



Titre: Identification and machine learning prediction of knee-point and
Title: knee-onset in capacity degradation curves of lithium-ion cells

Auteurs: Paula Fermín-Cueto, Euan McTurk, Michael Allerhand, Encarni
Authors: Medina-Lopez, Miguel F. Anjos, Joel Sylvester, & Gonçalo dos Reis

Date: 2020

Type: Article de revue / Article

Référence: Fermín-Cueto, P., McTurk, E., Allerhand, M., Medina-Lopez, E., Anjos, M. F.,
Citation: Sylvester, J., & dos Reis, G. (2020). Identification and machine learning prediction
of knee-point and knee-onset in capacity degradation curves of lithium-ion cells.
Energy and AI, 1, 100006 (10 pages). <https://doi.org/10.1016/j.egyai.2020.100006>

 **Document en libre accès dans PolyPublie**
Open Access document in PolyPublie

URL de PolyPublie:
PolyPublie URL: <https://publications.polymtl.ca/10631/>

Version: Version officielle de l'éditeur / Published version
Révisé par les pairs / Refereed

Conditions d'utilisation: Creative Commons Attribution-Utilisation non commerciale-Pas
Terms of Use: d'oeuvre dérivée 4.0 International / Creative Commons Attribution-
NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND)


 **Document publié chez l'éditeur officiel**
Document issued by the official publisher

Titre de la revue: Energy and AI (vol. 1)
Journal Title:

Maison d'édition: Science Direct
Publisher:

URL officiel: <https://doi.org/10.1016/j.egyai.2020.100006>
Official URL:

Mention légale: © 2020 The Author(s). Published by Elsevier Ltd. This is an open access article under
Legal notice: the CC BY-NC-ND license. (<http://creativecommons.org/licenses/by-nc-nd/4.0/>)



Identification and machine learning prediction of knee-point and knee-onset in capacity degradation curves of lithium-ion cells

Paula Fermín-Cueto^a, Euan McTurk^b, Michael Allerhand^a, Encarni Medina-Lopez^c, Miguel F. Anjos^a, Joel Sylvester^b, Gonçalo dos Reis^{a,d,*}

^a School of Mathematics, University of Edinburgh, The King's Buildings, Edinburgh, UK

^b Dukosi Ltd., Quantum Court, Research Avenue South, Heriot Watt University Research Park, Edinburgh, UK

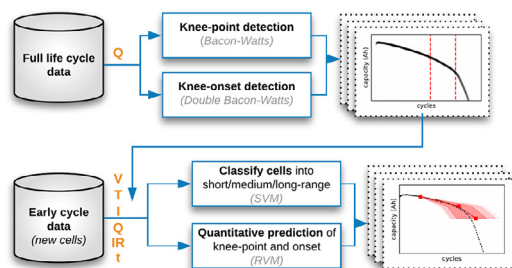
^c Institute for Infrastructure and Environment, School of Engineering, University of Edinburgh, The King's Buildings, Edinburgh, UK

^d Centro de Matemática e Aplicações (CMA), FCT, UNL, Portugal

HIGHLIGHTS

- A novel method for identification of the *Knee-point* in cell capacity degradation curves is developed.
- The new concept of *knee-onset*, for *Knee-point* early indication, is introduced along with robust identification algorithms.
- We show a strong linear relation between *knee-onset*, *knee-point* and *end-of-life*, where predicting one yields the others.
- Machine learning techniques for the early prediction of *knee-point* and *-onset* using only early-cycle data are used.
- The uncertainty of the predictions is methodologically quantified, providing reliable risk assessment for decision making.

GRAPHICAL ABSTRACT



ARTICLE INFO

Article history:

Received 14 February 2020

Received in revised form 11 April 2020

Accepted 11 April 2020

Available online 23 April 2020

Keywords:

Lithium-ion battery

Degradation

Knee-point

Knee-onset

Machine learning

Uncertainty quantification

ABSTRACT

High-performance batteries greatly benefit from accurate, early predictions of future capacity loss, to advance the management of the battery and sustain desirable application-specific performance characteristics for as long as possible. Li-ion cells exhibit a slow capacity degradation up to a knee-point, after which the degradation accelerates rapidly until the cell's End-of-Life. Using capacity degradation data, we propose a robust method to identify the knee-point within capacity fade curves. In a new approach to knee research, we propose the concept 'knee-onset', marking the beginning of the nonlinear degradation, and provide a simple and robust identification mechanism for it. We link cycle life, knee-point and knee-onset, where predicting/identifying one promptly reveals the others. On data featuring continuous high C-rate cycling (1C–8C), we show that, on average, the knee-point occurs at 95% capacity under these conditions and the knee-onset at 97.1% capacity, with knee and its onset on average 108 cycles apart.

After the critical identification step, we employ machine learning (ML) techniques for early prediction of the knee-point and knee-onset. Our models predict knee-point and knee-onset quantitatively with 9.4% error using only information from the first 50 cycles of the cells' life. Our models use the knee-point predictions to classify the cells' expected cycle lives as short, medium or long with 88–90% accuracy using only information from

* Corresponding author at: School of Mathematics, University of Edinburgh, The King's Buildings, Edinburgh, UK.

E-mail address: G.dosReis@ed.ac.uk (G. dos Reis).

the first 3–5 cycles. Our accuracy levels are on par with existing literature for End-of-Life prediction (requiring information from 100-cycles), nonetheless, we address the more complex problem of knee prediction.

All estimations are enriched with confidence/credibility metrics. The uncertainty regarding the ML model's estimations is quantified through prediction intervals. These yield risk-criteria insurers and manufacturers of energy storage applications can use for battery warranties. Our classification model provides a tool for cell manufacturers to speed up the validation of cell production techniques.

1. Introduction

The global market for lithium-ion cells is increasing with the uptake of electric vehicles and energy storage systems. It is common for manufacturers of electric vehicles and grid storage applications to provide a battery pack State of Health (SOH) warranty of eight years, which covers the cells dropping below 70–80% of their original capacity [1]. However, identical vehicles or energy storage systems may be subjected to very different duty cycles and ambient conditions, which affects the rate of degradation of the battery pack via its degradation mechanisms [1–4].

Li-ion cells exhibit a two-phase capacity fade behaviour: the capacity initially degrades at a low rate and then, starting at a certain onset point, the capacity goes through an accelerated degradation, displaying a so-called *knee* pattern, until the cell's End-of-Life (EoL). The IEEE Standard 485™-2010 [5] relates the “knee” with the transition to a stage of rapid decrease in capacity. The occurrence of the knee is a crucial factor of the cycle life of the cell. As such, the ability to detect and, more importantly, predict the occurrence of the knee and (if possible) its onset in each cell, depending on how it is cycled, is valuable to cell and battery manufacturers, who can adjust their specifications and warranties accordingly, and to the end user, who will have the option to adjust the duty cycles that the cell is subjected to in order to extend its useful life, and to schedule battery maintenance in a cost-effective manner.

Although this notion of the *knee* is well documented, [5–12], the literature on knee-point identification is sparse, with only a few attempts documented. Outside the battery domain, Satopaa et al. [10] define the knee as the point of maximum curvature and develop an algorithm based on this definition that can be applied to a wide range of systems. Diao et al. [6] define the knee as the intersection of two tangent lines to the capacity fade curve drawn at two significant points (an inflection point and the point of maximum slope changing ratio). The downside of this and the previous approach is that they rely on gradients and cannot handle raw data directly. To overcome this, Diao et al. [6] first characterise the capacity fade using the model they introduced in [13]. More recently, Zhang et al. [12] defined the knee-point as the intersection of two straight lines with different slopes and proposed an algorithm for online knee detection based on quantile regression (some parameter tuning is required). They fit a median regressor to the State of Health data and they define the knee as the first point at which the SOH data is outside a safety zone around the median regression line. They found that the knee-point in NMC cells appeared between 90–95% SOH. This method works well when applied to incoming data streams, but for the purposes of off-line identification in this work it is not as suitable, since the knee-points identified vary with the amount of training data used. Lastly, End-of-Life, knee-point and knee-onset are linked via linear regression methodologies, such that knowing one reveals the others.

Having access to a standard definition and methodology to determine the knees (-onset & -point) unlocks new opportunities for developing more effective battery prognosis systems, especially in the area of Remaining Useful Life (RUL) models. Accurate RUL predictions are of great importance for predictive maintenance, as they can reduce failure rates, safety issues and the maintenance costs of an application. Additionally, RUL predictions are advantageous for the effective administration of cell warranties, and can provide feedback to the end user to change the duty cycle of their application to prolong the cell's life. However, RUL prediction models documented in the literature define this remaining useful life based on the cell's cycle life, generally defined as the number of cycles before the cell reaches 70–80% of its nominal capacity. Shifting the focus to the knee (-onset & -point) enables an earlier

detection of accelerated health degradation, leading to more effective predictive maintenance; this is exactly what is proposed in the first part of this manuscript.

In the first part of this paper, a new method is proposed to identify the knee-point, and juxtapose it against other identification algorithms in the literature. The concept of knee-onset, and an algorithm to identify it, is also presented. To the best of our knowledge, this is the first time that the concept of *knee-onset*, or a method to identify it, appears in the literature. The methods proposed in this work are tested on a rapid cycling dataset, using real cells subjected to high C-rates ranging from 1C to 8C. The detection algorithms are simple and robust against noise without superimposing a degradation model. The identification methodologies used are derived from the Bacon–Watts model [14], are very general and of a wider interest than just battery degradation curves.

The second part of this paper focuses on early prediction. Classical approaches to prediction of State-of-Health focus on physical/chemistry models building up from micro- to macro-scale [15] (and references), simplified macro-scale models, or semi-empirical models [16,17]; see the review works [18,19]. The work in this paper falls within an alternative approach, fully focused on data discovery and its exploitation. Models for predicting Remaining Useful Life (RUL) or/and cycle life based on machine learning and statistical methods have gained increasing attention in the literature in the recent years: Support Vector Machine (SVM) and Relevance Vector Machine (RVM), are among the most widely used machine learning techniques to approach this problem [20–22]. Experiments with artificial neural networks (ANNs) have also been conducted [23,24]. These data-driven prediction models have the advantage of requiring little prior knowledge or assumptions on degradation mechanisms.

The early prediction problem aims to predict (qualitatively or quantitatively) a cell's RUL or cycle life using data from only the cell's early cycles, where significantly less degradation occurs. Making predictions within this framework is a challenge. To the best of our knowledge, Severson et al. [25] are the first to consider the early prediction problem, and their recent results for early cycle life prediction are very encouraging, although care is needed when interpreting their results (see Section 3 below). Drawing from domain knowledge to extract a small set of powerful predictors, they achieve very good results in terms of prediction error using early-cycle data. From a statistical point of view, they train a logistic regression model to classify cells into low-lifetime and high-lifetime using the first 5 cycles with an accuracy of ~ 95%, and a linear regression model to generate point estimates of the cycle life using data from the first 100 cycles with an accuracy of ~ 91%.

Access to relevant cycling data is one of the major challenges hindering the development of data-driven models. Generating a dataset covering a wide range of operating conditions and rare events is expensive and time-consuming. The NASA Ames Li-ion cell dataset¹, which contains data from 19 Li-ion cells, has been widely used in the literature [20,21,26–28]. Severson et al. [25] recently published one of the largest cell cycling datasets available, describing the degradation of 124 commercial Li-ion phosphate (LFP)/graphite cells. Other authors employed simulation or semi-empirical models to overcome the lack of relevant datasets: Finegan and Cooper [29] simulate data from a battery model and train a machine learning algorithm to predict the occurrence of an internal short circuit in a battery. D'Arpino et al. [30] base their analysis on a semi-empirical model of capacity and power fade combined with an electro-thermal model.

¹ Available at <http://ti.arc.nasa.gov/project/prognostic-data-repository>.

In the second part of this work focuses on data-driven modelling of complex dynamical systems. Two machine learning predictive models are presented with the ability to fully leverage the data from Severson et al. [25] to accurately predict/classify the knee-point and knee-onset of Li-ion cells using only early-cycle data. The manners in which this work differs in breadth and depth from [25] are described in Section 3. The machine learning models are benchmarked by providing measures of confidence for the models' predictions. This analysis "grounds" the machine learning agnostic predictions to tangible risk quantifiers for decision making.

Data description. The dataset generated by Severson et al. [25], denoted "A123 dataset," is used in this study. The dataset consists of 124 commercial lithium iron phosphate (LFP)/graphite cells cycled under fast-charging conditions until End-of-Life. For perspective, the cells underwent 4C discharge followed by varied fast-charging conditions ranging mainly from 3C to 8C with only a small number of cells being charged at 1C to 2C (less than 10%). The dataset contains in-cycle measurements of temperature, current, charge and discharge capacity, as well as per-cycle measurements of capacity, internal resistance and charge time. Data is recorded consistently from the second cycle until, at least, the cycle at which each cell reaches 80% of the nominal capacity. Detailed descriptions of the data and experiment can be found at <https://data.matr.io/1>.

2. Knee-point identification and the concept of knee-onset

There is consensus in the industry and the literature around the notion of the knee-point as the transition from a slow degradation rate to a rapid one. However, this transition is not abrupt and can be considered to take place over a number of cycles. For this reason, determining a single point in the capacity fade curve as the knee-point is a subjective task. Here, a new method to identify the knee-point is proposed and compared against existing methodologies. This is followed by the proposal of the novel concept of knee-onset and discussion of two different algorithms to identify it.

2.1. Knee-point identification

For the knee-point, the concept of Zhang et al. [12] is followed: the knee-point is the intersection of two lines with two different slopes that characterise the two stages in capacity degradation. Due to the noise in experimental cell capacity data, it is problematic to use gradients to determine the slope of these straight lines, so additional machinery is needed: Diao et al. [6] (slope changing ratio method) and Satopaa et al. [10] (maximum curvature method) as highlighted in the introduction.

The Bacon and Watts [14] model is used to identify the knee-point. The model is a straightforward and easy-to-use statistical model which does not rely on gradient methodologies and is robust against noise. Bacon and Watts proposed the model of Eq. (1), concretely, two straight line relationships, to the left and right of some unknown transition point x_1 , namely:

$$Y = \alpha_0 + \alpha_1(x - x_1) + \alpha_2(x - x_1) \tanh\{(x - x_1)/\gamma\} + Z, \quad (1)$$

where Z is a normally distributed and centred-in-zero random variable representing the residuals, α_1 and α_2 control the slopes of the intersecting lines, α_0 is a type of intercept of the leftmost segment (at $x = x_1$), and γ controls the abruptness of the transition. Parameters α_i and x_1 have been optimised (see Methods below) and γ fixed to a small value to obtain an abrupt transition around the change point x_1 . The knee-point is defined by x_1 . Fig. 1 shows the Bacon–Watts model applied to the degradation data of sample cell² b3c45 in the A123 dataset.




Fig. 1. Capacity degradation data for sample cell b3c45 in the A123 dataset and knee-point obtained applying the Bacon–Watts model. The average width of the 95% confidence interval (computed with the non-parametric bootstrap procedure) of the knee-points estimated with the Bacon–Watts model was 6.1 cycles.

In Fig 2(a), the knee-points identified in a sample of Li-ion cells covering a wide range of cycle lives are presented and compared with those proposed by Diao et al. [6] and Satopaa et al. [10]. These last two methods cannot be applied directly to capacity degradation data, so in both cases the capacity fade is first characterised using the model in Diao et al. [13], adjusted to overcome some issues encountered when applied to the A123 dataset. Specifically, the independent variable has been scaled to resolve numerical stability issues and the point of maximum slope changing ratio is forced to be greater than the first significant point. The knee-point detection method proposed by Zhang et al. [12] based on quantile regression has not been included in this analysis; although this method works well against an incoming data stream, it is not as convenient for off-line identification, as the knee-points identified vary with the amount of training data used.

The results displayed in Fig. 2(b) show that the maximum curvature and Bacon–Watts algorithms identify very similar knee-points. On average, Bacon–Watts knee-points occur 49 cycles after the maximum curvature knee-point. Both methods visually capture the midpoint of the transition from a slow degradation rate to a rapid one (Fig. 2(a)). The slope changing ratio method is less consistent and often detects the knee-points arbitrarily, either too early (e.g. b1c3, b1c0) or too late (e.g. b1c1). It is worth pointing out the robustness of the Bacon–Watts method, which is applied on raw experimental data, can cope with the large amount of noise present in cell b1c0, and generally works well even when the other two methods have failed.

2.2. Knee-onset concept and detection

The knee-point defined above can be seen as the midpoint of the knee. As such, the ability to detect the knee-point does not give the end user advanced warning of the transition to nonlinear capacity fade, and merely informs them that accelerated State-of-Health degradation is well underway. Therefore, the ability to identify/predict the onset of accelerated degradation is desirable from the perspective of the end user. For this reason, we propose a new concept, which we call the *knee-onset*. The *knee-onset* is defined as the point that marks the beginning of the accelerated degradation rate at which the capacity fade can no longer be approximated as a linear function. This definition is more conservative and might be preferred in applications where Remaining Useful Life (RUL) is important, as it enables an earlier warning (and response).

Two different algorithms for knee-onset detection are explored. As a *first method*, the knee-onset is defined as the intersection between the first segment of the Bacon–Watts model based on Eq. (1), with the capacity fade curve. This method captures the intuition that the knee-onset marks the end of the linear degradation phase. For the *second method*,

² Notation $bXcY$ refers to the cell in channel Y of batch X in the A123 dataset (see <https://data.matr.io/1>).

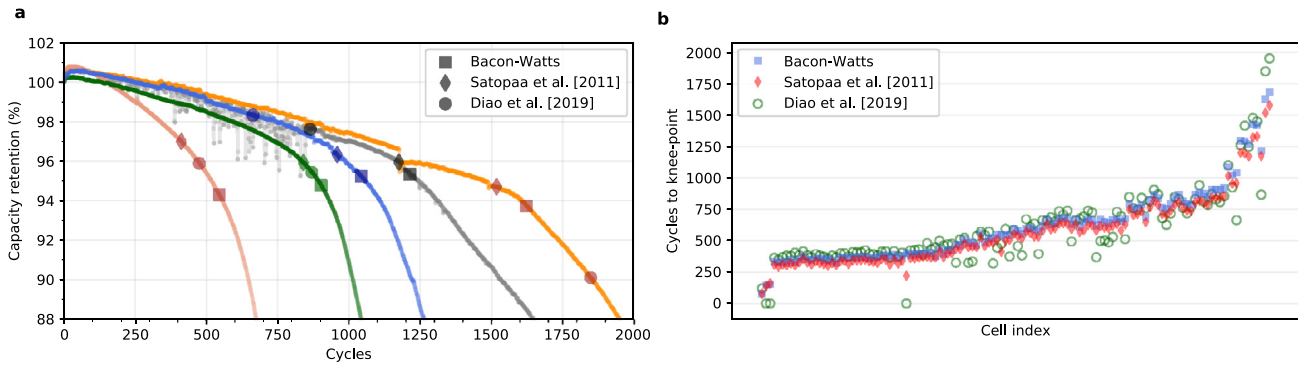


Fig. 2. (a) Comparison of knee-points obtained with Bacon–Watts, maximum curvature and slope changing ratio methods on a sample of cells from the A123 dataset (from left to right b2c47, b3c3, b1c3, b1c0, b1c1). See Supplementary Fig. 1 for results of a larger cell sample. (b) Comparison of knee-points in all cells in A123 dataset, sorted by their cycle life.

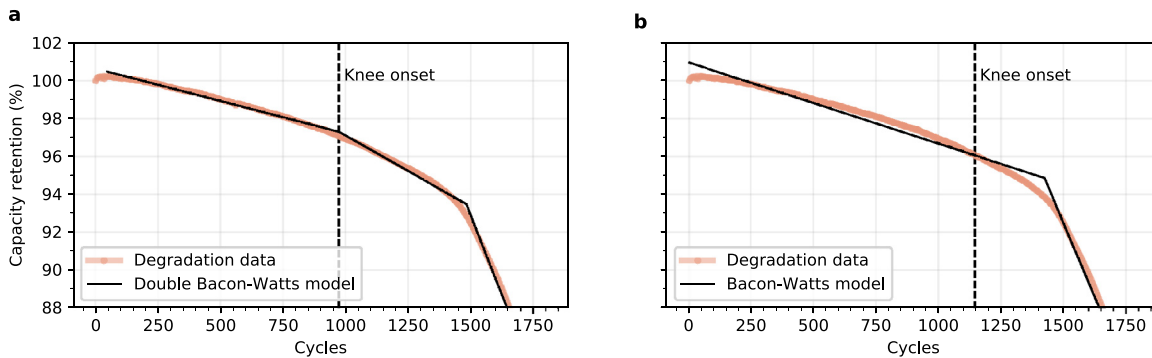


Fig. 3. Capacity degradation data for sample cell b3c45 in the A123 dataset and knee-onset corresponding to two different models: (a) double Bacon–Watts and (b) intersection of Bacon–Watts with capacity fade. The average width of the 95% confidence interval (computed with the non-parametric bootstrap procedure) of the knee-onsets estimated with the double Bacon–Watts was 13.8 cycles (the interval is very small that it is not depicted in the plots).

the Bacon–Watts model is adjusted to identify two transitions in the data instead of one, concretely:

$$Y = \alpha_0 + \alpha_1(x - x_0) + \alpha_2(x - x_0) \tanh\{(x - x_0)/\gamma\} + \alpha_3(x - x_2) \tanh\{(x - x_2)/\gamma\} + Z, \tag{2}$$

where the involved terms are in the same vein as those in Eq. (1) (with the parameters α_i and x_j to be estimated and γ fixed to a small value to obtain an abrupt transition around the change points x_0, x_2). This algorithm is denoted as the double Bacon–Watts, and the knee-onset defined as the change point x_0 in the fitted results. Both methods are illustrated in Fig. 3(a) and (b).

In the case of the knee-onset, both methods yield very similar results across the A123 dataset (Fig. 4(a) and (b)). However, the intersection method can fail if there is noise in the experimental data around the knee, as is the case for cell b1c1. For completeness, it is noted that, in the data analysis, the knee-point x_1 found via Eq. (1) always lies in between points x_0 and x_2 found via Eq. (2) (Supplementary Fig. 3).

These experimental tests lead to the selection of the standard Bacon–Watts and the double Bacon–Watts models for knee-point and knee-onset identification, respectively. The methods are robust, they provide visually acceptable results and they do not require a model of capacity degradation to be superimposed on the data. Computationally, the methods take less than 1 second to run. The double Bacon–Watts model for knee-onset identification provided a warning of accelerated degradation an average of 108 cycles before the single Bacon–Watts knee-point identification model, thus providing the end user with a valuable new metric for cell SOH that gives them more time to adjust the duty cycle of the cell or plan the maintenance of the battery pack.

2.3. Relations between knee-point, knee-onset, cycle life and capacity

As mentioned previously, for automotive applications, when a cell reaches 80% of its original capacity, the cell is considered as having reached its End-of-Life or cycle life. For the A123 dataset, it is found that, on average, the knee occurs at 95% capacity with its onset at 97.1% capacity. On average, the knee and its onset differ by 107.9 cycles (standard deviation of 66.4 cycles), and on average the knee-point and the cycle life differ by 187.8 cycles (standard deviation of 90.8 cycles).

Fig 5 (a) shows that, for the Li-ion cells in this dataset, the knee-point and knee-onset points determined with the Bacon–Watts models display a strong linear correlation with the cycle life.

This observation makes it possible to estimate the cycle life when the knee-point is known (or a prediction of it is available) and vice versa, using the linear regression model of Eq. (3):

$$Y = \beta_0 + \beta_1 X + Z, \tag{3}$$

where Y is the cycle life, X is either the knee-onset or the knee-point, Z is a normally distributed and centred-in-zero random variable representing the residuals, and β_0 and β_1 are the intercept and the slope, respectively, of this linear regression model. The coefficients' estimates and their confidence intervals are presented in Table 1.

The results of applying this model to extrapolate the cycle life when the knee-onset/point is available are displayed in Fig. 5(b). The cycle life can be forecast with an expected error of 31.4 cycles (MAPE = 4.0%) when the knee-point is known. Similarly, in the case of the knee-onset, the cycle life can be predicted with an average error of 49.3 cycles (6.1%).

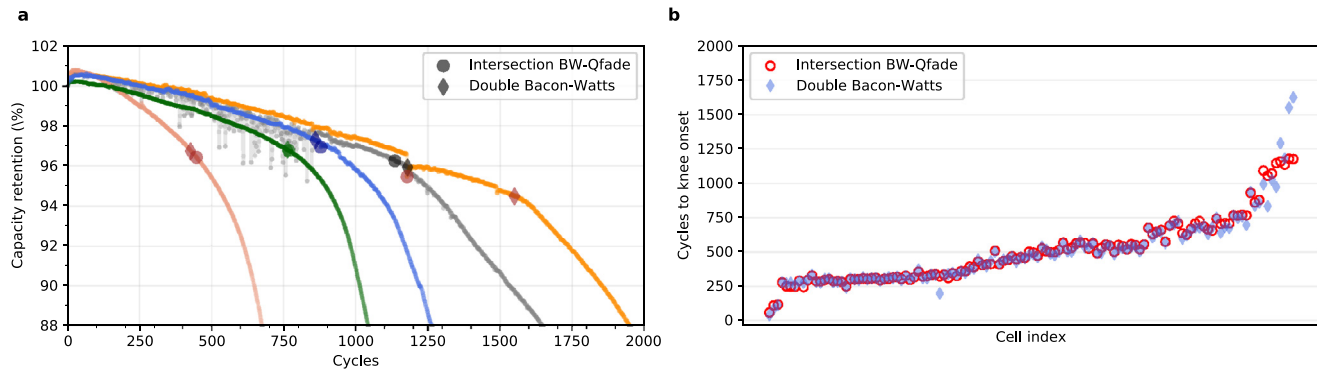


Fig. 4. (a) Comparison of knee-onsets obtained with double Bacon–Watts and the intersection between the standard Bacon–Watts and the capacity fade on the same sample of cells as in Fig. 2(a). See Supplementary Fig. 1 for results of a larger cell sample. (b) Comparison of knee-onsets in all cells in A123 dataset, sorted by their cycle life.

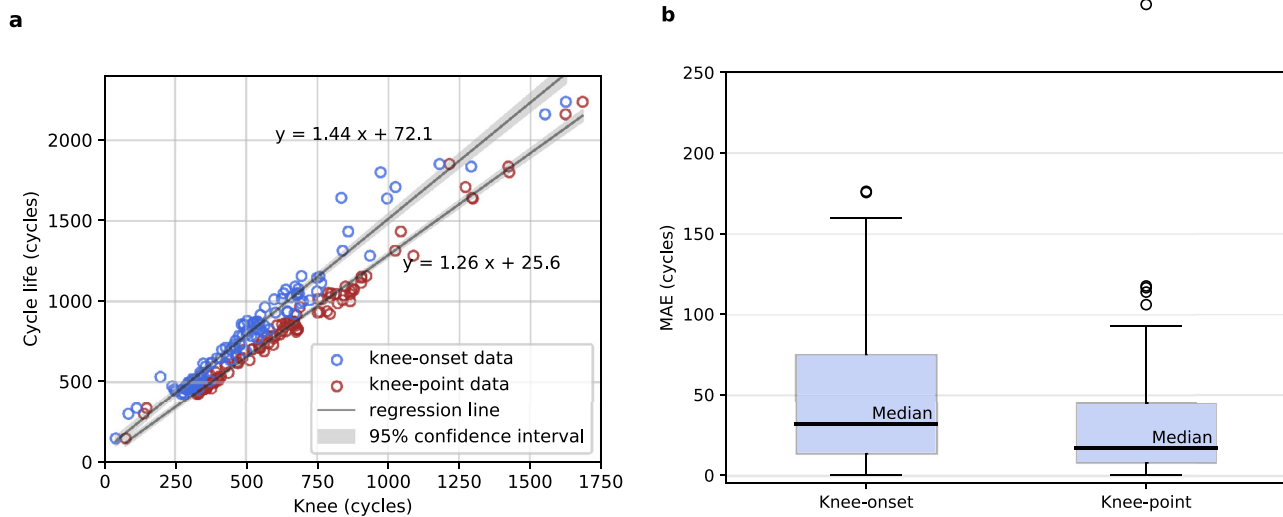


Fig. 5. (a) Linear regression models linking the knee-point to cycle life and and knee-onset to cycle life, both linear with 95% confidence intervals/bands around them. (b) Box-plot showing the distribution of prediction errors when the linear regression model is used to go from known knee-onset to prediction of cycle life (left) and from known knee-point to prediction of cycle life (right). The box-plots display the errors’ distributions through their quartiles (P25, median and P75); outliers are plotted individually.

Table 1

Coefficients of two linear regression models relating the knee-onset (a) and the knee-point (b) to the cycle life, respectively. The small p -values for coefficients β_1 , computed using the Wald test, allow the rejection of the null hypothesis that a linear relationship does not exist between the cycle life and the knee-point or knee-onset, with a significance level $\alpha = 0.05$. The confidence intervals capture the uncertainty around the estimated coefficients. The coefficient of determination, R^2 , of these linear regression models is 0.961 for the knee-onset and 0.983 for the knee-point, both very close to 1, showing a very strong agreement between the experimental data and the fitted values.

(a) Cycle life vs. knee-onset			
Coefficient	Estimate	p -value	Confidence interval ($\alpha = 0.05$)
Intercept (β_0)	72.13	3.4×10^{-6}	[42.80, 101.47]
Slope (β_1)	1.44	8.6×10^{-88}	[1.39, 1.49]
Model: cycle life = $72.13 + 1.44 \times \text{Knee-onset}$			
(b) Cycle life vs. knee-point			
Coefficient	Estimate	p -value	Confidence interval ($\alpha = 0.05$)
Intercept (β_0)	25.57	1.3×10^{-2}	[5.47, 45.67]
Slope (β_1)	1.26	5.2×10^{-110}	[1.23, 1.29]
Model: cycle life = $25.57 + 1.26 \times \text{Knee-point}$			

There is a further advantage to these linear relations, namely, as the majority of literature on the State-of-Health focuses on the End-of-Life for either identification or prediction, the results of this study mean that those methodologies can, in principle, also be used to predict knee-point and knee-onset. Unfortunately, such analysis has not yet been conducted, but the path is now open to do so.

Accurate estimates of when the capacity of cells or battery packs is expected to switch from decreasing linearly to decreasing at a much higher rate are valuable to end users, since this allows the pack cycling parameters, servicing and replacement schedules to be adjusted accordingly. This ensures that lifespan is maximised and cell(s) approaching End of Life are replaced at a convenient time, which minimises the downtime of the system. The knee-onset prediction also gives advanced warning of reduced State of Available Power and therefore reduced ability to meet performance criteria of the application, since decreasing capacity usually occurs in tandem with increased internal resistance and therefore reduced State of Available Power (SOAP). Cell manufacturers will find the ability to rapidly and reliably grade cells into cycle life categories useful since this ensures that customers will receive cells that are best suited to their applications, thus improving the grouping within battery packs of cells with the closest performance characteristics and further extending the lifespan of the pack.

3. Improved machine learning for early knee prediction

The identification of knee-onset and knee-point was a critical step towards enabling the development of predictive models for those quantities. Two problems are focused upon: *Problem 1: early classification* and *Problem 2: early quantitative prediction of lifespan*. Information on the early behaviour of the cell is used for both problems. For Problem 1, the data is used to classify a cell into one of three categories which indicate when the knee-midpoint is predicted to take place (short-life, medium-life or long-life) and, for Problem 2, the data is used to predict the cycle number at which the knee-point and knee-onset will occur. The closest related literature is the work of Severson et al. [25], which addresses variants of Problems 1 and 2 but for the End-of-Life problem (cycle in which the cell reaches 80% of its nominal capacity). The early prediction of knee-point and its onset is a more involved problem requiring deeper exploitation of the available data.

For completeness, the work of Severson et al. is juxtaposed with the work in this study. Severson et al. [25] used lasso logistic regression to address Problem 1 (classifying cells in two classes: low- and high-lifetime) and elastic net, a variation of linear regression, for Problem 2. Both problems targeted prediction/classification of End-of-Life and their choice of features or predictors was guided by domain expertise. *In this study*: (i) the much less tractable knee-point and knee-onset are targeted (instead of End-of-Life) to make predictions – the direct application of Severson et al.’s methodologies and feature choices fails to predict the knee-quantities satisfactorily; (ii) machine learning algorithms are employed with stronger predictive performance and which are less prone to overfitting to the training data while maintaining low complexity (*Support Vector Machine (SVM)* is used for Problem 1 and *Relevance Vector Machine (RVM)* is used for Problem 2); (iii) a feature generation pipeline is proposed that makes fewer assumptions on indicators of capacity degradation and, imperatively, can better leverage the information available in the data; (iv) lastly, and with a view to real-world market applications (insurance/warranties, predictive maintenance and manufacturing), uncertainty quantification is presented for the quality of the predictions from the machine learning algorithms (see next section) providing reliable risk assessment criteria for decision making.

There are two aspects of Severson et al.’s methodology that contributed to the high accuracy that they achieved in Problem 1 (92.7% and 97.5% in their primary and secondary test sets, respectively [25, Table 2]): (1) their definition of low-lifetime and high-lifetime cells and (2) the class imbalance in the secondary test dataset. In our view, choosing 550 cycles as the cycle life threshold is debatable: with such a threshold, 97.6% of the cells in the “low-lifetime” class correspond to cells in the second batch of the A123 dataset and 95.2% of the cells in the “high-lifetime” class correspond to batch 1. Cells in the second batch were subject to significantly higher C-rates (5 or 6C up to 80% SOC), which resulted in a substantially faster degradation compared to cells in batch 1. We therefore think that it is possible that their logistic regression model learned patterns that separate batch 1 from batch 2, instead of learning early indicators of future degradation. Regarding the secondary test set, which corresponds to the third batch of cells, it suffers from severe class imbalance, since only 1 of the 40 cells in this batch belongs to the low-lifetime class. As a consequence, it is hypothetically trivial to build a model that achieves 97.5% accuracy, by simply predicting “high-lifetime” for all cells. Therefore, using the accuracy metric to present the performance of the predictive model on the secondary test set is misleading.

Notwithstanding these observations, *Severson et al.’s contribution sets an inspirational baseline within data-driven prediction of End-of-Life*. Their machine learning models achieve an outstanding performance in the very hard task of predicting the long-term performance of a cell from early cycles, employing light interpretable models.

In the development of our methodology, different techniques were employed to obtain models with a good generalisation performance and

to present the robustness of the results. The *knee-point* threshold was set at 500 cycles (notice that this would correspond to a higher cycle life threshold), which ensures a higher presence of cells in batches 1 and 3 in the “short-lived” class (21%). A third class of very long-life cells (≥ 1100 cycles) is defined, which contains a balanced distribution of batches 1 and 3. Moreover, the accuracy metric is complemented with a confusion matrix to give additional insight on the model errors. All of the performance metrics, for both Problems 1 and 2, are presented with confidence intervals, a crucial good-practice step when working with small datasets. Lastly, by using leave-one-out validation to create training and test sets (see Methods section), it is ensured that all cells are used for testing, thus avoiding the class imbalance issue and maximising the data used in training. This makes our model more generalisable to unseen data.

Feature generation and extraction pipelines. Crucial to our two machine learning algorithms are the pipelines for generation and selection of descriptive variables. The variables available on a per-cycle basis (internal resistance, charge time and discharge capacity) are utilised in addition to different combinations of variables described in the in-cycle data: $Q(V)$, $dQ/dV(V)$, $dV/dQ(Q)$, $\Delta Q_{cycle-cycle_0}(V)$, $\Delta V_{cycle-cycle_0}(Q)$, $T(V)$ and $I(t)$, where $Q(V)$ is the discharge voltage curve as a function of voltage, $T(V)$ is the temperature curve during discharge as a function of voltage and $I(t)$ is the current as a function of time. $\Delta Q_{cycle-cycle_0}(V)$ variables, proposed by Severson et al. [25], represent the change in discharge voltage curves between a given cycle and a reference cycle $cycle_0$. The first cycle available in the data, cycle 2, is used for the classification problem and the cycle 10 for quantitative prediction. $\Delta V_{cycle-cycle_0}(Q)$ follow the same notation. The cycle-to-cycle evolution of these variables for a sample cell is illustrated in Supplementary Fig. 3. Time series analysis is used to extract predictors from these different variables in two steps: first, the high dimensional in-cycle data is aggregated at a cycle level by applying different summary metrics, to finally extract predictors from the cycle-to-cycle evolution of the resulting variables, using similar time series metrics – Supplementary Fig. 4 illustrates this process. Lastly, different feature selection and transformation techniques are applied to select around 100 predictors to train the machine learning models (Supplementary Figs. 6 and 7).

Problem 1: Early classification of cells by knee-point occurrence

The classification problem is about predictive categorisation of cells by the cycle-range in which the knee-point takes place, concretely, into short-range (knee-point in < 500 cycles), medium-range (500–1100 cycles) or long-range (> 1100 cycles), using only information from a cell’s early life. This information is relevant for predictive maintenance of batteries in grid-storage applications and for manufacturers. The latter can use the output of such models to grade cells during manufacturing after running as few as *three* charge-discharge cycles (see below), which in turn ensures that customers will receive cells that will meet and sustain the performance criteria of their applications.

The machine learning algorithm selected for this task is a *Support Vector Machine (SVM)* [31], a powerful yet simple classifier. The classifier, trained with data from just the first 3 cycles, categorises cells into ‘short’, ‘medium’ and ‘long-range’ with an accuracy of 88%. The confusion matrix in Fig. 6(a) provides additional insight on the type of the errors made by the model. The most common errors are those that underestimate the lifespan of the cell: there may be long-range cells classified as medium-range and medium classified as short, and there are no instances of short-range cells classified as long-range or vice versa. This is preferable from a business perspective, as it means that the likelihood of a cell failing to meet its categorised cycle life, and thus being subject to a warranty claim, is negligible.

For a comparative study of how much information is sufficient for the model to produce quality predictions, the model is trained using information from the first 3 up to the first 40 cycles (16 SVM models are built for this study) and compare the impact of the amount of training




Fig. 6. (a) Confusion matrix of predictions made at 3 cycles for the classification model. Each row of the matrix represents the percentage of each group of cells that were classified as ‘short’, ‘medium’ or ‘long-range’; overall accuracy of 88%. (b) Classification accuracy for increasing amounts of cycling data used to make predictions/train the algorithms. The Wilson score interval was used to compute the binomial proportion confidence interval with $\alpha = 0.1$. c, Mean error (MAE) and mean absolute percentage error (MAPE) of regression models to predict knee-point and knee-onset and impact of the number of cycles used to make predictions. The confidence intervals were implemented using the non-parametric bootstrap procedure, with significance level $\alpha = 0.1$.

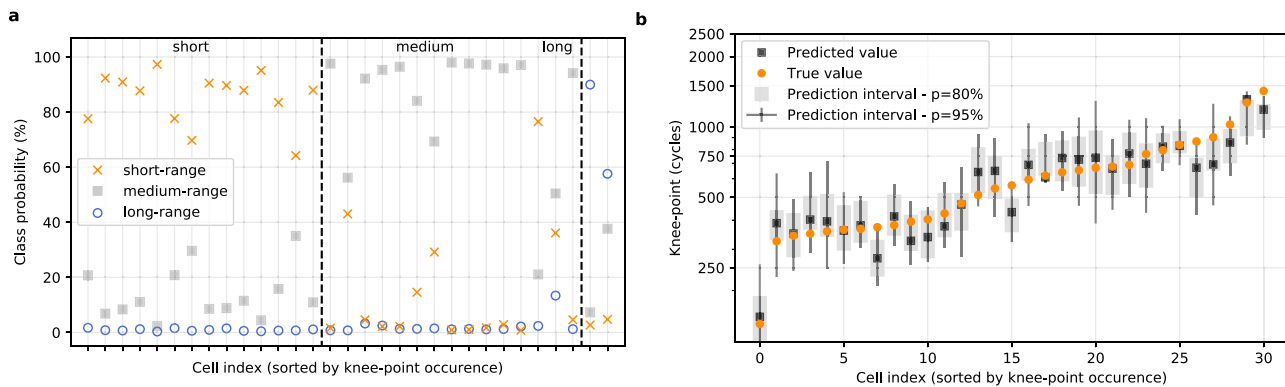


Fig. 7. (a) Probability of the cell belonging to each class (based on the knee-point) for a sample of 31 cells in the A123 dataset, using the first 3 cycles to make predictions. (b) Knee-point prediction intervals for 80% and 95% probability of including the true values, using conformal prediction, for a sample of 31 cells from the A123 dataset, using the first 50 cycles to make predictions.

data in the predictive performance. The sensitivity analysis displayed in Fig. 6(b) shows that, for less than 20 cycles, increasing the amount of cycling data does not bring an improvement in performance. Only after 25 cycles is it observed that more data results in higher classification accuracy. This model can deliver 88%-accurate predictions after running as few as 3 charge-discharge cycles. This would allow cell manufacturers to grade their cells during manufacturing in noticeably shorter time (2–3 h instead of the 3–5 h required to run 5 cycles for the model in [25]) and with greater confidence than with today’s processes, then assign appropriate prices and warranty terms more accurately reflecting the cycle life of the cell.

It is well-known in classifier literature that reducing the number of classes (in this case from 3 to 2) would increase the accuracy level, so our classifier (with its 88%-accuracy) can only improve by reducing classes. In fact, the larger inaccuracies in our classifier happen in cross miss-classification between short- and medium-lived (see the 2-by-2 submatrix in the top-left corner of the confusion matrix, Fig. 6(a)). Mathematically speaking, merging the short and medium classes would strongly improve the classifier’s accuracy. This latter classifier, would classify cells as ‘long-knee’ and ‘not long-knee’ which is of interest for battery manufacturers.

Both this classification model and the quantitative prediction model described below require little computational effort; both were trained and evaluated in approximately 0.1 seconds with the 124 cells in the A123 dataset.

Problem 2: Quantitative prediction of knee-point and knee-onset

This problem focuses on giving quantitative predictions of the cycle numbers in which the knee-point and the knee-onset occur, respectively, based only on information from the early life of the cell. Accurate estimations of knee (point & onset) occurrence from early-cycle data have a clear advantage for insurance and maintenance and, importantly, can accelerate the design of fast-charging policies and the validation of new technologies by eliminating the need to run long cycling experiments.

The machine learning model used here stems from *Relevance Vector Machine (RVM)* [31]. The predictive results using information on the first 50 cycles are summarised in Table 2 – 50 cycle information is a substantial advancement regarding Severson et al. [25] who require data from the first 100 cycles.

Fig. 6 (c) displays a sensitivity analysis of the model performance using different amounts of data (from 25 to 100 cycles, i.e. we train 7 RVM models and compare). It can be seen that increasing the number of cycles generally brings a significant reduction in prediction error. The largest improvements are obtained when increasing the number of cycles used from 25 to 50, whereas the error stays almost flat thereafter. This model would allow an end user to accurately predict the knee-point and knee-onset after running only 50 charge-discharge cycles. With the charging policies designed to generate the A123 dataset, each charge-discharge cycle takes between 40 and 60 min. Therefore, a reduction of 50 cycles (from 100 cycles in Severson et al. [25]) translates to reducing

Table 2

Result of RVM regressor for knee-point (a) and knee-onset (b) when predictions are made using the first 50 cycles: MAE and MAPE scores with 90% Bootstrap confidence intervals.

(a) Knee-onset prediction		
Metric	Score	Confidence interval ($\alpha = 0.1$)
MAE (cycles)	55.8	[47.3, 64.9]
MAPE (%)	12.0	[10.2, 13.9]
(b) Knee-point prediction		
Metric	Score	Confidence interval ($\alpha = 0.1$)
MAE (cycles)	57.8	[49.6, 66.7]
MAPE (%)	9.4	[8.2, 10.7]

the length of lab-based cycling experiments by between 35 and 50 h, which in turn brings associated cost savings.

Based on these results, the knee-point seems to be easier to predict than the knee-onset. Their mean absolute error (MAE) is similar throughout the sensitivity analysis, but the mean absolute percentage error (MAPE) is consistently higher for the knee-onset.

Quantifying uncertainty

Providing a measure of uncertainty with each prediction gives insight into how confident the model is in its predictions and allows the end user to control the level of risk they are willing to take in subsequent decision-making processes. Contrary to more standard statistical methods, most machine learning algorithms do not naturally provide uncertainty quantifiers for the predictions they make. The goal of this section is to address this issue in the context of the findings of this study. Such results are necessary for actuarial approaches to warranties/insurance.

In the case of the classification model (Problem 1), an SVM can be adjusted to provide the probability of a cell belonging to each class (short, medium or long), along with the model's decision [32]. This is illustrated in Fig 7(a). If the end user deems that there is a high cost associated with classifying a short-range cell as medium or long, they could decide to tune the model to categorise a cell as short-range whenever the model outputs a probability of the cell belonging to this group of e.g. 20%.

For the quantitative prediction of the knee-point and/or the knee-onset (Problem 2), *Conformal prediction* [33] intervals are used to quantify model uncertainty. Conformal prediction uses past experience to determine prediction intervals, and can be tuned to include the true values with probability p , where p is fixed by the user. Conformal prediction intervals differ from the standard confidence intervals, as the intervals are not centred around the predicted value. As can be seen in Fig. 7(b), the higher the parameter p the more frequently the interval includes the true knee-point, at the expense of having wider, less precise intervals. The end user can use this information to decide when to stop an experiment. The model can output predictions every few cycles and the user could continue cycling only those cells for which the prediction intervals are too broad, until these are reduced to an acceptable level, thus resulting in a more efficient use of laboratory resources and more reliable predictions.

4. Conclusions

Accurate estimates of a cell's lifespan from early-cycle data are of crucial importance for advancing the development of battery technology, improving and fast-tracking manufacturing processes, optimising fast-charging policies and securing a more robust insurance market for batteries. On the other hand, the knee-point of the capacity degradation of a cell provides a more meaningful measure of its lifespan than the cycle life, as it allows one to optimise the use of the battery.

The concept of knee-point is further refined with the novel concept of "knee-onset" to define the point at which the cell shows the first signs of the transition to the accelerated degradation phase. This provides an

earlier warning than the "knee-point", where the rapid degradation is already well underway. The knee-onset thus has significant commercial value regarding warranty provision. It leads to more effective predictive maintenance, since the ability to predict the occurrence of the knee-onset gives the end user enough time to schedule a replacement before the cell fails to meet performance requirements (108 cycles in advance on average). This is critical for energy grids, where batteries are used intensively and where replacements on-the-fly are costly. Novel methods are proposed to identify the knee-point and knee-onset, which are shown to provide consistent and visually accurate results, and can be applied directly without superimposing a degradation model, as they are robust against noise. Moreover, this methodology works even in scenarios where other known methodologies fail.

A side aspect of the "knee" concepts is that the analysis of the data in combination with the results of this study reveal a strong linear relation between cycle life, knee-point and knee-onset. The relation can be used to seamlessly pass from one quantity to the other – predicting one quantity promptly yields the remaining ones.

The "knee" prediction from early-cycle data is also addressed by employing machine learning models with methodologies to systematically extract relevant predictors from cycling data. The resulting models classify cells as short, medium or long-range with 87.9% accuracy and minimal overestimation of knee-point using as few as 3 cycles, and also provide estimates of the occurrence of knee-onset and knee-point with a 12.0% and 9.4% error, respectively, after 50 cycles. Considering that each charge-discharge cycle in the dataset used requires between 40 and 60 min, this means substantial cost and time savings in manufacturing and lab-based experiments with respect to current baselines.

The models' outputs are further enhanced with prediction uncertainty tools which help the end user to control the risk associated with making decisions based on the models' predictions. In other words, a benchmarking analysis tool is provided, which quantifies the reliability of the machine learning models. These error assessments are vital to give this modelling methodology credibility for its use for insurance/warranties, predictive maintenance and manufacturing.

The last point we comment on regards the robustness of the model's extrapolation power to other datasets. We mention that (1) we have taken steps in our methodology to reduce overfitting and achieve good generalisation performance, and (2) despite this, it is possible that our machine learning models, as they are, would not have the same good performance on unseen data that is massively different from the datasets we used (e.g. different temperatures, chemistry, etc.), nonetheless, the crucial novelty is the design of the machine learning pipeline, which can (and should be) retrained whenever new data is available. It is very common for Machine Learning models to become live models that get updated every time the distribution of the input data changes.

Methods

Bacon-Watts models. For both the Bacon-Watts and the double Bacon-Watts models, γ was set to a low value ($\gamma = 10^{-5}$) to obtain an abrupt transition. To learn parameters $\alpha_0, \alpha_1, \alpha_2$ (as well as α_3 in the double Bacon-Watts) and x_0 , the Levenberg–Marquardt nonlinear least squares algorithm is used to fit the model in Eq. (1) to the capacity degradation data. The optimal value of x_0 represents the identified knee-point of a given cell. The same methodology is used for Eq. (2).

Machine learning methods. The machine learning pipeline outlined in this document involves feature extraction and selection, and model fitting and validation. For the classification problem, features are extracted from the raw data using time series analysis and various feature reduction techniques described in Supplementary Figs. 6 and 7. The dimensionality of the resulting dataset is then reduced using recursive feature elimination to select the 100 most relevant features. Lastly, the selected features are standardised by removing the mean and scaling to zero variance, a common feature transformation applied when using linear models.

SVM, selected for the classification model, is a supervised machine learning model that uses the model in (4) to build a decision boundary between two classes by maximising the distance between the closest data points (the support vectors) and the decision boundary,

$$y(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + b. \quad (4)$$

SVMs are effective in high-dimensional spaces, they are fast both in training and testing time, they are robust against outliers and they can be adjusted to work in multi-class tasks, like the one at hand.

Although SVMs are inherently linear models, they can learn nonlinear relations by applying kernel functions which map the inputs to a different feature space. A Gaussian RBF (Radial Basis Function) was selected, a popular Kernel method that uses the euclidean distance in the original space to calculate the similarity between dimensions \mathbf{x} and \mathbf{x}' :

$$k(\mathbf{x}, \mathbf{x}') = e^{-\gamma \|\mathbf{x} - \mathbf{x}'\|^2}, \quad \gamma > 0. \quad (5)$$

The main evaluation metric used to analyse the performance of the resulting model is the accuracy score:

$$\text{accuracy}(\mathbf{y}, \hat{\mathbf{y}}) = \frac{1}{n_{\text{samples}}} \sum_{i=1}^{n_{\text{samples}}} \mathbb{1}(y_i = \hat{y}_i), \quad (6)$$

where y_i is the true class (short, medium or long-range) of sample (cell) i and \hat{y}_i is the predicted class.

Due to the limited number of samples available, leave-one-out [34] is used both for feature selection and model fitting: multiple training sets are created, leaving a different cell out each time. Each model is trained on the remaining samples and the performance is evaluated on the hidden cell. This facilitates the generation of predictions for all the samples available, leading to more statistically robust results whilst minimising the sacrificing of data in development. This ensures that the training sets are as large as possible and that the model is evaluated on all the samples, leading to better predictive performance and more statistically robust results.

For the quantitative prediction problem, a similar feature extraction pipeline is applied, described in Supplementary Fig. 6. This task uses RVM, a Bayesian sparse model suitable for regression tasks that is based on the same principles as SVM, with some enhancements: (i) the output of RVM is a probability distribution of possible outcomes, rather than a point estimate; (ii) it results in more sparse models, making RVM faster than SVM in testing; (iii) it can be less prone to overfitting. RVM model uses

$$p(y|\mathbf{x}, \mathbf{w}, \beta) = \mathcal{N}(t|\mathbf{w}^T \mathbf{x}, \sigma^2) \quad (7)$$

to determine the conditional distribution of a target variable y given a vector of inputs \mathbf{x} . The noise variance is represented by σ^2 and \mathbf{w} are the model coefficients that have to be learned.

This regression model is evaluated using a leave-one-out framework and the chosen performance scores are the mean absolute error (MAE) and the mean absolute percentage error (MAPE):

$$\text{MAE}(\mathbf{y}, \hat{\mathbf{y}}) = \frac{1}{n_{\text{samples}}} \sum_{i=1}^{n_{\text{samples}}} |\hat{y}_i - y_i|, \quad (8)$$

$$\text{MAPE}(\mathbf{y}, \hat{\mathbf{y}}) = \frac{100\%}{n_{\text{samples}}} \sum_{i=1}^{n_{\text{samples}}} \frac{|\hat{y}_i - y_i|}{y_i}, \quad (9)$$

where \mathbf{y} is the vector of true knee-points expressed in number of cycles and $\hat{\mathbf{y}}$ is the vector of predicted values.

Code availability

Code for knee-point/onset identification is available at https://github.com/pfermined/knee_identification.

Code for classification and quantitative prediction is available upon request.

Author contributions

P.F.C. performed data analysis. P.F.C. and G.d.R. performed the modelling and interpretation of results. M.A. proposed the initial Bacon-Watts model and contributed to the associated analysis. E.M., E.M.L., J.S., M.F.A. provided domain expertise and aided in result interpretation. P.F.C., E.M. and G.d.R. edited the manuscript. All authors reviewed the manuscript. G.d.R. supervised the work.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper. Public declaration of Software Disclosure with University of Edinburgh (Edinburgh Innovation) with Invention Id: ERI2019/119, ERI2019/120, Id ERI2019/121 and ERI2020/057.

Acknowledgments

This project was funded by an industry-academia partnership grant Reg-191072 by *The Data Lab – Innovation Centre* (Edinburgh (UK)) and further supported by *Dukosi Ltd*. The authors thank the following colleagues from *Dukosi Ltd* for the helpful discussions, insights and expertise that greatly assisted the research: M. Buhari, J. Darroch, A. Jones and L. Ross. We also thank I. Tudela-Montes and M. de Carvalho (both at the University of Edinburgh) for the helpful discussions.

G. dos Reis acknowledges support from the *Fundação para a Ciência e a Tecnologia* (Portuguese Foundation for Science and Technology) through the project UIDB/00297/2020 (Centro de Matemática e Aplicações CMA/FCT/UNL).

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.egyai.2020.100006

References

- [1] Waldmann T, Wilka M, Kasper M, Fleischhammer M, Wohlfahrt-Mehrens M. Temperature dependent ageing mechanisms in lithium-ion batteries – a post-mortem study. *J Power Sources* 2014;262:129–35. doi:10.1016/j.jpowsour.2014.03.112.
- [2] Dubarry M, Baure G, Devie A. Durability and reliability of EV batteries under electric utility grid operations: path dependence of battery degradation. *J Electrochem Soc* 2018;165(5):A773–83.
- [3] Somerville L, Bareo J, Trask S, Jennings P, McGordon A, Lyness C, et al. The effect of charging rate on the graphite electrode of commercial lithium-ion cells: a post-mortem study. *J Power Sources* 2016;335:189–96. doi:10.1016/j.jpowsour.2016.10.002.
- [4] Gao Y, Jiang J, Zhang C, Zhang W, Ma Z, Jiang Y. Lithium-ion battery aging mechanisms and life model under different charging stresses. *J Power Sources* 2017;356:103–14. doi:10.1016/j.jpowsour.2017.04.084.
- [5] IEEE Recommended practice for sizing lead-Acid batteries for stationary applications. IEEE Std 485–2010 (Revision of IEEE Std 485–1997)2011;:1–90.
- [6] Diao W, Saxena S, Han B, Pecht M. Algorithm to determine the knee point on capacity fade curves of lithium-ion cells. *Energies* 2019;12:2910. doi:10.3390/en12152910.
- [7] Neubauer J, Pesaran A. The ability of battery second use strategies to impact plug-in electric vehicle prices and serve utility energy storage applications. *Lancet* 2011;196:10351–8. doi:10.1016/j.jpowsour.2011.06.053.
- [8] Ecker M, Nieto N, Kbitz S, Schmalstieg J, Blanke H, Warnecke A, et al. Calendar and cycle life study of Li(NiMnCo)O₂-based 18650 lithium-ion batteries. *J Power Sources* 2014;248:839–51. doi:10.1016/j.jpowsour.2013.09.143.
- [9] Han X, Ouyang M, Lu L, Jianqiu L. Cycle life of commercial lithium-ion batteries with lithium titanium oxide anodes in electric vehicles. *Energies* 2014;7:4895–909. doi:10.3390/en7084895.
- [10] Satopaa V, Albrecht J, Irwin D, Raghavan B. Finding a “kneede” in a haystack: Detecting knee points in system behavior. In: Proceedings of the 2011 thirty-first international conference on distributed computing systems workshops. IEEE; 2011. p. 166–71. doi:10.1109/ICDCSW.2011.20.
- [11] Schuster SF, Bach T, Fleder E, Miller J, Brand M, Sextl G, et al. Nonlinear aging characteristics of lithium-ion cells under different operational conditions. *J Energy Storag* 2015;1:44–53. doi:10.1016/j.est.2015.05.003.

- [12] Zhang C, Wang Y, Gao Y, Wang F, Mu B, Zhang W. Accelerated fading recognition for lithium-ion batteries with nickel-cobalt-manganese cathode using quantile regression method. *Appl Energy* 2019;256:113841. doi:10.1016/j.apenergy.2019.113841.
- [13] Diao W, Saxena S, Pecht M. Accelerated cycle life testing and capacity degradation modeling of licoo2-graphite cells. *J Power Sources* 2019;435:226830. doi:10.1016/j.jpowsour.2019.226830.
- [14] Bacon D, Watts D. Estimating the transition between two intersecting straight lines. *Biometrika* 1971;58(3):525–34. doi:10.1093/biomet/58.3.525.
- [15] Hunt MJ, Planella FB, Theil F, Widanage WD. Derivation of an effective thermal electrochemical model for porous electrode batteries using asymptotic homogenisation. arXiv preprint arXiv:1911004762019.
- [16] Bloom I, Cole B, Sohn J, Jones S, Polzin E, Battaglia V, et al. An accelerated calendar and cycle life study of li-ion cells. *J Power Sources* 2001;101(2):238–47. doi:10.1016/S0378-7753(01)00783-2.
- [17] Broussely M, Herreyre S, Biensan P, Kasztejna P, Nechev K, Staniewicz R. Aging mechanism in Li ion cells and calendar life predictions. *J Power Sources* 2001;97:13–21. doi:10.1016/S0378-7753(01)00722-4.
- [18] Ramadesigan V, Northrop PW, De S, Santhanagopalan S, Braatz RD, Subramanian VR. Modeling and simulation of lithium-ion batteries from a systems engineering perspective. *J Electrochem Soc* 2012;159(3):R31–45.
- [19] Waag W, Fleischer C, Sauer DU. Critical review of the methods for monitoring of lithium-ion batteries in electric and hybrid vehicles. *J Power Sources* 2014;258:321–39. doi:10.1016/j.jpowsour.2014.02.064.
- [20] Patil M, Tagade P, Hariharan K, Kolake S, Song T, Yeo T-j, et al. A novel multi-stage support vector machine based approach for li ion battery remaining useful life estimation. *Appl Energy* 2015;159:285–97. doi:10.1016/j.apenergy.2015.08.119.
- [21] Zhou J, Liu D, Peng Y, Peng X. An optimized relevance vector machine with incremental learning strategy for lithium-ion battery remaining useful life estimation. In: Proceedings of the 2013 IEEE international instrumentation and measurement technology conference (I2MTC). IEEE; 2013. p. 561–5. doi:10.1109/I2MTC.2013.6555479.
- [22] Xue Z, Zhang Y, Cheng C, Ma G. Remaining useful life prediction of lithium-ion batteries with adaptive unscented Kalman filter and optimized support vector regression. *Neurocomputing* 2019. doi:10.1016/j.neucom.2019.09.074.
- [23] Zheng S, Ristovski K, Farahat A, Gupta C. Long short-term memory network for remaining useful life estimation. In: Proceedings of the 2017 IEEE international conference on prognostics and health management (ICPHM). IEEE; 2017. p. 88–95. doi:10.1109/ICPHM.2017.7998311.
- [24] Babu GS, Zhao P, Li X-L. Deep convolutional neural network based regression approach for estimation of remaining useful life. In: Proceedings of the international conference on database systems for advanced applications. Springer; 2016. p. 214–28. doi:10.1007/978-3-319-32025-0_14.
- [25] Severson K, Attia P, Jin N, Perkins N, Jiang B, Yang Z, et al. Data-driven prediction of battery cycle life before capacity degradation. *Nat Energy* 2019;4:1–9. doi:10.1038/s41560-019-0356-8.
- [26] Ng S, Xing Y, Tsui K-L. A naive bayes model for robust remaining useful life prediction of lithium-ion battery. *Appl Energy* 2014;118:114–23. doi:10.1016/j.apenergy.2013.12.020.
- [27] Goebel K, Saha B, Saxena A, Celaya J, Christophersen J. Prognostics in battery health management. *IEEE Instrumentation & Measurement Magazine* 2008;11:33–40. doi:10.1109/MIM.2008.4579269.
- [28] Saha B, Goebel K, Poll S, Christophersen J. Prognostics methods for battery health monitoring using a Bayesian framework. *IEEE Trans Instrum Meas* 2009;58:291–6. doi:10.1109/TIM.2008.2005965.
- [29] Finegan DP, Cooper SJ. Battery safety: data-driven prediction of failure. *Joule* 2019;3(11):2599–601. doi:10.1016/j.joule.2019.10.013.
- [30] D’Arpino M, Cancian M, Sergeant A, Canova M, Perullo C. A simulation tool for battery life prediction of a turbo-hybrid-electric regional jet for the NASA ULI program. In: *AIAA Propulsion and Energy 2019 Forum*; 2019. p. 4469.
- [31] Bishop CM. Springer; 2006. p. 325–53.
- [32] Wu T-F, Lin C-J, Weng RC. Probability estimates for multi-class classification by pairwise coupling. *J Mach Learn Res* 2004;5(Aug):975–1005.
- [33] Shafer G, Vovk V. A tutorial on conformal prediction. *J Mach Learn Res* 2008;9:371–421.
- [34] James G, Hastie T, Robert T, Witten D. Springer; 2013. p. 178–81.