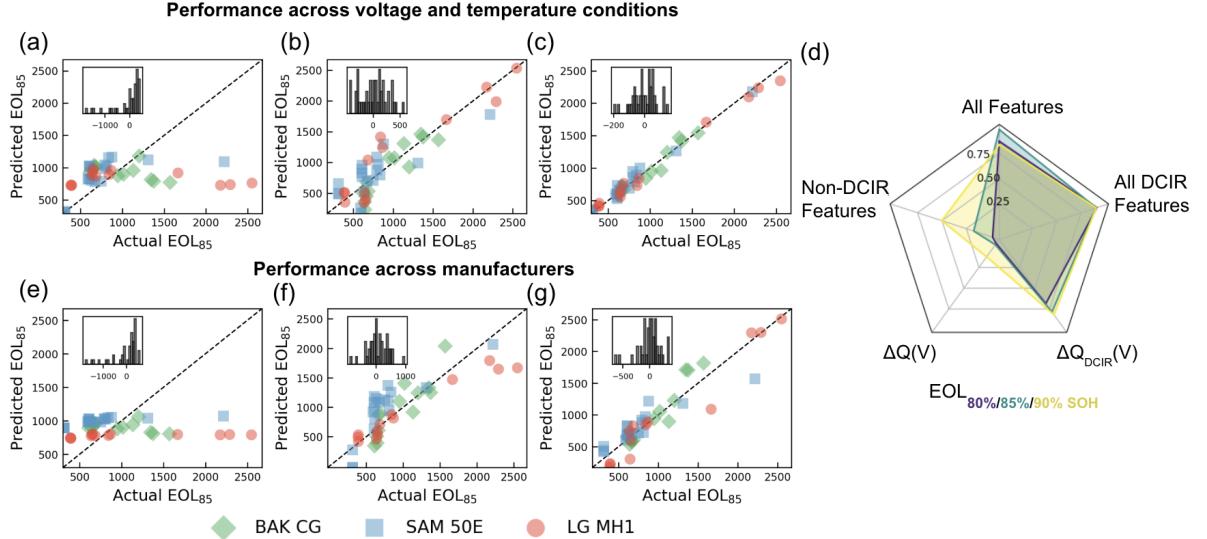


Figure 3: **Predicted vs. actual cycles to EOL for different feature sets from an elastic net regressor across varying conditions.** Panels (a) and (e) depict predictions using the feature set from [Severson et al. \(2019\)](#), while panels (b) and (f) show predictions using a downselected set of 20 features derived from novel DCIR-based $\Delta Q(V)$ data, comparing cycles 102 and 2. Panels (c) and (g) demonstrate results from step-wise backward sequential feature selection, reducing all features to 20 for validation across different testing conditions (top row) and to 10 features for validation across different manufacturers (bottom row), respectively. (d) R^2 scores for different feature sets, corresponding to the predicted number of cycles to reach EOL₈₀, EOL₈₅, and EOL₉₀.

This baseline model is directly compared with an equivalent model trained on an augmented feature set, which includes univariate time-series features computed from DCIR pulses. The expanded feature set consists of 60 attributes, significantly increasing the dimensionality and predictive power of the model. The inclusion of augmented DCIR $\Delta Q(V)$ features captures intricate degradation patterns that are not discernible using the standard feature set alone. As demonstrated in Fig. 3b this approach results in more accurate predictions, with a notable improvement in MAE of 161 cycles. Further, using an amalgamation of the downselected features from the entire feature set, including resistance values from the 10 DCIR pulses, conventional $\Delta Q(V)$ univariate features, and linear fit parameters (p_1, p_2) further enhances the EOL prediction (Fig. 3c). Predictably, models exhibit lower errors at higher SOH thresholds, reflecting the reduced complexity of predicting less degraded states, as presented in Table 3.

The errors and scores reported in Table 3 and Fig. 3a-d are based on the custom CV strategy outlined in the previous section. For reference, using a more conventional 5-fold CV strategy, as detailed in the SI, results in even better performance. However, given the uniqueness of the dataset, which includes cell triplicates, the custom CV strategy is argued to better showcase the model's generalizability across testing conditions.

4.2. Cross manufacturer performance

Traditional cycle life prediction using $\Delta Q(V)_{101-1}$ features, though simple and widely used, fail to generalize across manufacturers, as shown in Fig. 3e with an MAE of 416 cycles. These features capture broad degradation trends but overlook critical indicators like internal resistance growth, which results from processes such as SEI formation and active mass loss due to electrode cracking and structural transformations. Additionally, variations in electrode design and manufacturing processes further limit the predictive power of $\Delta Q(V)$ alone.

SOH (%)	$\Delta Q(V)$	non-DCIR	$\Delta Q_{DCIR}(V)$	All features
80	424.8	425.9	294.3	179.6
85	390.6	319.8	160.6	113.4
90	197.1	142.1	107.4	72.6

Table 3: MAE for different feature combinations for EOL defined at different SOH thresholds. Features based on DCIR measurements provide significant improvements in prediction accuracy compared to non-DCIR features.

The incorporation of early-cycle DCIR-based features, as illustrated in Fig. 3f with an MAE 265 cycles, substantially improves the model’s generalization across manufacturers. DCIR tracks the evolution of internal resistance across various states of charge. This provides a more comprehensive assessment of the cell’s internal electrochemical processes, which enables the model to capture nuanced changes in degradation behavior.

Furthermore, the combination of both DCIR and traditional $\Delta Q(V)$ features, as shown in Fig. 3g with an MAE of 141 cycles, results in a marked improvement in prediction accuracy. By leveraging both the broad degradation trends captured by $\Delta Q(V)$ and the detailed internal resistance information provided by DCIR, the model can more accurately predict state-of-health across a wide range of cell designs and manufacturing processes. This hybrid feature set is particularly effective for bridging the gap between cells with different electrode materials, as evidenced by the model’s ability to predict the EOL of LG graphite-only cells based solely on training data from silicon-graphite cells produced by Samsung and BAK, as well as the accurate prediction of Samsung’s NCA-based cells based the NMC cells by LG and BAK.

5. Electrochemical Indicators and Their Role in Generalization

Variations in manufacturing practices, such as electrode composition, coating uniformity, and assembly precision, introduce significant complexity into cycle life prediction. These subtle differences manifest as changes in internal resistance, capacity fade, and overall cell degradation. Internal resistance, in particular, serves as a critical indicator of underlying degradation mechanisms, including solid electrolyte interphase (SEI) growth, particle fracture, and lithium plating. These degradation processes influence the electrochemical behavior of the cell, which in turn impacts its performance and lifetime.

As shown in Fig. 4, resistance-based features (R_{102} , R_2) and other impedance-related indicators play a critical role in predicting cycle life, contributing substantially to model performance under various operating conditions (Fig. 4a) and across different manufacturers (Fig. 4b). These features, which track resistance behavior at different states of charge, likely capture degradation mechanisms, such as solid electrolyte interphase (SEI) growth and lithium plating, that can manifest across different chemistries and manufacturing processes, though the specific behavior of these mechanisms may still be influenced by these factors. As illustrated in Fig. 4c, cells with greater number of cycles at EOL_{85} , as indicated by the darker colors, exhibit smaller $\Delta Q_{DCIR,102-2}^{(0)}$. This suggests that impedance-based features can effectively differentiate degradation behaviors even when the underlying cell chemistries and manufacturing processes vary significantly.

In contrast, capacity-based features primarily track shifts in capacity fade but may fail to capture early-stage degradation mechanisms such as internal resistance growth, which can occur well before significant capacity loss is observed. This limitation becomes

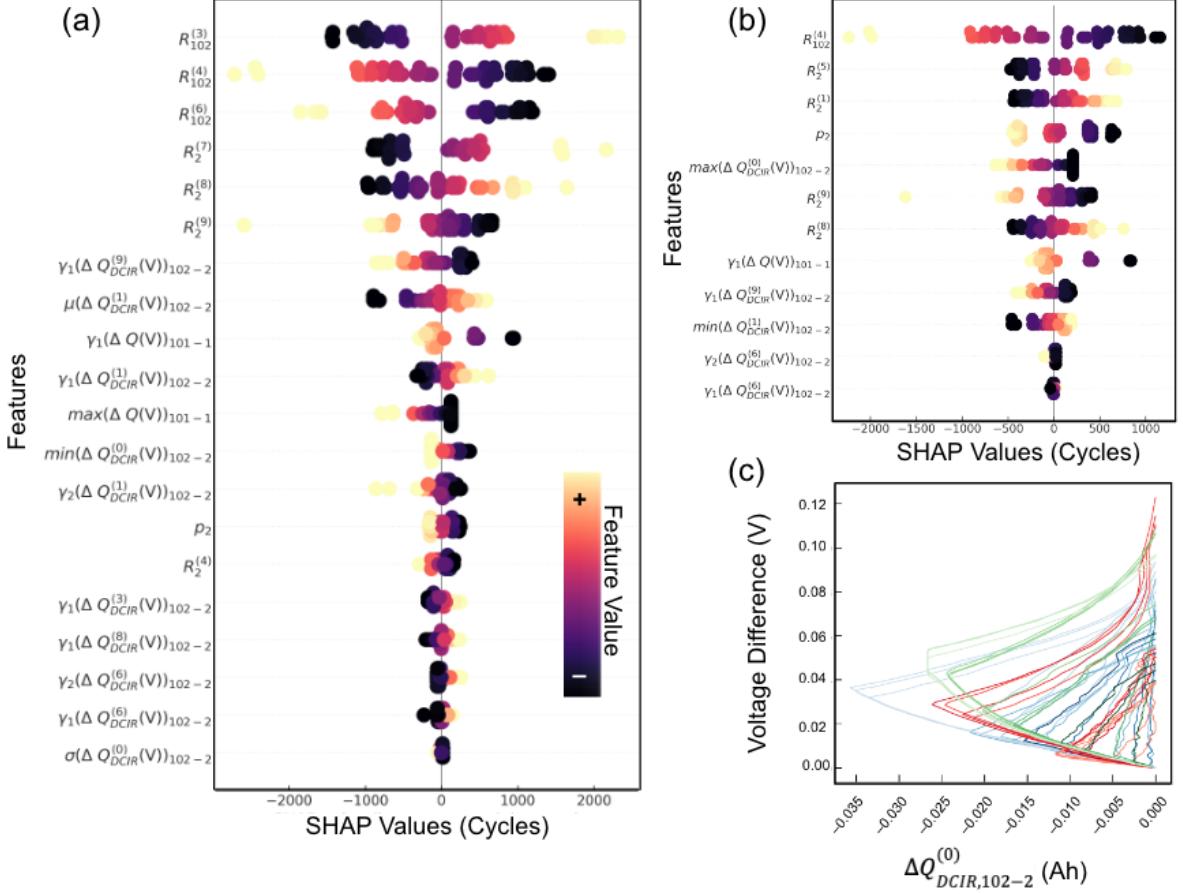


Figure 4: **Feature-based drivers of model performance.** SHAP values (in cycles) demonstrating the contribution of various features to the prediction of EOL₈₅ across (a) operating conditions and (b) manufacturers. Positive SHAP values indicate a positive contribution to longer cycle life, while negative values imply a reduction in predicted cycle life. The color gradient corresponds to the magnitude of feature values, where darker colors indicate higher values. (c) Voltage difference plotted against $\Delta Q_{DCIR,102-2}^{(0)} (Ah)$ for cells from three different manufacturers (LG - red, Samsung - blue, BAK - green). Darker colors correspond to a greater number of cycles to EOL.

particularly evident when attempting to generalize models across different manufacturers, as subtle variations in manufacturing processes (e.g., electrode coating uniformity or electrolyte formulation) can result in different onset points and rates of degradation. The ability of impedance-based features to generalize across these variations highlights their importance in achieving reliable lifetime predictions.

Overall, the ability of DCIR-based features to reflect internal resistance growth, irrespective of specific manufacturing differences, contributes significantly to the model's robustness in generalizing across a wide range of cell designs. These features effectively capture the microscopic and macroscopic degradation processes, ensuring a more reliable prediction of battery life despite variations in cell chemistry or manufacturing techniques.

6. Conclusion

This work demonstrates the efficacy of incorporating internal-impedance-based features from pulse measurements into data-driven models for the prediction of Li-ion battery state-of-health, addressing limitations associated with traditional capacity-based features.

While these conventional features capture broad degradation trends, they fail to generalize effectively across different manufacturers due to variations in cell design, electrode composition, and manufacturing processes. The inclusion of DCIR data, particularly in early cycles, allows for the detection of internal resistance growth, which is closely tied to key degradation mechanisms. Results presented in this work show that this integrated approach significantly improves prediction accuracy and generalizes well across different manufacturers and chemistries. Moreover, this methodology has the potential to allow manufacturers to adapt their battery compositions and leverage pre-existing data for predictive modeling without the need to build entirely new datasets, facilitating rapid innovation and optimization of new cell designs. Incorporating DCIR-based features into predictive models could thus provide a more robust and efficient framework for EOL estimation, advancing battery management systems and their performance in diverse real-world applications.

7. Data Availability and Future Revisions

The data supporting the findings of this study will be made publicly available in the coming weeks. Additionally, minor edits and further content updates are expected as the manuscript undergoes refinement based on ongoing review and feedback.

Acknowledgements

The authors acknowledge Jacob Lokshin, Brian Wee, Yunyun Wang, Brenda Miao, Dan Zhao and Joshua Douglas for their invaluable contributions and support throughout this investigation.

References

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, et al. Closed-loop optimization of fast-charging protocols for batteries with machine learning. *Nature*, 578(7795):397–402, 2020.

Peter M Attia, Kristen A Severson, and Jeremy D Witmer. Statistical learning for accurate and interpretable battery lifetime prediction. *Journal of The Electrochemical Society*, 168(9):090547, 2021.

Peter M Attia, Alexander Bills, Ferran Brosa Planella, Philipp Dechent, Goncalo Dos Reis, Matthieu Dubarry, Paul Gasper, Richard Gilchrist, Samuel Greenbank, David Howey, et al. “knees” in lithium-ion battery aging trajectories. *Journal of The Electrochemical Society*, 169(6):060517, 2022.

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, 2014.

Daoquan Chen, Weicong Hong, and Xiuze Zhou. Transformer network for remaining useful life prediction of lithium-ion batteries. *Ieee Access*, 10:19621–19628, 2022.

Weiping Diao, Jonghoon Kim, Michael H Azarian, and Michael Pecht. Degradation modes and mechanisms analysis of lithium-ion batteries with knee points. *Electrochimica Acta*, 431:141143, 2022.

Mohamed Elmahallawy, Tarek Elfouly, Ali Alouani, and Ahmed M Massoud. A comprehensive review of lithium-ion batteries modeling, and state of health and remaining useful lifetime prediction. *Ieee Access*, 10:119040–119070, 2022.

Zicheng Fei, Zijun Zhang, Fangfang Yang, Kwok-Leung Tsui, and Lishuai Li. Early-stage lifetime prediction for lithium-ion batteries: A deep learning framework jointly considering machine-learned and handcrafted data features. *Journal of Energy Storage*, 52:104936, 2022.

Paula Fermín-Cueto, Euan McTurk, Michael Allerhand, Encarni Medina-Lopez, Miguel F Anjos, Joel Sylvester, and Goncalo Dos Reis. Identification and machine learning prediction of knee-point and knee-onset in capacity degradation curves of lithium-ion cells. *Energy and AI*, 1:100006, 2020.

Donal P Finegan and Samuel J Cooper. Battery safety: data-driven prediction of failure. *Joule*, 3(11):2599–2601, 2019.

Donal P Finegan, John Darst, William Walker, Qibo Li, Chuanbo Yang, Rhodri Jervis, Thomas MM Heenan, Jennifer Hack, James C Thomas, Alexander Rack, et al. Modelling and experiments to identify high-risk failure scenarios for testing the safety of lithium-ion cells. *Journal of Power Sources*, 417:29–41, 2019.

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, 2021.

Alexis Geslin, Bruis Van Vlijmen, Xiao Cui, Arjun Bhargava, Patrick A Asinger, Richard D Braatz, and William C Chueh. Selecting the appropriate features in battery lifetime predictions. *Joule*, 7(9):1956–1965, 2023.

Md Sazzad Hosen, Rekabra Youssef, Theodoros Kalogiannis, Joeri Van Mierlo, and Mai-tane Berecibar. Battery cycle life study through relaxation and forecasting the lifetime via machine learning. *Journal of Energy Storage*, 40:102726, 2021.

Xiaoyu Li, Lei Zhang, Zhenpo Wang, and Peng Dong. Remaining useful life prediction for lithium-ion batteries based on a hybrid model combining the long short-term memory and elman neural networks. *Journal of Energy Storage*, 21:510–518, 2019.

Chen Ling. A review of the recent progress in battery informatics. *npj Computational Materials*, 8(1):33, 2022.

Yanying Ma, Lifeng Wu, Yong Guan, and Zhen Peng. The capacity estimation and cycle life prediction of lithium-ion batteries using a new broad extreme learning machine approach. *Journal of Power Sources*, 476:228581, 2020.

Man-Fai Ng, Jin Zhao, Qingyu Yan, Gareth J Conduit, and Zhi Wei Seh. Predicting the state of charge and health of batteries using data-driven machine learning. *Nature Machine Intelligence*, 2(3):161–170, 2020.

Noah Paulson, Joseph Kubal, and Susan Babinec. Multivariate prognosis of battery advanced state of health via transformers. *Cell Reports Physical Science*, 5:101928, 2024.

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, 2022.

Hartmut Popp, Ningxin Zhang, Marcus Jahn, Mikel Arrinda, Simon Ritz, Matthias Faber, Dirk Uwe Sauer, Philippe Azais, and Iosu Cendoya. Ante-mortem analysis, electrical, thermal, and ageing testing of state-of-the-art cylindrical lithium-ion cells. *Elektrotech. Informationstechnik*, 137(4):169–176, 2020.

Fuzhan Rahamanian, Robert M Lee, Dominik Linzner, Kathrin Michel, Leon Merker, Balazs B Berkes, Leah Nuss, and Helge Sören Stein. Attention towards chemistry agnostic and explainable battery lifetime prediction. *npj Computational Materials*, 10 (1):100, 2024.

Daniel Rangel-Martinez, KDP Nigam, and Luis A Ricardez-Sandoval. Machine learning on sustainable energy: A review and outlook on renewable energy systems, catalysis, smart grid and energy storage. *Chemical Engineering Research and Design*, 174:414–441, 2021.

Samveg Saxena, Caroline Le Floch, Jason MacDonald, and Scott Moura. Quantifying ev battery end-of-life through analysis of travel needs with vehicle powertrain models. *Journal of Power Sources*, 282:265–276, 2015.

Saurabh Saxena, Logan Ward, Joseph Kubal, Hong-Keun Kim, Wenquan Lu, Susan Babinec, and Noah Paulson. A recurrent neural network model for battery capacity fade curve prediction using early life data. In *Electrochemical Society Meeting Abstracts* 239, number 1, pages 65–65. The Electrochemical Society, Inc., 2021.

Saurabh Saxena, Logan Ward, Joseph Kubal, Wenquan Lu, Susan Babinec, and Noah Paulson. A convolutional neural network model for battery capacity fade curve prediction using early life data. *Journal of Power Sources*, 542:231736, 2022.

Christina Schmitt, Martina Gerle, Dennis Kopljar, and K Andreas Friedrich. Full parameterization study of a high-energy and high-power li-ion cell for physicochemical models. *Journal of The Electrochemical Society*, 170(7):070509, 2023.

Kristen A Severson, Peter M Attia, Norman Jin, Nicholas Perkins, Benben Jiang, Zi Yang, Michael H Chen, Muratahan Aykol, Patrick K Herring, Dimitrios Fraggedakis, et al. Data-driven prediction of battery cycle life before capacity degradation. *Nature Energy*, 4(5):383–391, 2019.

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, et al. The challenge and opportunity of battery lifetime prediction from field data. *Joule*, 5 (8):1934–1955, 2021.

William Q Walker, John J Darst, Donal P Finegan, Gary A Bayles, Kenneth L Johnson, Eric C Darcy, and Steven L Rickman. Decoupling of heat generated from ejected and non-ejected contents of 18650-format lithium-ion cells using statistical methods. *Journal of Power Sources*, 415:207–218, 2019.

Fujin Wang, Zhi Zhai, Bingchen Liu, Shiyu Zheng, Zhibin Zhao, and Xuefeng Chen. Open access dataset, code library and benchmarking deep learning approaches for state-of-health estimation of lithium-ion batteries. *Journal of Energy Storage*, 77:109884, 2024.

Yongzhi Zhang, Rui Xiong, Hongwen He, and Michael G Pecht. Long short-term memory recurrent neural network for remaining useful life prediction of lithium-ion batteries. *IEEE Transactions on Vehicular Technology*, 67(7):5695–5705, 2018.

Appendix A. Model Architecture Summary

While we only present results from a simple elastic net model in the manuscript, it is worth exploring the performance of other models — both simple and complex — for the data introduced in this study. Here, we perform the same analysis as in Section 3.1 in the manuscript using the cross-validation strategy detailed in the same section. Fig. A.1 shows parity plots for different architectures using (i) $\Delta Q(V)$ features, (ii) DCIR-based features, and (iii) all features. The same cycle pairing as in the manuscript has been used.

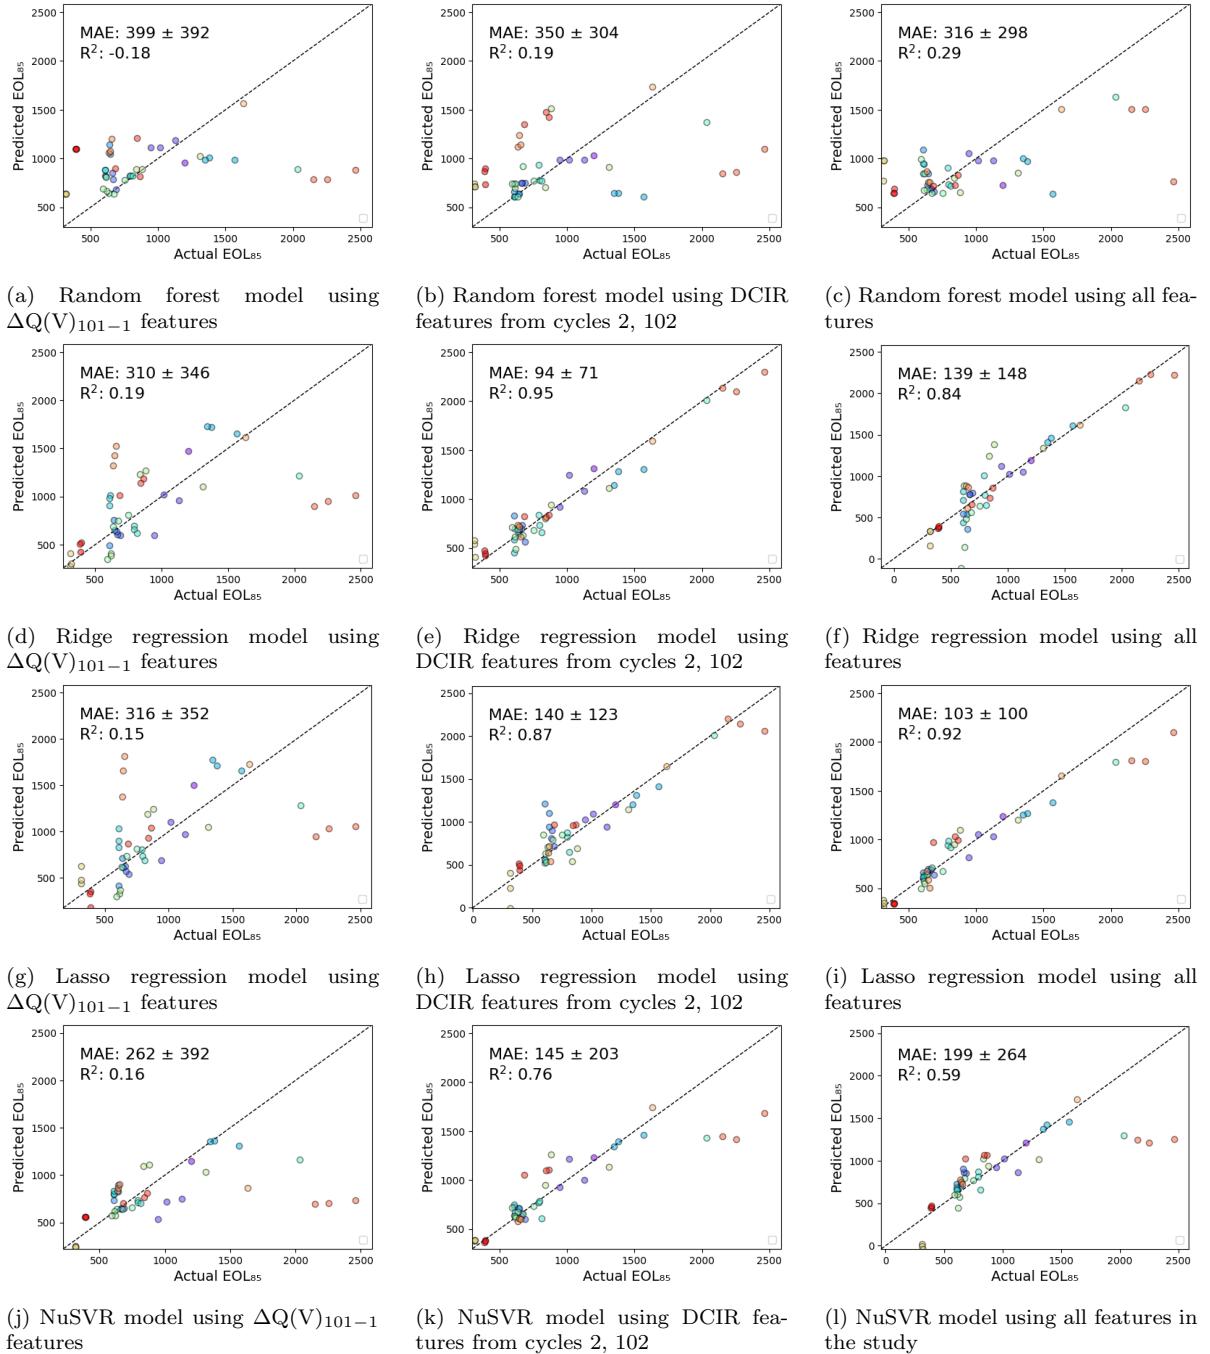
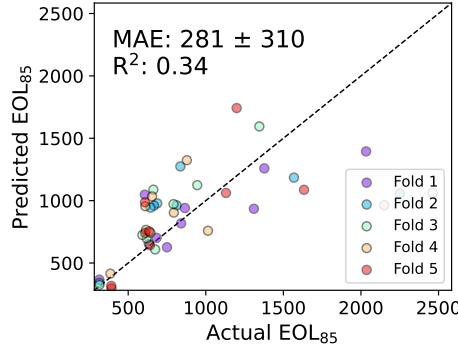


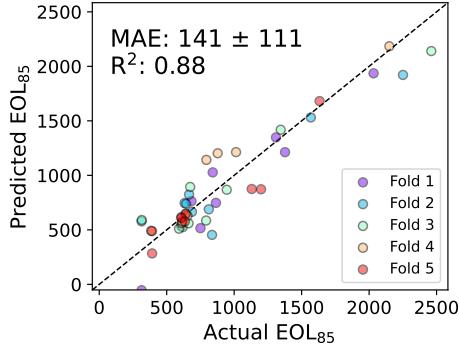
Figure A.1: A parity plot comparison for multiple architectures. The first column of figures uses the original six feature $\Delta Q(V)$ set from Severson et al., the second column of figures uses the novel augmented DCIR $\Delta Q(V)$ features developed in this study, and the third column of figures uses all features. Each row represents a new architecture.

Appendix B. Model Validation

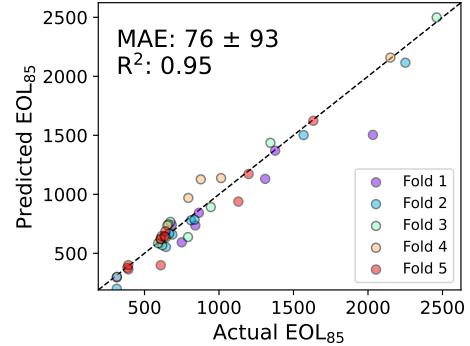
In addition to the "Leave One Triplicate Out" cross validation strategy employed in the manuscript, we perform traditional 5 fold cross-validation for direct comparison. The following parity plots and table are equivalent analyses to that presented in the main text, but with sklearn's standard KFold cross validator with 5 splits (80/20 training/testing split). Shown in the parity plots of Fig. B.1 are the compiled test set predictions of all 5 test folds, and the reported MAE and R2 are the averaged values of all test folds.



(a) Model performance using $\Delta Q(V)_{101-1}$ features with 5fold cross-validation.



(b) Model performance using DCIR $\Delta Q(V)_{102-2}$ features with 5fold cross-validation.



(c) Model performance using full feature set with 5fold cross-validation.

Figure B.1: A parity plot comparison for elastic net model performance using (a) the original six feature $\Delta Q(V)$ set from Severson et al, (b) the novel augmented DCIR $\Delta Q(V)$ feature set, and (c) the full feature set described in the manuscript to predict EOL₈₅.

Features	EOL ₈₀	EOL ₈₅	EOL ₉₀
$\Delta Q(V)_{101-1}$	$359 \pm 433 (0.06)$	$259 \pm 333 (-0.01)$	$134 \pm 136 (0.26)$
Non-DCIR features	$366 \pm 438 (0.03)$	$256 \pm 331 (0.01)$	$122 \pm 113 (0.44)$
DCIR-based features	$202 \pm 168 (0.80)$	$109 \pm 119 (0.85)$	$40 \pm 38 (0.94)$
All features	$163 \pm 127 (0.87)$	$145 \pm 165 (0.73)$	$46 \pm 63 (0.88)$

Table B.1: Model performance for Elastic Net architecture using a standard 5-fold cross validation technique

Appendix C. Feature correlation

Feature correlation (on a log scale) for DCIR-based univariate features is shown in Fig C.1. For simplicity, only data for one pulse has been shown. The Pearson correlation

coefficient for each has been provided as well. Generally speaking, we do not observe high feature correlation with the end of life of a battery, suggesting that principal component analysis to reduce the feature space is a key component of the machine learning workflow.

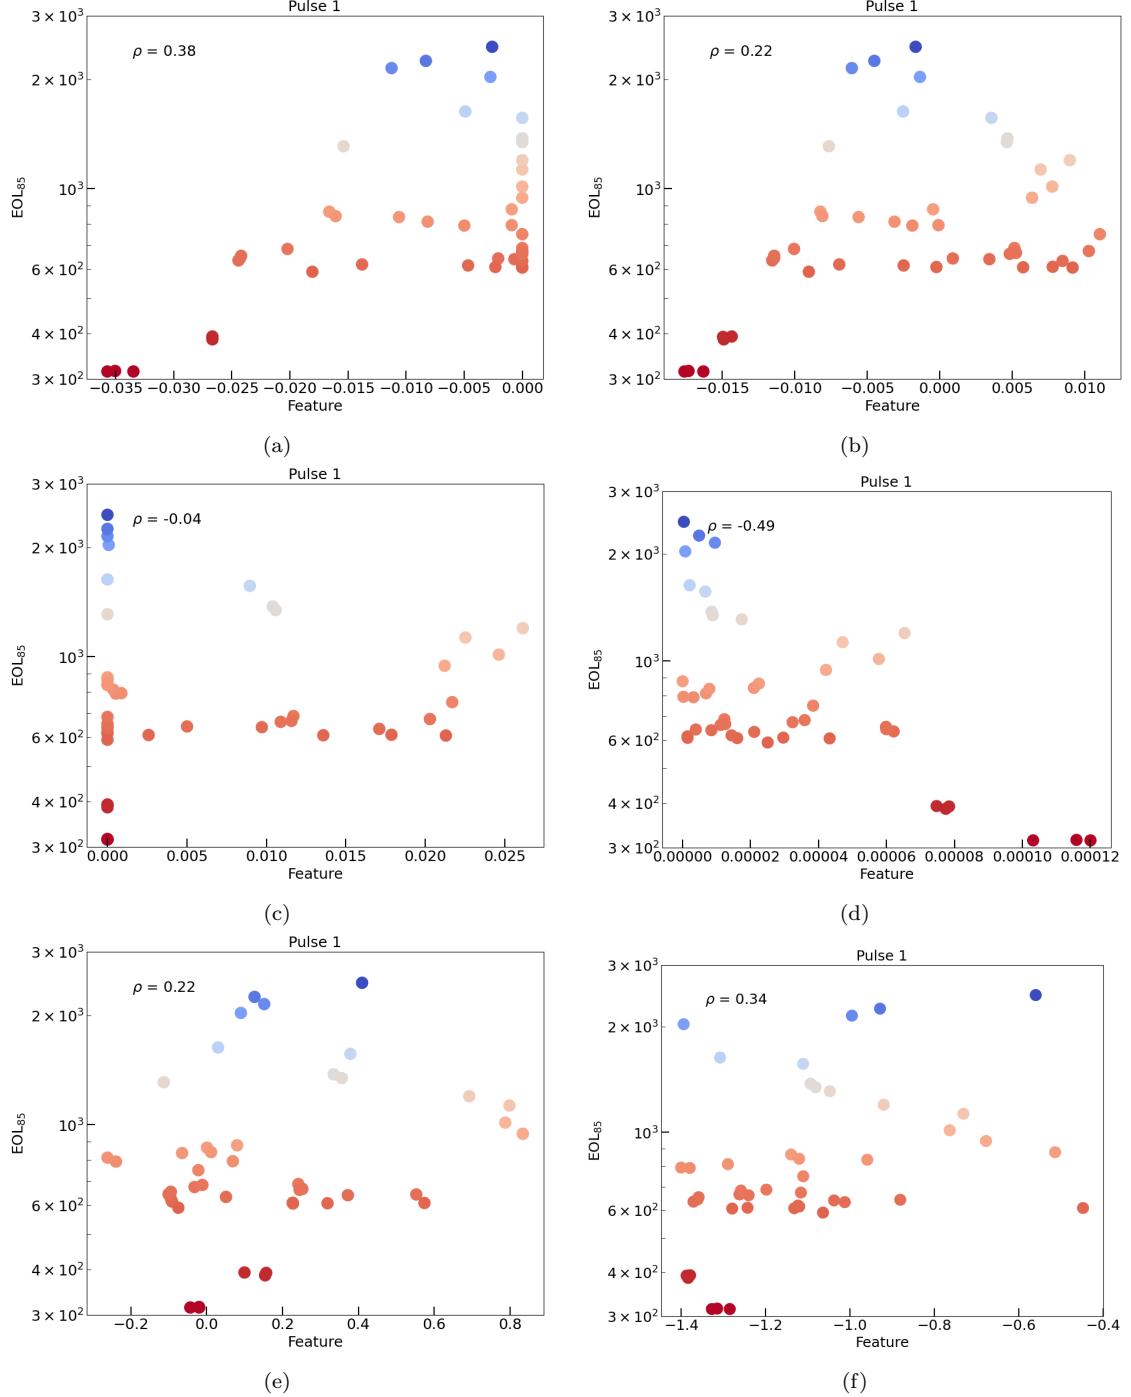


Figure C.1: Feature correlation (log scale) of univariate DCIR-based features used in the study. For simplicity, only features from the first DCIR pulse have been shown, computed from cycles 2 and 102. The six univariate features are (a) minimum, (b) maximum, (c) mean, (d) variance, (e) skew, and (f) kurtosis of the mini $\Delta Q(V)$ curve curved obtained from the first DCIR pulse, as explained in the main manuscript. The Pearson's correlation coefficient has been displayed on each plot as well.