

Impedance-based forecasting of battery performance amid uneven usage

Penelope Jones,^{1,2} Ulrich Stimming,³ and Alpha A. Lee^{1,*}

¹*Department of Physics, University of Cambridge, Cambridge, UK*

²*The Alan Turing Institute, London, UK*

³*Chemistry, School of Natural and Environmental Sciences,
Newcastle University, Newcastle upon Tyne, UK*

Abstract

Accurate forecasting of lithium-ion battery performance is important for easing consumer concerns about the safety and reliability of electric vehicles. Most research on battery health prognostics focuses on the R&D setting where cells are subjected to the same usage patterns, yet in practice there is great variability in use across cells and cycles, making forecasting much more challenging. Here, we address this challenge by combining electrochemical impedance spectroscopy (EIS), a non-invasive measurement of battery state, with probabilistic machine learning. We generated a dataset of 40 commercial lithium-ion coin cells cycled under multistage constant current charging/discharging, with currents randomly changed between cycles to emulate realistic use patterns. We show that future discharge capacities can be predicted with calibrated uncertainties, given the future cycling protocol and a single EIS measurement made just before charging, and without any knowledge of usage history. Our method is data-efficient, requiring just eight cells to achieve a test error of less than 10%, and robust to dataset shifts. Our model can forecast well into the future, attaining a test error of less than 10% when projecting 32 cycles ahead. Further, we find that model performance can be boosted by 25% by augmenting EIS with additional features derived from historical capacity-voltage curves. Our results suggest that battery health is better quantified by a multidimensional vector rather than a scalar State of Health, thus deriving informative electrochemical ‘biomarkers’ in tandem with machine learning is key to predictive battery management and control.

* Correspondence email address: aal44@cam.ac.uk

I. INTRODUCTION

Electrification of the transportation industry is now taking place at an increasingly rapid pace, enabling significant strides towards a carbon neutral future. Fundamental to this transition has been the development of the lithium-ion battery, which powers the majority of electric vehicles on the road today. Notwithstanding the environmental benefits of this transition, reliance on the lithium-ion battery poses novel challenges, with consumer concerns including range anxiety, fear of battery failure and charging time. Easing these concerns demands the ability to accurately forecast battery performance, and specifically when usage conditions are variable.

The key challenge is the heterogeneity of the battery. Each user uses their car differently, and even across a single battery pack not all cells are necessarily charged or discharged with identical current [1–3]. These differences mean that each cell’s internal state, including the extent of lithium plating or electrode cracking, can vary significantly both at an intra-pack and inter-pack level [4, 5].

To quantify the extent of degradation within cells, and to identify cells that have reached their ‘End of Life’, the scalar State of Health (SoH) metric is typically adopted, measured using previous cycle discharge capacity or internal resistance [6, 7]. The problem with this approach is that batteries with the same numerical SoH do not necessarily exhibit identical levels of each degradation process (for example, lithium plating or electrode cracking), yet the impact of future cell usage on the cell’s future performance and degradation pathway depends significantly on the type of degradation that has already occurred [8–10]. In order to forecast battery performance, we need a non-invasive way to acquire information about the cell state at a microscopic level.

Previous work primarily focused on forecasting future battery performance in the laboratory setting, where cells are charged and discharged in the same way over the entirety of their lifetimes, thus the impact of variable cell usage on future performance can be ignored (see Figure 1). These studies applied machine learning on features extracted from the charging or discharging curve to predict discharge capacity [11], remaining useful life [12], and abrupt capacity decays [13, 14]. Innovations in extracting features from charge/discharge curves [15] and machine learning approaches for modelling time-series data [16, 17] have enabled significant improvements in the accuracy of predictions. Going beyond charging and discharging curves, approaches such as electrochemical impedance spectroscopy (EIS) [18] and acoustic time-of-flight analysis [19, 20] have been used for degradation forecasting. These approaches provide a fuller description of battery state – for example, EIS captures the response of the cell over a broad frequency range, with different frequencies correlating to distinct physical, chemical and mechanical changes in the active material [21]. However, extrapolating the




Figure 1: **Schematic of our approach compared to previous work.** Previous approaches for degradation prediction focused on constant charging protocols (the blue/red curve denotes the charge/discharge phase), and used features from capacity-voltage curves as input. This necessitates knowledge of historic charging data. Our approach considers variable charging protocols (the shaded blue/red region denotes the range of currents that the charge/discharge protocols are drawn from), which is more realistic for EV settings. Further, we employ the electrochemical impedance spectrum measured just before charging as input, without any knowledge of historic data, and predict the impact of different future usage protocols on the discharge capacity.

models developed for laboratory setting to field data, where cells are cycled in vastly different ways over their lifetimes, has proved a major challenge [22].

In this work, we seek to identify whether there exists a sufficiently informative ‘biomarker’ of cell health that can be used to forecast future performance, amid uneven historical and future cell usage. Figure 1 provides an illustration of our approach, and how it differs from previous approaches.

We find that upon acquisition of an EIS spectrum just before charging, both next cycle and longer term cell capacity can be predicted with a test error of less than 10%. We observe that our model is data-efficient, requiring just eight cells to attain a test error of less than 10%. Crucially, our approach is robust to dataset shift, attaining a test error of less than 7% on a dataset with a different distribution of cycling patterns to the training set. This is vital for deployment in the field where driving patterns may be different from those used to train the model. Finally, we demonstrate that, if available, using additional features based on historical capacity-voltage data can serve to augment the state representation and reduce average test error by up to 25%.

Our work departs from the NASA randomised usage dataset [23], which randomly cycles cells for 50 cycles before measuring the next cycle discharging capacity after charging via a 'reference' protocol. Although several models for forecasting degradation under randomised conditions have been built based on this data [6, 11, 24], the effect of a single protocol on next cycle discharge capacity cannot be disentangled, and there is a need for a reference charge / discharge protocol every few cycles which is not realistic for the EV setting.

II. RESULTS

A. Data generation

We generate training data by subjecting 24 Powerstream LiR2032 coin cells (of nominal capacity $1C = 35\text{mAh}$) to a sequence of randomly selected charge and discharge currents at room temperature for 110-120 full charge/discharge cycles. Each cycle consists of an initial diagnosis of battery state, involving acquisition of the galvanostatic EIS spectrum, followed by usage, involving a charging and discharging stage. We collect impedance measurements at 57 frequencies uniformly distributed in the log domain in the range 0.02Hz-20kHz; a vector representation of battery state is formed from the real and imaginary parts of cell impedance at these frequencies. Charging consists of a two stage Constant Current (CC) protocol; currents are randomly selected in the ranges 70mA-140mA (2C-4C) and 35mA-105mA (1C-3C) in stages 1 and 2 respectively. If the safety threshold voltage of 4.3V is reached before the time limit then charging is stopped. During discharging, a single constant discharge current, randomly selected in the range 35mA-140mA (1C-4C), is applied, until the voltage drops to 3.0V. To test the model's robustness to domain shift, we additionally cycle 16 cells of the same chemistry and manufacturer under the same conditions as above, except now fixing the discharge current for all cells and cycles at 52.5mA (1.5C) instead of randomly changing the




Figure 2: **Our charge-discharge protocol.** We generate battery cycling data by subjecting cells to a sequence of random charge and discharge currents. We apply two stages of Constant Current (CC) charging for up to 15 minutes each, with currents drawn from the ranges 70mA-140mA (2C-4C) and 35mA-105mA (1C-3C) respectively (the blue shaded region). If the safety threshold voltage of 4.3V is reached before the time limit, then charging is stopped. During discharging, a single constant discharge current, randomly selected in the range 35mA-140mA (1C-4C), is applied (the red shaded region), until the voltage drops to 3.0V.

discharge current at each cycle. The space of protocols considered is illustrated in Figure 2.

B. Capacity forecasting using EIS.

We first consider the setting in which we want to predict the next cycle discharge capacity, for a cell whose usage history (including for example, cycle or calendar age, or historical capacity-voltage data) is completely unknown, if we apply a particular charging and discharging profile. We frame the problem as a regression task, and train a probabilistic machine learning model to learn the mapping $Q_n = f(\mathbf{s}_n, \mathbf{a}_n)$, with uncertainty estimates, where \mathbf{s}_n is the battery state at the start of the n th cycle (formed from the EIS spectrum acquired just before charging commences), \mathbf{a}_n is the n th cycle charge/discharge protocol (formed from the concatenation of the n th cycle charge and discharge currents), and Q_n is the discharge capacity measured at the end of the cycle. We use an ensemble of 10 XGBoost models, each with 500 estimators and a maximum depth of 100 [25]. To test model performance we use the median R^2 score and median percentage error. To obtain test metrics from 24 cells, we randomly leave two cells out, train on the remaining cells and repeat this process 12 times leaving different cells out each time.




Figure 3: **Predicting next cycle discharge capacity.** (a) Given knowledge of the battery state (as characterised by the EIS spectrum) and next cycle charge/discharge protocol, our model predicts next cycle discharge capacity with an error of 8.2%. Both state and action are needed to accurately forecast performance; using (b) state or (c) action alone is insufficient.

Figure 3a illustrates the accuracy of our model. Using both state and action as input, the next cycle discharge capacity is predicted with an average error of 8.2%. Importantly, both state and action are found to be necessary to predict future cell performance: if state (3b) or action (3c) alone are used as inputs, the test error approximately doubles to 20.7% and 15.4% respectively. This demonstrates the importance of both the cell’s internal health and the externally selected usage in determining realised cell performance.

For applications such as optimised charging, repurposing triaging and cell insurance calculations, it is important that a model of battery life trajectory can forecast not only the immediate next cycle discharge capacity, but also capacity several cycles into the future. With this in mind, we next investigate how the predictive accuracy of the model changes as we push the model to predict capacity further into the future. In each case, the input comprises the concatenation of the state representation at the start of the n th cycle, \mathbf{s}_n , with the ‘action’ vector $\mathbf{a}_{n\dots n+j}$ comprising all charging and discharging currents that will be applied between cycle n and cycle $n + j$.

Figure 4 shows how the coefficient of determination R^2 changes with j . As expected, the accuracy of the model generally decreases as the forecasting interval increases. However, the model still attains $R^2 = 0.75$ when projecting 40 cycles into the future.




Figure 4: **Multi-step forecasting.** Our model can also forecast longer term battery performance, as quantified by (a) % test error, and (b) R^2 value. Given the EIS spectrum and knowledge of the next protocols that will be applied to the cell, the discharge capacity is predicted with a test error of less than 10% up to 32 cycles in advance.

C. Model robustness.

We next test the robustness of our method by investigating data efficiency and model generalisability. To test data efficiency, we measure how performance changes as the number of cells used to train the model increases. As seen in Figure 5, there is a marked reduction in test error from 23.8% to 8.2% as the number of cells is increased from two to 22. Nevertheless, the model is demonstrably data-efficient, with just eight cells needed to obtain a test error of less than 10%.

An important test of model generalisability is to study model accuracy when the domain distribution changes, i.e. when the model is being deployed in settings that are different to the training data. This is important for deployment in the field as the approach needs to be robust to driving patterns that might be different to the training data. We test model robustness by cycling an additional 16 cells of the same chemistry, but now adjusting the cycling protocol by fixing the discharge current to 1.5C for each cell throughout its life. We use a model trained using only cells that were subjected to random discharge currents over their lifetime, to predict next-cycle discharge capacity of cells subjected to fixed discharging. To illustrate the difference in training and test datasets, the distribution of discharge capacities is shown for each in Figure 6a.

The predictive accuracy of the model on the fixed discharge dataset is illustrated in Figure 6b. Promisingly, the model attains a test error of just 6.3% on this domain-shifted dataset, which corresponds to $R^2 = 0.76$.




Figure 5: **Data efficiency.** The model performance, as quantified by (a) % test error and (b) R^2 , improves as the number of cells used to train increases. The model is data-efficient, achieving a test error of less than 10% with just 8 cells in the training set.

Our model also outputs predictive uncertainty, i.e. how certain is the model about its predictions. It is especially important in the domain-shifted setting that the model ‘knows what it does not know’ and estimates high predictive uncertainty about data points that it is likely to obtain a high error on. We can test the model’s ability to estimate its uncertainty by observing how the average test error changes as the number of data points is reduced to include only the data points that the model is most confident about. If a model can successfully estimate its level of certainty, the average test error should reduce as the proportion of data is reduced to include only the most confidently predicted points. Figure 6c shows a 32% reduction in RMSE as the proportion of data is reduced from 100% to the most confident 25%, demonstrating that our model has learnt which predictions it should be confident about.

D. Comparison of state representations

Having demonstrated the ability of the EIS spectrum to capture battery state, we now consider whether this state representation could be further improved by augmenting the EIS spectrum with other physics-based features. Specifically, we consider the following additional features:

- Capacity-voltage discharge curve features (CVF): Following Severson et al. [12], we form a state representation at the start of cycle n by extracting features from the capacity-voltage discharge curve after cycle $n - 1$. We fit each curve to a spline function, linearly interpolating to measure




Figure 6: **Robustness to domain shift.** (a) The distribution of discharge capacity is different for cells cycled under variable discharge rates (blue) compared to a fixed discharge rate (red); the overlap region of the two distributions appears darker in colour. (b) Our model, trained on the variable discharge rate cells, accurately predicts the discharge capacities of cells cycled under a fixed discharge rate. The colour of the plotted points denotes predicted uncertainty (see colour bar). (c) Our predicted uncertainty tracks actual model error.

capacity at 1000 evenly spaced voltages from V_{\min} to V_{\max} . This 1000-dimensional capacity vector \mathbf{Q}_{n-1} is normalised by subtracting the equivalent vector from cycle 0, \mathbf{Q}_0 . The following features are then used as inputs: V_{\max} , V_{\min} , $\log(\text{var}(\mathbf{Q}_{n-1} - \mathbf{Q}_0))$, $\log(\text{IQR}(\mathbf{Q}_{n-1} - \mathbf{Q}_0))$. Additionally, we fit the capacity to a sigmoid $Q(\tilde{V}) = \frac{p_0}{1.0 + \exp(p_1(\tilde{V} - p_2))}$ where \tilde{V} is the normalised voltage and use the parameters p_0 , p_1 , p_2 as features.

- Capacity throughput (CT) since cycling commenced, as defined by the sum of cell charge and discharge capacities from cycles 0 to $n - 1$.
- Previous cycle discharge capacity Q_{n-1} .
- State of Health (SoH), as defined by Q_{n-1}/Q_0 .

We note that in contrast to EIS features, formation of a state representation using the aforementioned features demands access to historical current-voltage data, over at least the entirety of the previous discharge and for some features, over the entire cell lifetime.

Table I shows how the state representation impacts test error and model goodness of fit. In all cases, the model is trained to predict the next cycle discharge capacity, given the next cycle protocol and the chosen state representation. We see that whilst the EIS spectrum is the ‘optimal’

Input	R²	Test error (%)
Protocol only	0.32	15.4
EIS only	0.05	20.7
EIS + Protocol	0.83	8.2
SoH + Protocol	0.61	12.4
CVF + Protocol	0.75	8.8
CT + Protocol	0.52	13.3
Q_{n-1} + Protocol	0.56	12.1
EIS + CVF + CT	0.91	6.2
+ Q_{n-1} + Protocol		

Table I: **Comparison of state representations.** We consider whether augmenting the EIS-based state representation with additional features can enhance model performance. The State of Health (SoH), as defined by the ratio between the previous discharge capacity and the initial discharge capacity, does not fully capture the battery’s current health; using this scalar state representation instead of the multidimensional EIS representation increases the test error by 51%. The most informative representation of battery state is obtained by adding previous capacity-voltage discharge curve features (CVF), capacity throughput (CT), and previous discharge capacity Q_{n-1} .

feature (giving the best model performance if only one source of features is used), the best model performance is attained by combining all of the above features to form the state representation. In this case the average test error is just 6.2%. In this way, we see that historical capacity-voltage curve data, if available, can act as a rich source of supplementary information about the cell’s current health, even if the historical cycling conditions are not fixed or known. Crucially, information from capacity-voltage curve data is complementary to EIS - i.e. combining EIS with these features leads to a significant increase in accuracy. This is perhaps unsurprising, as EIS probes the impedance of the single ‘static’ cell discharged state (with high information content per instant state), whilst capacity-voltage curves probe how the cell state evolves continuously over the path from charged to discharged (with low information content per instant state).

III. DISCUSSIONS AND CONCLUSION

In this paper, we showed that the electrochemical impedance spectrum accurately characterises the internal state of a cell, and a machine learning model can be trained to accurately forecast both immediate and longer term cell performance with predictive uncertainty, even amid uneven and unknown historical cell usage. Our method is data-efficient, achieving a next-cycle test error of 9.9% with training data from just eight cells, and is robust to shifts in dataset distributions. Finally, we find that there is scope to boost model performance by 25% if historical cycling data is available; such data can be used to derive features that augment the cell state representation.

Our approach differentiates from the prior art in two important ways: First, we employ an information-rich electrical signal – EIS – which captures the response of the cell across different timescales without any knowledge of the cycling history. This is in contrast to most existing methods which employ features from the charging-discharging curve – a significantly more coarse-grained signal – as input to machine learning models. Our results suggests significant improvements in battery management systems abound by incorporating circuitries that measure electrochemical impedance.

Second, we focus on uneven cycling, where the charging and discharging rates vary from cycle to cycle. This departs from previous studies on machine learning for battery degradation which focused on constant charge/discharge conditions, which are typical in battery testing. Our results problematise the concept of a single scalar ‘State of Health’, as the state of the battery is dependent on the extent of the myriad different degradation mechanisms, which in turn depends on the sequence of historic charge/discharge protocols. Rather, we suggest that a cell can be described by a multidimensional state vector, captured using informative high-dimensional measurements like EIS, and a machine learning approach can be used to predict future capacities given the state vector and future charge/discharge protocols.

We note that the general framework that we have laid out for predicting future battery performance given current cell state and future actions has scope to be applied in a broad range of battery diagnostic and control settings. For example, predicting the effect of a proposed charging protocol on next cycle discharge capacity as well as long term degradation is important for optimising rapid charging applications [26], where a balance must be achieved between charging time and rate of cell degradation [27].

ACKNOWLEDGEMENTS

PKJ and AAL acknowledge the support of the Winton Programme for the Physics of Sustainability. PKJ acknowledges support from the Ernest Oppenheimer Fund and the Alan Turing Institute (EPSRC EP/W001381/1). AAL acknowledges support from the Royal Society. The authors thank Dr Yunwei Zhang for helpful discussions.

DATA AND CODE AVAILABILITY

The data and code required to reproduce this manuscript are available at <https://doi.org/10.5281/zenodo.5704796> and <https://github.com/PenelopeJones/battery-forecasting> respectively.

-
- [1] R. Gogoana, M. B. Pinson, M. Z. Bazant, and S. E. Sarma, *Journal of Power Sources* **252**, 8 (2014).
 - [2] M. J. Brand, M. H. Hofmann, M. Steinhardt, S. F. Schuster, and A. Jossen, *Journal of Power Sources* **334**, 202 (2016).
 - [3] T. Bruen and J. Marco, *Journal of Power Sources* **310**, 91 (2016).
 - [4] F. An, L. Chen, J. Huang, J. Zhang, and P. Li, **6**, 35051 (2021).
 - [5] M. Schindler, J. Sturm, S. Ludwig, J. Schmitt, and A. Jossen, **8**, 100102 (2021).
 - [6] R. R. Richardson, M. A. Osborne, and D. A. Howey, *Journal of Power Sources* **357**, 209 (2017).
 - [7] X. Hu, L. Xu, X. Lin, and M. Pecht, *Joule* **4**, 310 (2020).
 - [8] M. Broussely, P. Biensan, F. Bonhomme, P. Blanchard, S. Herreyre, K. Nechev, and R. J. Staniewicz, *Journal of Power Sources Selected Papers Presented at the 12th International Meeting on Lithium Batteries*, **146**, 90 (2005).
 - [9] U. R. Koleti, T. Q. Dinh, and J. Marco, *Journal of Power Sources* **451**, 227798 (2020).
 - [10] U. R. Koleti, T. N. M. Bui, T. Q. Dinh, and J. Marco, *Journal of Energy Storage* **39**, 102573 (2021).
 - [11] D. Yang, X. Zhang, R. Pan, Y. Wang, and Z. Chen, *Journal of Power Sources* **384**, 387 (2018).
 - [12] K. A. Severson, P. M. Attia, N. Jin, N. Perkins, B. Jiang, Z. Yang, M. H. Chen, M. Aykol, P. K. Herring, D. Fraggedakis, M. Z. Bazant, S. J. Harris, W. C. Chueh, and R. D. Braatz, *Nature Energy* **4**, 383 (2019).

- [13] P. Fermín-Cueto, E. McTurk, M. Allerhand, E. Medina-Lopez, M. F. Anjos, J. Sylvester, and G. Dos Reis, *Energy and AI* **1**, 100006 (2020).
- [14] W. Li, N. Sengupta, P. Dechent, D. Howey, A. Annaswamy, and D. U. Sauer, *Journal of Power Sources*, 230024 (2021).
- [15] D. Roman, S. Saxena, V. Robu, M. Pecht, and D. Flynn, *Nature Machine Intelligence* **3**, 447 (2021).
- [16] R. Zhao, P. J. Kollmeyer, R. D. Lorenz, and T. M. Jahns, in *2017 IEEE Energy Conversion Congress and Exposition* (2017) pp. 5234–5241.
- [17] W. Li, N. Sengupta, P. Dechent, D. Howey, A. Annaswamy, and D. U. Sauer, *Journal of Power Sources* **482**, 228863 (2021).
- [18] Y. Zhang, Q. Tang, Y. Zhang, J. Wang, U. Stimming, and A. A. Lee, *Nature communications* **11**, 1 (2020).
- [19] G. Davies, K. W. Knehr, B. Van Tassell, T. Hodson, S. Biswas, A. G. Hsieh, and D. A. Steingart, *Journal of The Electrochemical Society* **164**, A2746 (2017).
- [20] C. Bommier, W. Chang, Y. Lu, J. Yeung, G. Davies, R. Mohr, M. Williams, and D. Steingart, *Cell Reports Physical Science* **1**, 100035 (2020).
- [21] N. Meddings, M. Heinrich, F. Overney, J.-S. Lee, V. Ruiz, E. Napolitano, S. Seitz, G. Hinds, R. Raccichini, M. Gaberšček, and J. Park, *Journal of Power Sources* **480**, 228742 (2020).
- [22] V. Sulzer, P. Mohtat, A. Aitio, S. Lee, Y. T. Yeh, F. Steinbacher, M. U. Khan, J. W. Lee, J. B. Siegel, A. G. Stefanopoulou, and D. A. Howey, *Joule* (2021).
- [23] K. C. S. . D. M. Bole, B., *Annual Conference of the Prognostics and Health Management Society*. (2014).
- [24] R. R. Richardson, C. R. Birkl, M. A. Osborne, and D. A. Howey, *IEEE Transactions on Industrial Informatics* **15**, 127 (2019).
- [25] T. Chen and C. Guestrin, *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2016).
- [26] P. M. Attia, A. Grover, N. Jin, K. A. Severson, T. M. Markov, Y.-H. Liao, M. H. Chen, B. Cheong, N. Perkins, Z. Yang, P. K. Herring, M. Aykol, S. J. Harris, R. D. Braatz, S. Ermon, and W. C. Chueh, *Nature* **578**, 397–402 (2020).
- [27] P. Keil and A. Jossen, *Journal of Energy Storage* **6**, 125 (2016).