

Title: Data-driven modeling enables cycle life prediction for lithium-ion batteries using early-cycle data yet to exhibit capacity degradation

Authors: Kristen A. Severson^{1*}, Peter M. Attia^{2*}, Norman Jin², Zi Yang³, Nicholas Perkins², Michael H. Chen⁴, Muratahan Aykol⁵, Patrick K. Herring⁵, Dimitrios Fraggedakis¹, Martin Z. Bazant¹, Stephen J. Harris⁶, William C. Chueh^{2†}, Richard D. Braatz^{1†}

* These authors contributed equally to this work

† Corresponding authors: wchueh@stanford.edu, braatz@mit.edu

Affiliations:

1 Department of Chemical Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA

2 Department of Materials Science and Engineering, Stanford University, Stanford, CA, USA

3 Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, USA

4 Department of Mechanical Engineering, Stanford University, Stanford, CA, USA

5 Toyota Research Institute, Los Altos, CA, USA

6 Materials Science Division, Lawrence Berkeley National Lab, Berkeley, CA, USA

Abstract

Accurately predicting lifetime of complex systems like lithium-ion batteries is crucial for accelerating technology development. However, diverse aging

mechanisms, significant device variability, and varied operating conditions have remained major challenges. To study this problem, we generated a dataset consisting of 124 commercial lithium-iron-phosphate/graphite cells cycled under fast charging conditions. The cells exhibited widely varied cycle lives spanning from 150 to 2,300 cycles, with end-of-life defined as 20% degradation from nominal capacity. Using discharge voltage curves from early cycles yet to exhibit capacity degradation, we apply machine learning tools to predict cycle life with less than 15% error on average, which is improved to ~8% error by incorporating additional data. Our work represents a significant improvement over previous predictions that generally required data corresponding to >5% capacity degradation, without needing specialized diagnostics. Additionally, it highlights the promise of combining data generation with data-driven modeling to predict the behavior of complex and variable systems.

Main

Lithium-ion batteries are deployed in a wide range of applications due to their low and falling costs, high energy densities, and long cycle lives.¹⁻³ However, as is the case with many chemical, mechanical, and electronics systems, long battery cycle life implies delayed feedback of performance during development and manufacture, often many months to years. Accurately predicting cycle life using early-cycle data would accelerate this feedback loop as well as enable estimation of battery life expectancy for use in consumer electronics, electric vehicles, and second-life applications.⁴⁻⁶

However, the task of predicting capacity fade and/or cycle life for lithium-ion batteries is challenging because of nonlinear degradation with cycling and wide variability, even when controlling for operating conditions.⁷⁻¹¹

Many studies have modeled cycle life in lithium-ion batteries. Bloom *et al.*¹² and Broussely *et al.*¹³ performed early work that fit semi-empirical models to predict power and capacity loss. Since then, many authors have proposed physical and semi-empirical battery degradation models that account for diverse mechanisms such as growth of the solid-electrolyte interphase¹⁴⁻¹⁷, lithium plating^{18,19}, active material loss^{20,21}, and impedance increase²²⁻²⁴. Predictions of remaining useful life in battery management systems, summarized in these reviews^{5,6}, often rely on these mechanistic and semi-empirical models. Specialized diagnostic measurements such as coulombic efficiency^{25,26} and impedance spectroscopy²⁷⁻²⁹ can also estimate cycle life. While these chemistry and/or mechanism-specific models have shown predictive success, developing models that describe full cells cycled under relevant conditions (*e.g.*, fast charging) remains challenging, given the many degradation modes and their coupling to the thermal^{30,31} and mechanical^{30,32} heterogeneities within a cell³²⁻³⁴.

Approaches using statistical and machine learning techniques to predict cycle life are attractive, mechanism-agnostic alternatives. Recently, advances in computational power and data generation have enabled these techniques to accelerate progress in a variety of fields, including materials discovery for energy storage³⁵⁻³⁷ and catalysis³⁸, and prediction of material

properties^{39,40}. A growing body of literature^{6,41,42} applies machine learning techniques for predicting remaining useful life using data collected in controlled laboratory environments and during real world operation. Generally, predictions are made after accumulating data corresponding to >5% of degradation from the initial capacity⁴³⁻⁴⁹ or using specialized measurements at the beginning of life¹¹. Accurate early prediction of cycle life with significantly less degradation is challenging because of the typically nonlinear degradation process (with negligible capacity degradation in early cycles) as well as the relatively small datasets used to date that span a limited range of degradation rates⁴⁸. For example, Harris *et al.*¹⁰ correlated capacity values at cycle 80 to capacity values at cycle 500 for 24 cells exhibiting nonlinear degradation profiles, identifying only a weak correlation ($\rho=0.1$). In short, opportunities for improving upon state-of-the-art prediction models include higher accuracy, earlier prediction, interpretability, and validation over a wider range of lifetimes.

In this work, we develop data-driven models that accurately predict the cycle life of commercial lithium-iron-phosphate (LFP)/graphite cells using early-cycle data, with no prior knowledge or assumption of degradation mechanisms. Here, cycle life (or equivalently, end-of-life) is defined as the number of cycles until 80% of nominal capacity. We generated a dataset consisting of 124 cells with cycle lives ranging from 150 to 2,300 by using 72 different fast-charging conditions. Utilizing information from early cycles yet to exhibit capacity degradation, our feature-based models achieve prediction

errors of 15% using discharge voltage curves using the 10th and 100th cycles. The error further reduces to ~8% by incorporating data from additional cycles and data streams. These results illustrate the power of combining data generation with data-driven modeling to predict the behavior of complex systems far into the future.

Data generation

Because of the large number of capacity fade mechanisms and manufacturing variability of lithium-ion batteries, we expect the space that parameterizes capacity fade to be high dimensional. To probe this space, commercial LFP/graphite cells (A123 Systems, model APR18650M1A, 1.1 Ah nominal capacity) were cycled in a temperature-controlled environmental chamber (30 °C) under different fast charging conditions but identical discharging conditions (4C to 2.0V, where 1C is 1.1A; see Methods for details). By deliberately varying the charging conditions, we generate a dataset that captures a wide range of cycle lives, from approximately 150 to 2,300 cycles (average cycle life of 806 with a standard deviation of 377). Voltage, current, cell casing temperature, and internal resistance are measured during cycling. A full description of the experimental details can be found in the Methods section. The dataset contains approximately 96,700 cycles; to the best of the authors' knowledge, our dataset constitutes the largest publicly available one consisting of nominally identical commercial lithium-ion batteries cycled under controlled yet varied conditions (see data availability section for access information).

Figure 1 shows the discharge capacity as a function of cycle number for the first 1,000 cycles, where the color denotes the cycle life. The capacity fade is negligible in the first one hundred cycles and accelerates near the end of life, as is often observed in lithium-ion batteries.^{8,10,11} The fact that the capacity fade trajectories cross each other illustrates the weak relationship between initial capacity and lifetime. Indeed, we find weak correlations between the log of cycle life and the discharge capacity at the 2nd cycle ($\rho=-0.06$, Figure 1d) and the 100th cycle ($\rho=0.27$, Figure 1e), as well as between the log of cycle life and the capacity fade rate near cycle 100 ($\rho=0.47$, Figure 1f). These weak correlations are expected because capacity degradation in these early cycles is negligible; in fact, the capacities at cycle 100 increased from the initial values for 81% of cells in our dataset (Figure 1c). Given the limited predictive power of these correlations based on the capacity fade curves, we employ an alternative data-driven approach that considers a larger set of cycling data which includes the full voltage curves of each cycle, as well as additional measurements including cell internal resistance and temperature.

Machine learning approach

We develop a feature-based approach for building an early prediction model. In this paradigm, features, which are transformations of the raw data, are generated and used in a regularization framework. The final model uses a linear combination of a subset of the proposed features to predict the logarithm of cycle life. This subset is identified using the elastic net⁵¹. A

description of the computational framework can be found in the Methods section.

We propose features (Table 1) based on domain knowledge of lithium-ion batteries (though agnostic to chemistry and degradation mechanisms), such as initial discharge capacity, charge time, and cell can temperature. To capture the electrochemical evolution of individual cells during cycling, several features are calculated based on the discharge voltage curve (Figure 2a). Specifically, we consider the cycle-to-cycle evolution of $Q(V)$, the discharge voltage curve as a function of voltage for a given cycle. As the voltage range is identical with every cycle, we consider capacity as a function of voltage, as opposed to voltage as a function of capacity, in order to maintain a uniform basis for comparing cycles. For instance, we can consider the change in discharge voltage curve between cycles 20 and 30, denoted $\Delta Q_{30-20}(V) = Q_{30}(V) - Q_{20}(V)$, where the subscripts indicate the cycle number. This data transformation, $\Delta Q(V)$, is of particular interest because voltage curves and their derivatives are a rich data source that have been effective in degradation diagnosis⁵²⁻⁵⁸.

The $\Delta Q(V)$ curves for our dataset are shown in Figure 2b using the 100th and 10th cycles, *i.e.*, $\Delta Q_{100-10}(V)$. Summary statistics, *e.g.* minimum, mean, and variance, were then calculated for the $\Delta Q_{100-10}(V)$ curves of each cell. Each summary statistic is a scalar quantity that captures the change in the voltage curves between two given two cycles. In our data-driven approach, these summary statistics are selected based on their predictive

ability, not their physical meaning. Immediately, a clear trend emerges between the cycle life of a cell and a summary statistic, specifically variance, applied to $\Delta Q_{100-10}(V)$ (Figure 2c).

Because of the high predictive power of features based on $\Delta Q_{100-10}(V)$, we investigate models (1) using only the variance of $\Delta Q_{100-10}(V)$, (2) considering additional candidate features obtained during discharge, and (3) considering features including both charging and discharging. In all cases, data were taken from the first 100 cycles. These three models, each with progressively more features, were chosen to evaluate both the cost-benefit of acquiring additional data streams such as temperature and the limits of prediction accuracy. The complete set of 20 candidate features is shown in Table 1 and is described in detail in the Supplemental Information. The training data (41 cells) is used to select the model form and set the values of the coefficients, and the primary testing data (43 cells) is used to evaluate the model performance. We then evaluate the model on a secondary, unseen testing dataset of 40 cells generated after model development. Two metrics, defined in the Computational Methods section, are used to evaluate performance of the cycle life prediction: root-mean-squared error (RMSE), with units of cycles, and average percent error.

Results

We present three models to predict cycle life using increasing candidate feature set sizes. The first model, denoted as the “variance” model, does not consider subset selection and uses only the log variance of

$\Delta Q_{100-10}(V)$ for prediction. Surprisingly, using only this single feature results in a model that has approximately 15% average percent error on the primary test dataset and approximately 11% average percent error on the secondary test dataset. We stress the error metrics of the secondary test dataset, as these data had not been generated at the time of model development and are thus a rigorous test of model generalizability. The second, “discharge” model, considers additional information derived from measurements of voltage and current during discharge in the first 100 cycles (row blocks 1 and 2 of Table 1). Six out of thirteen features were selected. Finally, the third, “full” model, considers all available features (all rows blocks of Table 1). In this model, nine out of twenty features were selected. As expected, by adding additional features, the test average percent error decreases to 7.5% and the additional test average percent error decreases slightly to 10.7%. In all cases, the average percent error is less than 15% and reduces to as low as ~8% in the full model, excluding an anomalous cell. Table 2 and Figure 3 display the performance of the “variance”, “discharge”, and “full” models applied to the train, primary test, and secondary test datasets, and the specific features and model coefficients used in the full model are displayed in Figure 4.

We benchmark the performance of our cycle life prediction using early-cycle data against both prior literature and naïve models. A relevant metric is the extent of degradation that has to occur before an accurate prediction can be made. In our work, accurate prediction was achieved using voltage

curves from early cycles corresponding to capacity increase of 0.2% (median) relative to initial values (with the first and third quartile percentiles being 0.06% and 0.34%, respectively). We are not aware of previous early prediction demonstrations that do not require degradation in the battery capacity nor specialized measurements. In fact, published models generally require data corresponding to at least 5% capacity degradation before making a prediction at an accuracy comparable to this work⁴³⁻⁴⁹. We also benchmark our model performance using naïve models, e.g. univariate models and/or models that only utilize information from the capacity fade curve (see Supplementary Information, Benchmarking models section). Notably, if the average cycle life of the training data is used for prediction, the average percent error is approximately 30% and 36% for the primary and secondary test data, respectively. The best benchmark model has errors of 25% and 34% for the primary and secondary test data, respectively.

While models that include features based on additional data streams such as internal resistance and casing temperature generally have the lowest errors, the primary predictive ability comes from the variance of $\Delta Q_{100-10}(V)$ feature, as evidenced by the performance of the single-feature “variance” model. This feature is consistently selected in both models with feature selection (“discharge” and “full”). Other transformations of this trajectory can also be used to predict cycle life, alone or in combination with variance. For example, the full model selects the minimum and variance of the $\Delta Q_{100-10}(V)$ features. The physical meaning of the variance feature is

associated with the dependence of the discharged energy dissipation on voltage, which is indicated by the green region between the voltage curves in Fig. 2a. The integral of this region is the total change in energy dissipation between cycles under galvanostatic conditions and is linearly related to the mean of $\Delta Q(V)$. Zero variance would indicate energy dissipations that do not depend on voltage. Thus, the variance of $\Delta Q(V)$ reflects the extent of non-uniformity in the energy dissipation with voltage, due to either open-circuit or kinetic processes, a point that we return to later.

Discussion

We observe that features derived using early-cycle discharge voltage curve have excellent predictive performance, even before the onset of capacity fade. We rationalize this observation by investigating degradation modes that do not immediately result in capacity fade yet still manifest in the discharge voltage curve, and are also linked to rapid capacity fade near the end-of-life.

While our data-driven approach has successfully revealed predictive features of cycle life from early cycle discharge curves, identification of the degradation modes using only high rate data is challenging because of the convolution of kinetics with open-circuit behavior. Dubarry *et al.*⁵⁵ used low-rate diagnostic cycles to remove these kinetic effects and mapped degradation modes in LFP/graphite cells to their resultant shift in dQ/dV and dV/dQ derivatives for diagnostic cycles at C/20. One degradation mode – loss of active material of the delithiated negative electrode (LAM_{deNE}) – results in a

shift in discharge voltage with no change in capacity. This behavior is observed when the negative electrode is oversized relative to the positive electrode, as is the case in the LFP/graphite cells examined in this work. Thus, a loss of delithiated negative electrode material changes the potentials at which lithium ions are stored without changing the overall capacity.^{55,56} As proposed by Anséan *et al.*⁵⁶, at high rates of LAM_{deNE} , the negative electrode capacity will eventually fall below the lithium-ion inventory remaining in the cell. At this point, the negative electrode will not have enough sites to accommodate lithium ions during charging, inducing lithium plating.⁵⁶ Since plating is an additional source of irreversibility, the capacity loss accelerates. Thus, in early cycles, LAM_{deNE} shifts the voltage curve without affecting the capacity fade curve and induces rapid capacity fade at high cycle number. This degradation mode, in conjunction with loss of lithium inventory (LLI), is widely observed in commercial LFP/graphite cells operated under similar conditions.^{34,54,56,57,59,60}

To investigate the contribution of LAM_{deNE} to degradation in our experiments, additional experiments were performed for cells cycled with varied charging rates (4C, 6C, and 8C) and a constant discharge rate (4C), incorporating slow cycling at the 1st, 100th, and end of life cycles. Derivatives of diagnostic discharge curves at C/10 (Figure 5, rows 1 and 2) are compared to those, and $\Delta Q(V)$, at 4C at the 10th, 101st, and end of life cycles (rows 3 and 4). The shifts in dQ/dV and dV/dQ observed in diagnostic cycling correspond to a shift of the potentials at which lithium is stored in graphite

during charging and are consistent with LAM_{deNE} and LLI operating concurrently.⁵⁵⁻⁵⁷ The magnitude of these shifts increases with charging rate. These observations rationalize why models using features based on discharge curves have lower errors than models using only features based on capacity fade curves, since LAM_{deNE} does not manifest in capacity fade in early cycles. Other degradation modes that do not initially manifest in capacity fade have been reported, such as high-voltage cathode materials undergoing voltage fade.^{61,62} We also note that LAM_{deNE} alters a fraction of, rather than the entire, discharge voltage curve, consistent with the observed correlation between the variance of $\Delta Q_{100-10}(V)$ and cycle life.

We recognize that the above rationalization uses low-rate diagnostic cycling, which is largely not affected by kinetics. However, our predictions were developed using high-rate discharge data. As such, these discharge voltage curves can reflect kinetic degradation modes that are not observed in dQ/dV and dV/dQ derivatives at $C/10$. Because we perform a constant-voltage hold at the end of both charge and discharge, kinetic degradation modes may manifest in the discharge voltage curve but not in the capacity fade curve. We note that the change in the discharge energy between the diagnostic cycles (1st and 100th cycles) is 53% to 66% of the change between the 10th and 101st high-rate cycles, indicating the degradation is influenced by both low-rate and high-rate degradation modes (see Supplemental Information, Kinetic degradation section). These kinetic contributions during early cycles may also be linked to cycle life, especially nonlinearities in

reaction kinetics that could skew the voltage curves non-uniformly⁶³, and are part of an ongoing investigation.

As noted above, differential methods like dQ/dV and dV/dQ are used extensively to pinpoint degradation mechanisms⁵²⁻⁵⁷. These approaches require low-rate diagnostic cycles, as higher rates smear out features due to overpotential, as seen by comparing row 1 to row 3 in Figure 5. These diagnostic cycles often induce a temporary capacity recovery that interrupts the trajectory of capacity fade (see Supplemental Information, Diagnostic cycling section), complicating the history of the battery^{64,65}. Therefore, by applying summary statistics to $\Delta Q(V)$ collected at high rates, we simultaneously avoid both low-rate diagnostic cycles and numerical differentiation, which decreases the signal-to-noise ratio⁶⁶.

Finally, additional analysis during model development was performed to understand the impact of the cycle indices chosen for $\Delta Q(V)$ features. Linear models using only the variance of the difference $Q_i(V) - Q_j(V)$ for the training and testing datasets were investigated and are displayed in Figure 6. We find that the model is relatively insensitive to the indexing scheme for $i > 80$. This trend is further validated by the model coefficients shown in Figure S8. We hypothesize that the insensitivity of the model to the indexing scheme implies linear degradation with respect to cycle number, which is often assumed for LAM modes^{55,56}. Relative indexing schemes based on cycles in which a specified capacity fade was achieved were also investigated and did not result in improved predictions. Furthermore,

because the discharge capacity initially increases, specified *decreases* in capacity take longer to develop in terms of cycles than fixed indexing (see Supplemental Information, Relative indexing schemes section).

Conclusion

Data-driven modeling is critical in diagnostics and prognostics of lithium-ion batteries cycled under relevant conditions. We develop cycle life prediction models using early-cycle discharge data yet to exhibit degradation, generated from commercial LFP/graphite batteries cycled under fast charging conditions. The models achieve prediction errors of 15% using data from only the 10th and 100th cycles and errors as low as 8% using the first 100 cycles for batteries with lifetimes ranging from 150 to 2,300 cycles. This level of accuracy is achieved by extracting features from high-rate discharge voltage curves as opposed to only from the capacity fade curves, and without using data from slow diagnostic cycles nor assuming prior knowledge of cell chemistry and degradation mechanisms. The success of the model is rationalized by demonstrating consistency with degradation modes that do not manifest in capacity fade during early cycles but do impact the voltage profiles. In general, we expect our early prediction models to be especially effective for degradation modes that do not initially contribute to capacity fade, such as voltage fade in high-voltage cathode materials. Our approach can complement approaches based on physical and semi-empirical models and on specialized diagnostics. Broadly speaking, this work highlights the promise of combining data generation and data-driven

modeling for understanding and developing complex systems such as lithium-ion batteries.

Methods:

Experimental

Commercial high-power LFP/graphite A123 APR18650M1A cells were used in this work. The cells have a nominal capacity of 1.1 Ah and a nominal voltage of 3.3V. The manufacturer's recommended fast-charging protocol is 3.6C CC-CV. Rate capability of these cells is shown in Figure S9.

All cells were tested in cylindrical fixtures with 4-point contacts on a 48-channel Arbin LBT battery testing cycler. The tests were performed at a constant temperature of 30°C in an environmental chamber (Amerex Instruments). Cell can temperatures were recorded by stripping a small section of the plastic insulation and contacting a Type T thermocouple to the bare metal casing using thermal epoxy (OMEGATHERM 201) and Kapton tape.

The cells were cycled with various charging policies but identically discharged. Cells were charged from 0% to 80% SOC with various single-step and two-step charging policies. The charging time from 0% to 80% SOC ranged from 9 to 13.3 minutes. An internal resistance measurement was obtained during charging at 80% SOC by averaging 10 pulses of $\pm 3.6C$ with a pulse width of 30 ms, where 1C is 1.1 A, or the current required to fully (dis)charge the nominal capacity (1.1 Ah) in 1 hour. All cells then charged

from 80% to 100% SOC with a uniform 1C CC-CV charging step to 3.6V and a current cutoff of C/50. All cells were subsequently discharged with a CC-CV discharge at 4C to 2.0V with a current cutoff of C/50. The voltage cutoffs used in this work follow those recommended by the manufacturer.

To standardize the voltage-capacity data across cells and cycles, all 4C discharge curves were fit to a spline function and linearly interpolated (see Figure S10). Capacity was fit as a function of voltage and evaluated at 1000 linearly-spaced voltage points from 3.5V to 2.0V. These uniformly-sized vectors enabled straightforward data manipulations such as subtraction.

Computational

This study involved both model fitting, selection of the coefficient values, and model selection (selection of the model structure). To perform both of these tasks simultaneously, a regularization technique was employed. A linear model of the form

$$\hat{y}_i = \hat{w}^T x_i \quad (1)$$

was proposed, where \hat{y}_i is the predicted number of cycles for battery i , x_i is a p -dimensional feature vector for battery i , and \hat{w} is a p -dimensional model coefficient vector. When applying regularization techniques, a penalty term is added to the least-squares optimization formulation to avoid over-fitting. Two regularization techniques, the lasso⁶⁷ and the elastic net⁵¹, simultaneously perform model fitting and selection by finding sparse coefficient vectors. The formulation is

$$\hat{w} = \underset{w}{\operatorname{argmin}} \|y - Xw\|_2^2 + \lambda P(w) \quad (2)$$

where y is the n -dimensional vector of observed battery lifetimes, X is the $n \times p$ matrix of features, λ is a non-negative scalar. The term

$$\|y - Xw\|_2^2 \quad (3)$$

is found in ordinary least squares and is also referred to as squared loss because the optimization is minimizing the squared error. The formulation of the second term, $P(w)$, depends on the regularization technique being employed. For the lasso,

$$P(w) = \|w\|_1, \quad (4)$$

and

$$P(w) = \frac{1-\alpha}{2} \|w\|_2^2 + \alpha \|w\|_1 \quad (5)$$

for the elastic net. Both formulations will result in sparse models. The elastic net has been shown to perform better when $p \gg n$,⁵¹ as is often the case in feature engineering applications, but it requires fitting an additional hyper-parameter (α and λ , as opposed to only λ in the lasso). The elastic net is also preferred when there are high correlations between the features, as is the case in this application. To choose the value(s) of the hyper-parameter(s), we apply 4-fold cross validation and Monte Carlo sampling.

The model development dataset is divided into two equal sections, referred to as the training and primary testing data. The training data is used to choose the hyper-parameters α and λ and determine the values of the

coefficients, w . The training data is further subdivided into calibration and validation sets for cross-validation. The testing data is then used as a measure of generalizability because this data has not been used to learn the model coefficients or form.

Root-mean-squared error (RMSE) and average percent error are chosen to evaluate model performance. RMSE is defined as

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (6)$$

where y_i is the observed cycle life, \hat{y}_i is the predicted cycle life, and n is the total number of samples. Average percent error is defined as

$$\%err = \frac{1}{n} \sum_{i=1}^n \frac{|y_i - \hat{y}_i|}{y_i} \times 100 \quad (7)$$

where all variables are defined as above.

References

1. Nykvist, B. & Nilsson, M. Rapidly falling costs of battery packs for electric vehicles. *Nat. Clim. Change* **5**, 329–332 (2015).
2. Schmuch, R., Wagner, R., Hörpel, G., Placke, T. & Winter, M. Performance and cost of materials for lithium-based rechargeable automotive batteries. *Nat. Energy* **3**, 267–278 (2018).
3. Cano, Z. P. *et al.* Batteries and fuel cells for emerging electric vehicle markets. *Nat. Energy* **3**, 279–289 (2018).
4. Peterson, S. B., Apt, J. & Whitacre, J. F. Lithium-ion battery cell degradation resulting from realistic vehicle and vehicle-to-grid utilization. *J. Power Sources* **195**, 2385–2392 (2010).
5. Ramadesigan, V. *et al.* Modeling and Simulation of Lithium-Ion Batteries from a Systems Engineering Perspective. *J. Electrochem. Soc.* **159**, R31–R45 (2012).
6. Waag, W., Fleischer, C. & Sauer, D. U. Critical review of the methods for monitoring of lithium-ion batteries in electric and hybrid vehicles. *J. Power Sources* **258**, 321–339 (2014).
7. Paul, S., Diegelmann, C., Kabza, H. & Tillmetz, W. Analysis of ageing inhomogeneities in lithium-ion battery systems. *J. Power Sources* **239**, 642–650 (2013).
8. Schuster, S. F. *et al.* Nonlinear aging characteristics of lithium-ion cells under different operational conditions. *J. Energy Storage* **1**, 44–53 (2015).

9. Schuster, S. F., Brand, M. J., Berg, P., Gleissenberger, M. & Jossen, A. Lithium-ion cell-to-cell variation during battery electric vehicle operation. *J. Power Sources* **297**, 242–251 (2015).
10. Harris, S. J., Harris, D. J. & Li, C. Failure statistics for commercial lithium ion batteries: A study of 24 pouch cells. *J. Power Sources* **342**, 589–597 (2017).
11. Baumhöfer, T., Brühl, M., Rothgang, S. & Sauer, D. U. Production caused variation in capacity aging trend and correlation to initial cell performance. *J. Power Sources* **247**, 332–338 (2014).
12. Bloom, I. *et al.* An accelerated calendar and cycle life study of Li-ion cells. *J. Power Sources* **101**, 238–247 (2001).
13. Broussely, M. *et al.* Aging mechanism in Li ion cells and calendar life predictions. *J. Power Sources* **97-98**, 13–21 (2001).
14. Christensen, J. & Newman, J. Effect of anode film resistance on the charge/discharge capacity of a lithium-ion battery. *J. Electrochem. Soc.* **150**, A1416–A1420 (2003).
15. Ploehn, H. J., Ramadass, P. & White, R. E. Solvent Diffusion Model for Aging of Lithium-Ion Battery Cells. *J. Electrochem. Soc.* **151**, A456 (2004).
16. Broussely, M. *et al.* Main aging mechanisms in Li ion batteries. *J. Power Sources* **146**, 90–96 (2005).
17. Pinson, M. B. & Bazant, M. Z. Theory of SEI formation in rechargeable batteries: Capacity fade, accelerated aging and lifetime prediction. *J. Electrochem. Soc.* **160**, A243–A250 (2013).

18. Arora, P. Mathematical Modeling of the Lithium Deposition Overcharge Reaction in Lithium-Ion Batteries Using Carbon-Based Negative Electrodes. *J. Electrochem. Soc.* **146**, 3543 (1999).
19. Yang, X.-G., Leng, Y., Zhang, G., Ge, S. & Wang, C.-Y. Modeling of lithium plating induced aging of lithium-ion batteries: Transition from linear to nonlinear aging. *J. Power Sources* **360**, 28–40 (2017).
20. Christensen, J. & Newman, J. Cyclable lithium and capacity loss in li-ion cells. *J. Electrochem. Soc.* **152**, A818–A829 (2005).
21. Zhang, Q. & White, R. E. Capacity fade analysis of a lithium ion cell. *J. Power Sources* **179**, 793–798 (2008).
22. Wright, R. . *et al.* Power fade and capacity fade resulting from cycle-life testing of Advanced Technology Development Program lithium-ion batteries. *J. Power Sources* **119-121**, 865–869 (2003).
23. Ramadesigan, V. *et al.* Parameter estimation and capacity fade analysis of lithium-ion batteries using reformulated models. *J. Electrochem. Soc.* **158**, A1048–A1054 (2011).
24. Cordoba-Arenas, A., Onori, S., Guezennec, Y. & Rizzoni, G. Capacity and power fade cycle-life model for plug-in hybrid electric vehicle lithium-ion battery cells containing blended spinel and layered-oxide positive electrodes. *J. Power Sources* **278**, 473–483 (2015).
25. Burns, J. C. *et al.* Evaluation of effects of additives in wound Li-ion cells through high precision coulometry. *J. Electrochem. Soc.* **158**, A255–A261 (2011).

26. Burns, J. C. *et al.* Predicting and extending the lifetime of Li-ion batteries. *J. Electrochem. Soc.* **160**, A1451-A1456 (2013).
27. Chen, C. H., Liu, J. & Amine, K. Symmetric cell approach and impedance spectroscopy of high power lithium-ion batteries. *J. Power Sources* **96**, 321-328 (2001).
28. Tröltzsch, U., Kanoun, O. & Tränkler, H.-R. Characterizing aging effects of lithium-ion batteries by impedance spectroscopy. *Electrochimica Acta* **51**, 1664-1672 (2006).
29. Love, C. T., Virji, M. B. V., Rocheleau, R. E. & Swider-Lyons, K. E. State-of-health monitoring of 18650 4S packs with a single-point impedance diagnostic. *J. Power Sources* **266**, 512-519 (2014).
30. Waldmann, T. *et al.* A mechanical aging mechanism in lithium-ion batteries. *J. Electrochem. Soc.* **161**, A1742-A1747 (2014).
31. Waldmann, T. *et al.* Influence of Cell Design on Temperatures and Temperature Gradients in Lithium-Ion Cells: An In Operando Study. *J. Electrochem. Soc.* **162**, A921-A927 (2015).
32. Bach, T. C. *et al.* Nonlinear aging of cylindrical lithium-ion cells linked to heterogeneous compression. *J. Energy Storage* **5**, 212-223 (2016).
33. Harris, S. J. & Lu, P. Effects of inhomogeneities - nanoscale to mesoscale - on the durability of Li-ion batteries. *J. Phys. Chem. C* **117**, 6481-6492 (2013).
34. Lewerenz, M., Marongiu, A., Warnecke, A. & Sauer, D. U. Differential voltage analysis as a tool for analyzing inhomogeneous aging: A case

- study for LiFePO₄|graphite cylindrical cells. *J. Power Sources* **368**, 57–67 (2017).
35. Jain, A. *et al.* Commentary: The Materials Project: A materials genome approach to accelerating materials innovation. *APL Mater.* **1**, 011002 (2013).
36. Aykol, M. *et al.* High-throughput computational design of cathode coatings for Li-ion batteries. *Nat. Commun.* **7**, 13779 (2016).
37. Sendek, A. D. *et al.* Holistic computational structure screening of more than 12000 candidates for solid lithium-ion conductor materials. *Energy Environ. Sci.* **10**, 306–320 (2017).
38. Ulissi, Z. W. *et al.* Machine-learning methods enable exhaustive searches for active bimetallic facets and reveal active site motifs for CO₂ reduction. *ACS Catal.* **7**, 6600–6608 (2017).
39. Raccuglia, P. *et al.* Machine-learning-assisted materials discovery using failed experiments. *Nature* **573**, 73–77 (2016).
40. Ward, L., Agrawal, A., Choudhary, A. & Wolverton, C. A general-purpose machine learning framework for predicting properties of inorganic materials. *Npj Comput. Mater.* **2**, 16028 (2016).
41. Si, X.-S., Wang, W., Hu, C.-H. & Zhou, D.-H. Remaining useful life estimation – A review on the statistical data driven approaches. *Eur. J. Oper. Res.* **213**, 1–14 (2011).

42. Wu, L., Fu, X. & Guan, Y. Review of the remaining useful life prognostics of vehicle lithium-ion batteries using data-driven methodologies. *Appl. Sci.* **6**, 166 (2016).
43. Saha, B., Goebel, K. & Christophersen, J. Comparison of prognostic algorithms for estimating remaining useful life of batteries. *Trans. Inst. Meas. Control* **31**, 293–308 (2009).
44. Nuhic, A., Terzimehic, T., Soczka-Guth, T., Buchholz, M. & Dietmayer, K. Health diagnosis and remaining useful life prognostics of lithium-ion batteries using data-driven methods. *J. Power Sources* **239**, 680–688 (2013).
45. Hu, C., Jain, G., Tamirisa, P. & Gorka, T. Method for estimating the capacity and predicting remaining useful life of lithium-ion battery. *Appl. Energy* **126**, 182–189 (2014).
46. Miao, Q., Xie, L., Cui, H., Liang, W. & Pecht, M. Remaining useful life prediction of lithium-ion battery with unscented particle filter technique. *Microelectron. Reliab.* **53**, 805–810 (2013).
47. Hu, X., Jiang, J., Cao, D. & Egardt, B. Battery Health Prognosis for Electric Vehicles Using Sample Entropy and Sparse Bayesian Predictive Modeling. *IEEE Trans. Ind. Electron.* 1–1 (2015).
doi:10.1109/TIE.2015.2461523
48. Zhang, Y., Xiong, R., He, H. & Pecht, M. Lithium-ion battery remaining useful life prediction with Box-Cox transformation and Monte Carlo

simulation. *IEEE Trans. Ind. Electron.* 1-1 (2018).

doi:10.1109/TIE.2018.2808918

49. Zhang, Y., Xiong, R., He, H. & Pecht, M. Long short-term memory recurrent neural network for remaining useful life prediction of lithium-ion batteries. *IEEE Trans. Veh. Technol.* 1-1 (2018).
doi:10.1109/TVT.2018.2805189
50. Saha, B. & Goebel, K. Battery Data Set. *NASA Ames Progn. Data Repos.* (2007).
51. Zou, H. & Hastie, T. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Methodol.* **67**, 301-320 (2005).
52. Bloom, I. *et al.* Differential voltage analyses of high-power, lithium-ion cells: 1. Technique and application. *J. Power Sources* **139**, 295-303 (2005).
53. Smith, A. J., Burns, J. C. & Dahn, J. R. High-precision differential capacity analysis of LiMn₂O₄/graphite cells. *Electrochem. Solid-State Lett.* **14**, A39-A41 (2011).
54. Safari, M. & Delacourt, C. Aging of a commercial graphite/LiFePO₄ cell. *J. Electrochem. Soc.* **158**, A1123-A1135 (2011).
55. Dubarry, M., Truchot, C. & Liaw, B. Y. Synthesize battery degradation modes via a diagnostic and prognostic model. *J. Power Sources* **219**, 204-216 (2012).
56. Anseán, D. *et al.* Fast charging technique for high power LiFePO₄ batteries: A mechanistic analysis of aging. *J. Power Sources* **321**, 201-209 (2016).

57. Anseán, D. *et al.* Operando lithium plating quantification and early detection of a commercial LiFePO₄ cell cycles under dynamic driving schedule. *J. Power Sources* **356**, 36–46 (2017).
58. Birkl, C. R., Roberts, M. R., McTurk, E., Bruce, P. G. & Howey, D. A. Degradation diagnostics for lithium ion cells. *J. Power Sources* **341**, 373–386 (2017).
59. Liu, P. *et al.* Aging mechanisms of LiFePO₄ batteries deduced by electrochemical and structural analyses. *J. Electrochem. Soc.* **157**, A499–A507 (2010).
60. Sarasketa-Zabala, E. *et al.* Understanding lithium inventory loss and sudden performance fade in cylindrical cells during cycling with deep-discharge steps. *J. Phys. Chem. C* **119**, 896–906 (2015).
61. Thackeray, M. M. *et al.* Li₂MnO₃-stabilized LiMO₂ (M = Mn, Ni, Co) electrodes for lithium-ion batteries. *J. Mater. Chem.* **17**, 3112–3125 (2007).
62. Gent, W. E. *et al.* Coupling between oxygen redox and cation migration explains unusual electrochemistry in lithium-rich layered oxides. *Nat. Commun.* **8**, 2091 (2017).
63. Bazant, M. Z. Thermodynamic stability of driven open systems and control of phase separation by electro-autocatalysis. *Faraday Discuss.* **199**, 423–463 (2017).
64. Rashid, M. & Gupta, A. Effect of relaxation periods over cycling performance of a Li-ion battery. *J. Electrochem. Soc.* **162**, A3145–A3153 (2015).

65. Lewerenz, M. *et al.* Systematic aging of commercial LiFePO₄|Graphite cylindrical cells including a theory explaining rise of capacity during aging. *J. Power Sources* **345**, 254–263 (2017).
66. Richardson, R. R., Birkl, C. R., Osborne, M. A. & Howey, D. Gaussian Process Regression for In-situ Capacity Estimation of Lithium-ion Batteries. *IEEE Trans. Ind. Inform.* 1–1 (2018). doi:10.1109/TII.2018.2794997
67. Tibshirani, R. Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B Methodol.* **58**, 267–288 (1996).

Contributions: P.M.A., N.J., N.P., M.C., and W.C.C. conceived of and conducted the experiments. K.A.S. and Z.Y. performed the modeling. M.A., Z.Y., P.K.H. performed data management. P.M.A., K.A.S., N.J., S.J.H., W.C.C., and R.D.B. interpreted the results. All authors edited and reviewed the manuscript. W.C.C. and R.D.B. supervised the work.

Competing interests: K.A.S., R.D.B, W.C.C., P.M.A., N.J, S.J.H., and N.P. have filed a provisional patent related to this work: US Application No. 62/575,565, dated October 17, 2018.

Data availability: The datasets generated during and/or analyzed during the current study are available in the figshare repository at [DOI pending].

Acknowledgements: This work was supported by Toyota Research Institute through the Accelerated Materials Design and Discovery program. P.M.A. acknowledges support from the Thomas V. Jones Stanford Graduate Fellowship and the National Science Foundation Graduate Research Fellowship under Grant No. DGE-114747. N.P. was supported by SAIC

Innovation Center through Stanford Energy 3.0. industry affiliates program. S.J.H. acknowledges support from the Assistant Secretary for Energy Efficiency, Vehicle Technologies Office of the U.S. Department of Energy (U.S. DOE) under the Advanced Battery Materials Research (BMR) Program. We thank Prof. Evan Reed, Prof. Stefano Ermon, Dr. Yiyang Li, Dr. Christoph Bauemer, Aditya Grover, Todor Markov, Dean Deng, Antonio Baclig, and Harry Thaman for insightful discussions.




Figure 1 | a, Discharge capacity for the first 1,000 cycles of LFP/graphite cells. The color of each curve is scaled based on the battery's cycle life, as is done throughout the manuscript. **b**, A detailed view of **a**, showing only the first 100 cycles. A clear ranking of cycle life has not emerged by cycle 100. **c**, Histogram of the state of health at cycle 100. The cell with the largest degradation (90%) is excluded to show the detail of the rest of the distribution. **d**, Cycle life as a function of discharge capacity at cycle 2. The correlation coefficient of capacity at cycle 2 and log cycle life is -0.06. **e**, Cycle life as a function of discharge capacity at cycle 100. The correlation coefficient of capacity at cycle 100 and log cycle life is 0.27. **f**, Cycle life as a

function of the slope of the discharge capacity curve for cycles 95 through 100. The correlation coefficient of this slope and log cycle life is 0.47.

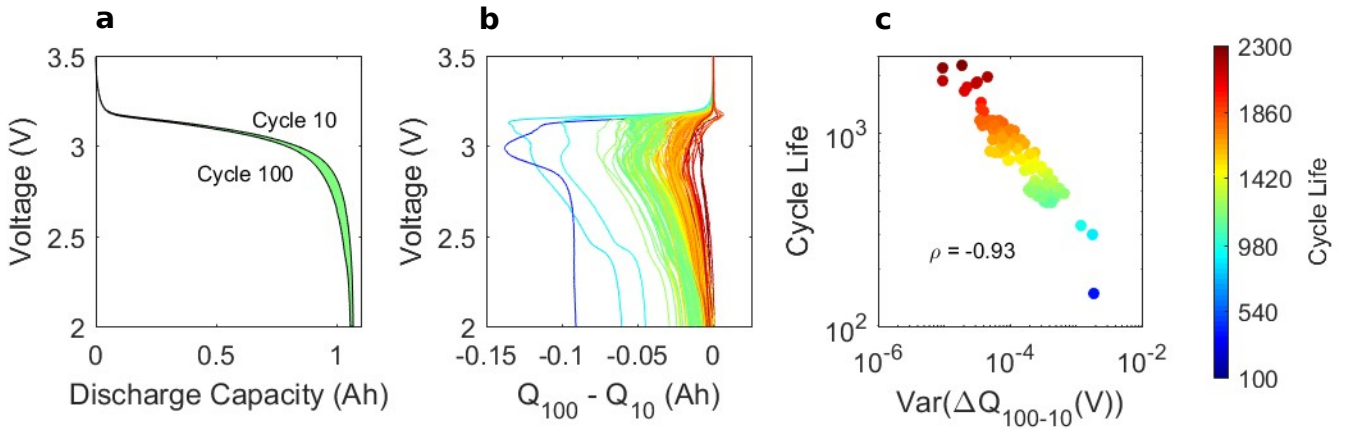


Figure 2 | **a**, Discharge capacity curves for 100th and 10th cycles for a representative cell. **b**, Difference of the discharge capacity curves as a function of voltage between the 100th and 10th cycles, $\Delta Q_{100-10}(V)$, for 124 cells. **c**, Cycle life plotted as a function of the variance of $\Delta Q_{100-10}(V)$ on a log-log axis, with a correlation coefficient of -0.93. In all plots, the colors are determined based on the final cycle lifetime. In **c**, the color is redundant with the y-axis.

Table 1 | Features considered for the various model implementations. The simplest model uses only the log variance of $\Delta Q_{100-10}(V)$ and does not consider model selection. More complex models are considered using only discharge information (first two sections) as well as additional measurements (all sections).

	Features	“Variance”	“Discharge”	“Full”
$\Delta Q_{100-10}(V)$ features	Minimum		✓	✓
	Mean			
	Variance	✓	✓	✓
	Skewness		✓	
	Kurtosis		✓	
	Value at 2V			
Discharge capacity fade curve features	Slope of the linear fit to the capacity fade curve, cycles 2 to 100			✓
	Intercept of the linear fit to capacity fade curve, cycles 2 to 100			✓
	Slope of the linear fit to the capacity fade curve, cycles 91 to 100			
	Intercept of the linear fit to capacity fade curve, cycles 91 to 100			
	Discharge capacity, cycle 2		✓	✓
	Difference between max discharge capacity and cycle 2		✓	
	discharge capacity, cycle 100			
Other features	Average charge time, first 5 cycles			✓
	Maximum temperature, cycles 2 to 100			
	Minimum temperature, cycles 2 to 100			
	Integral of temperature over time, cycles 2 to 100			✓
	Internal resistance, cycle 2			
	Minimum internal			✓

	resistance, cycles 2 to 100			
	Internal resistance, difference between cycle 100 and cycle 2			✓

Table 2 | Model metrics for the results shown in Figure 3. Train and test refer to the data used to learn the model and evaluate model performance, respectively. One battery in the test set reaches 80% SOH rapidly and does not match other observed patterns. Therefore, the parenthetical primary test results correspond to the exclusion of this battery.

	RMSE (cycles)			Mean Percent Error		
	Train	Primary Test	Secondary Test	Train	Primary Test	Secondary Test
“Variance” model	103	138 (138)	196	14.1%	14.7% (13.2%)	11.4%
“Discharge” model	76	91 (86)	173	9.8%	13.0% (10.1%)	8.6%
“Full” model	51	118 (100)	214	5.6%	14.1% (7.5%)	10.7%

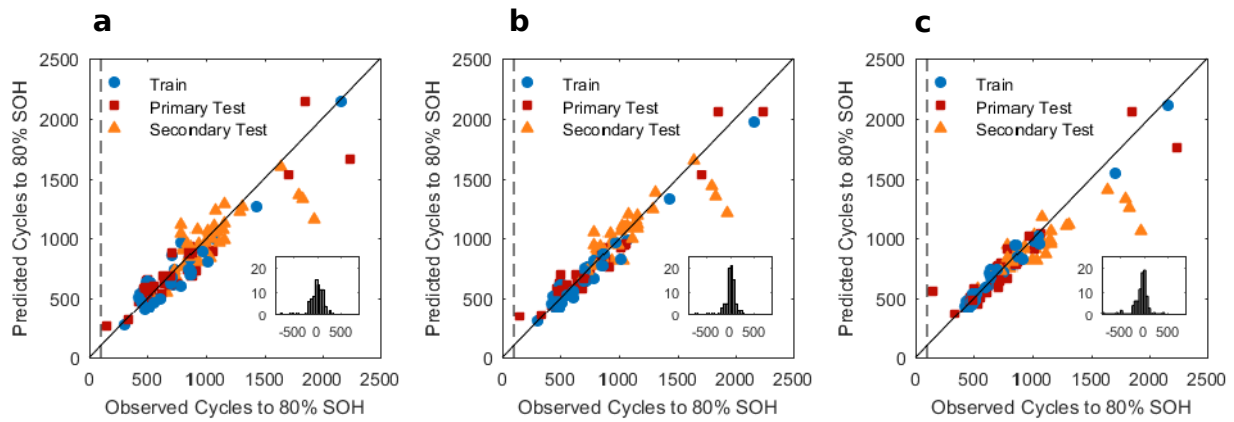


Figure 3 | Observed and predicted cycles to 80% SOH for several implementations of the feature-based model. The training data are used to learn the model structure and coefficient values. The testing data are used to assess generalizability of the model. We differentiate the primary test and secondary test datasets because the latter was generated after model development. The vertical dotted line indicates when the prediction is made in relation to the observed cycle life. The inset shows the histogram of residuals (predicted - observed) for the primary and secondary test data. **a, “variance” model using only the log variance of $\Delta Q_{100-10}(V)$. **b**, “discharge” model using six features based only on discharge cycle information, described in Table 1. **c**, “full” model using the nine features described in Table 1. Because some temperature probes lost contact during experimentation, four cells are excluded from the full model analysis.**

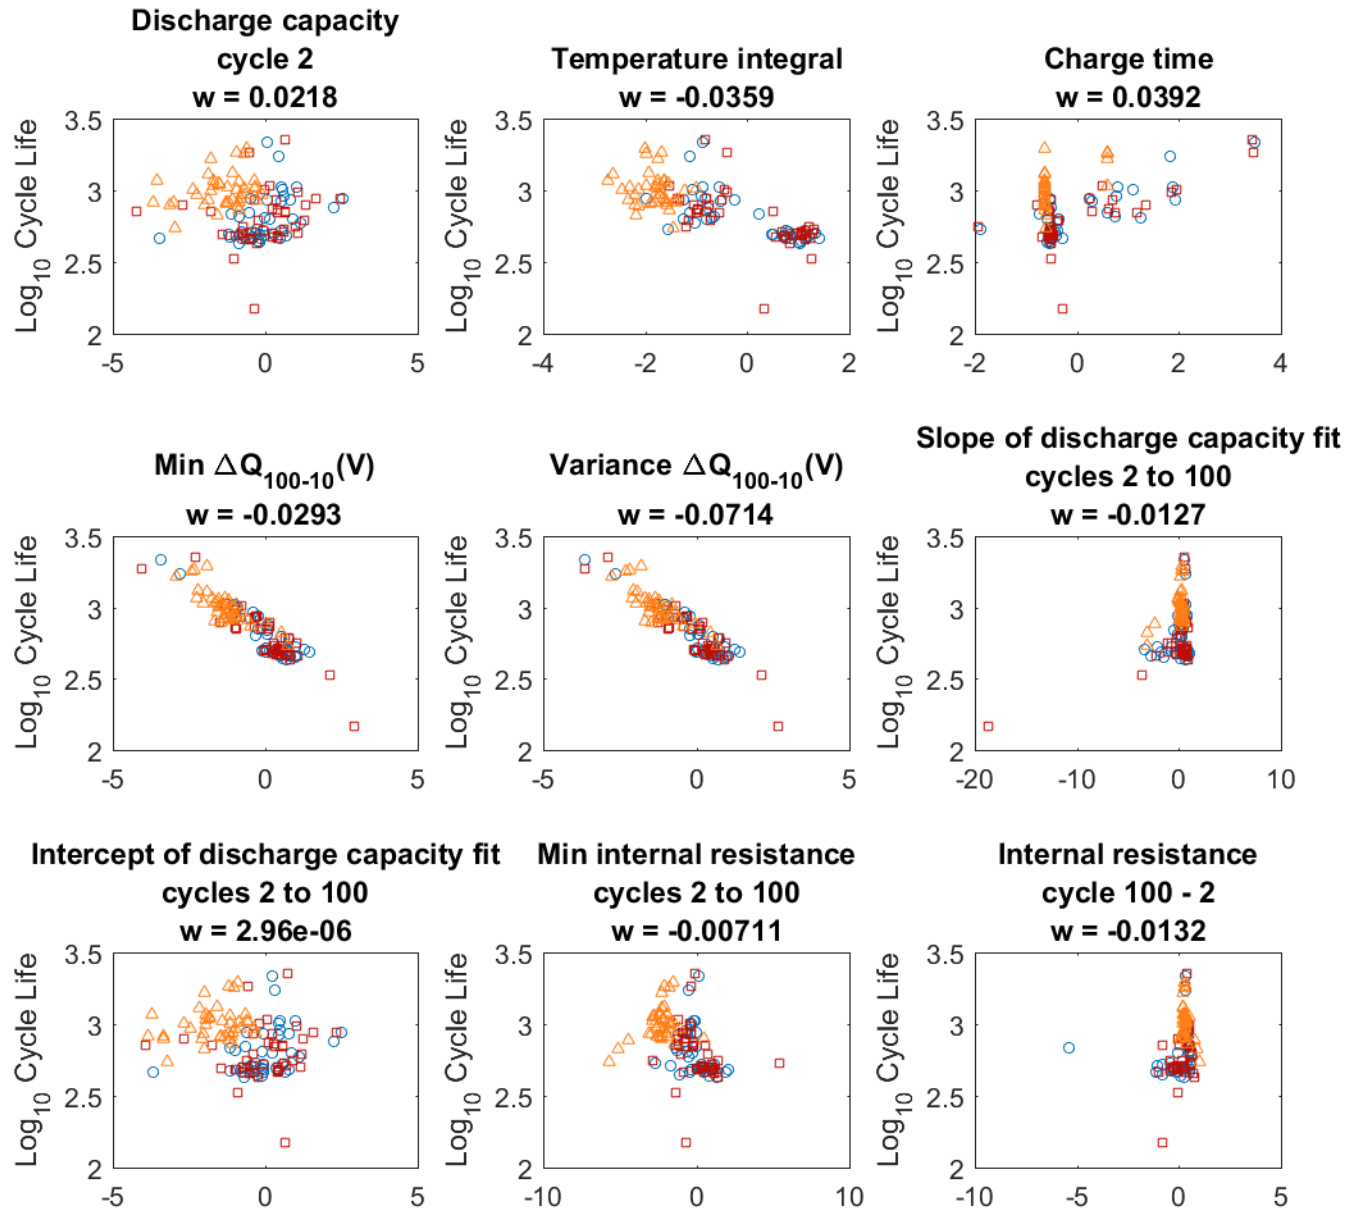


Figure 4 | Nine features used in the full model described in Table 1. The coefficient value for the feature in the linear model is in the title of plot. The train, primary test, and secondary test cells are represented by blue circles, red squares, and orange triangles, respectively. Each of the features has been z-scored based on the training data.

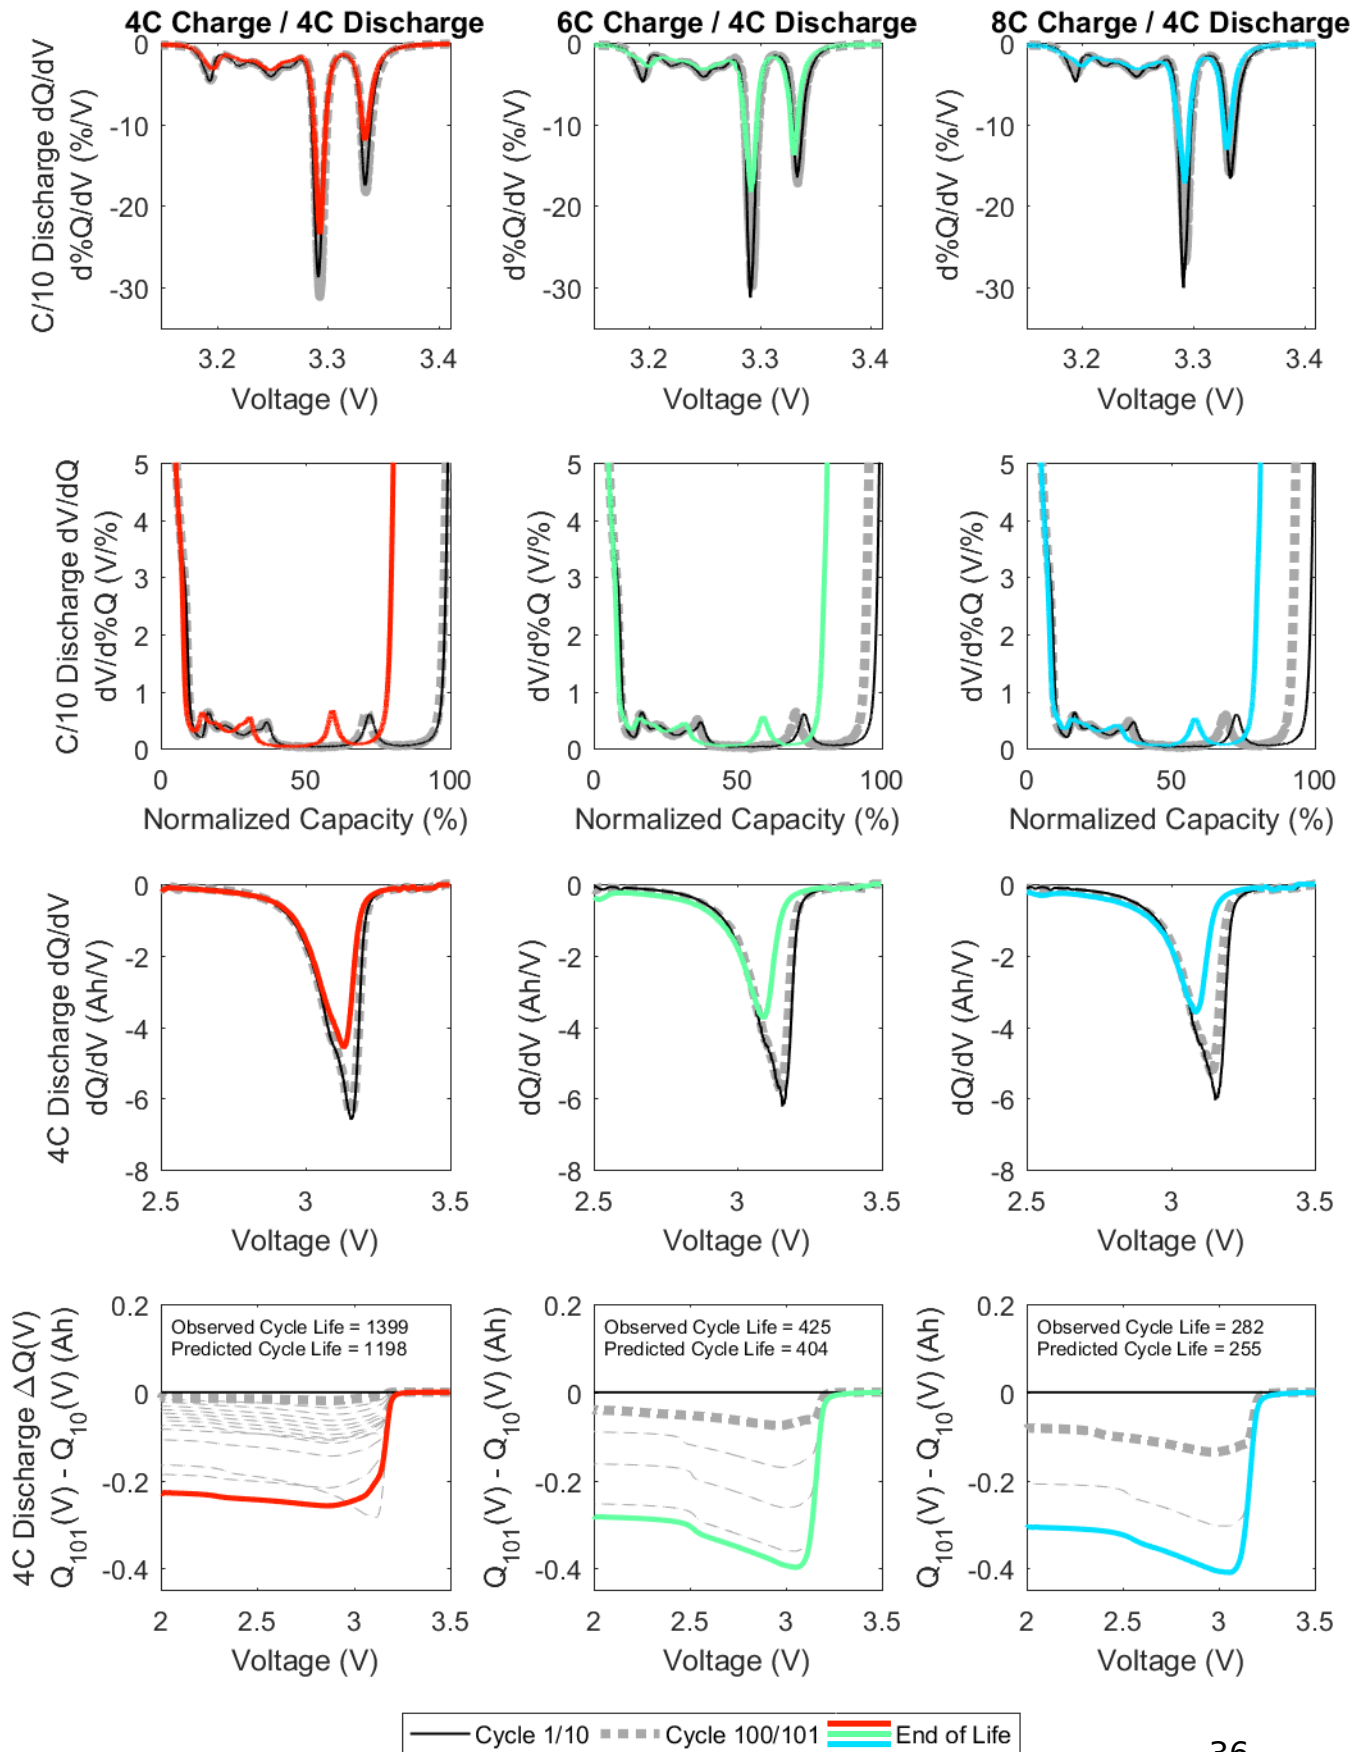


Figure 5 | Results of three cells that were tested with periodic slow diagnostic cycles. From top to bottom, the plots are dQ/dV using slow cycling, dV/dQ using slow cycling, dQ/dV using fast cycling, and $\Delta Q(V)$ using fast cycling. The solid black line is the first cycle (cycle 10 for fast cycling), the dotted grey line is cycle 101 or 100 (fast and slow, respectively), and the colored thick line is the end of life cycle (80% SOH). For $\Delta Q(V)$, a thin dotted grey line is added every 100 cycles. The patterns observed using slow cycling are consistent with LAM_{deNE} and LLI. The features are smeared during fast charging. The log variance $\Delta Q(V)$ model trained using the high-throughput dataset is able to predict lifetime within 15%.




Figure 6 | RMSE error, in cycles, for training and testing datasets using only the log variance of $\Delta Q_{i-j}(V)$, where the discharge cycles that are used in analysis are varied. These errors are averaged over 20 random partitions of the data into equal training and testing datasets. The errors are relatively flat after cycle 80. The increases in error around cycles $j = 55$ and $i = 70$ are due to temperature fluctuations of the chamber (see Figure S6 for information on experimental temperature).