Journal of Power Sources 592 (2024) 233910

Contents lists available at ScienceDirect

Journal of Power Sources

journal homepage: www.elsevier.com/locate/jpowsour

ICFormer: A Deep Learning model for informed lithium-ion battery diagnosis and early knee detection

N. Costa [a,*], D. Anseán [b], M. Dubarry [c], L. Sánchez [a]

[a] Computer Science Department, Polytechnic School of Engineering, University of Oviedo, Gijon, 33204, Asturias, Spain
[b] Department of Electrical Engineering, Polytechnic School of Engineering, University of Oviedo, Gijon, 33204, Asturias, Spain
[c] Hawaii Natural Energy Institute, University of Hawaii at Manoa, Honolulu, HI, USA

HIGHLIGHTS

- The approach integrates incremental capacity analysis with a Transformer-based model.
- The self-attention mechanism identifies most relevant changes in data through cycling.
- The model estimates degradation modes and early detects knee points.
- The method outperforms state-of-the-art approaches in degradation modes diagnosis.
- Results are validated on synthetic data and real data from commercial batteries.

ARTICLE INFO

Keywords:
Battery diagnosis
Degradation modes
Knee detection
Deep learning
Transformers

ABSTRACT

Accurate diagnosis of lithium-ion battery (LIB) degradation is critical for safe and reliable operation in real-world applications. In recent years, data-driven approaches powered by Machine Learning algorithms emerged as a promising solution, among which Deep Learning methods were proven to be effective for various tasks such as State of Charge and State of Health estimation and Remaining Useful Life prediction, however, their application for knee point estimation is still at an early stage. While experimental techniques, were demonstrated to be successful for detecting knees, current methods focus on tracking individual cycles rather than analyzing the overall trajectory of degradation, which could potentially offer better performance. In this paper, we introduce ICFormer, a novel Deep Learning model based on a Transformer encoder that leverages self-attention on the evolution of incremental capacity curves to accurately identify relevant changes in LIB degradation trajectories. The proposed model not only detects knees, but also anticipates them while also outperforming state-of-the-art approaches in diagnosing degradation modes. The effectiveness of the model is validated using both synthetic and experimental data. We demonstrate that the method can provide valuable knowledge on the factors contributing to capacity loss and offer advanced insights for battery management and predictive maintenance strategies.

1. Introduction

Lithium-ion battery (LIB) lifetime requirements are becoming increasingly challenging to accelerate decarbonizing energy systems such as stationary energy storage and electric vehicles [1–3]. As these systems are used and designed for a vast range of applications worldwide, their batteries will exhibit performance losses that are path-dependent and influenced by the usage conditions (e.g., temperature, power rates, duty cycles, etc.) [4–6]. As a result, LIBs can exhibit degradation trajectories that are linear, sublinear, or superlinear, with the latter being commonly referred to as a "knee" curve [7–9]. The accurate modeling and prediction of knees present notable challenges necessitating a dedicated and comprehensive framework since, as highlighted in a recent review by Attia et al. [9], knee "pathways" are intrinsically linked to the battery degradation mechanisms. Thus, degradation identification must serve as a key initial stage in detecting and mitigating the formation of knees to ensure optimal performance and longevity of LIBs. This feature is lacking in most knee detection methods proposed so far in the literature [9–15].

To date, several approaches have been proposed to estimate battery degradation mechanisms, ranging from simple empirical models to

* Corresponding author.
E-mail address: costanahuel@uniovi.es (N. Costa).

https://doi.org/10.1016/j.jpowsour.2023.233910
Received 8 June 2023; Received in revised form 13 October 2023; Accepted 28 November 2023
Available online 9 December 2023

more sophisticated data-driven and physics-based methods. Physics-based models were developed to describe the behavior of LIBs based on physical principles and parameters [16,17]. Despite their advantages, these approaches require accurate knowledge of the parameters and assumptions used in the models, which is challenging to obtain in practice without significant experience and experiments. For diagnosis purposes, a potential workaround is to focus on degradation modes rather than mechanisms and to use a mechanistic electrode model [18,19] for high-fidelity emulation of the voltage response of an aged cell to the detriment of stress specific information. Data-driven approaches rely on data collected from LIBs to train models that can make predictions about battery behavior [20]. These approaches do not require a priori knowledge of the parameters used in physics-based models and can be trained using large amounts of data to provide accurate predictions [21]. They are also well-suited for online applications because they can be implemented in real-time without the need for extensive upfront planning or modeling. Within the data-driven approaches, Machine Learning (ML), and specifically Deep Learning (DL) approaches are becoming more prominent. In particular, Recurrent Neural Networks (RNN) or Convolutional Neural Networks (CNN) have been proposed to estimate SOH [22,23], predicting remaining useful life (RUL) [24,25] and quantify degradation modes [26,27]. However, the application of DL methods for knee investigation is still at an early stage [9,28], although some promising works are starting to emerge. In [29] the authors introduced a RNN-based S2S (sequence-to-sequence) method for predicting the whole degradation trajectory of capacity till the end of life using 100 cycles of data. In [30] the knee point detection was modeled as a classification problem, where each training sample belongs to the positive class if it happens before the knee point, and to the negative if it is after. Then, they employed recurrent neural networks trained with synthetic data for detecting the knee. In [31] a CNN-based model was proposed to extract temporal features from time-series data under variable battery usage to predict the number of cycles left to reach the knee-point. In [10] the complete degradation curve is estimated using the data of the first cycle with the help of a CNN.

Although existing DL approaches for knee detection represent a valuable advancement in the field, they often fall short in providing a comprehensive understanding of the internal degradation mechanisms of the battery, which are essential for unraveling why knees occur. Moreover, many of these methods rely on common variables such as current, voltage, or temperature as inputs for their models. While these variables offer insights, they may not provide the depth and precision needed for an efficient performance. To shed light on this issue, it is imperative to delve deeper into the degradation trajectories, which are currently categorized into knee "pathways" [9]. Some of these pathways are linked to different combinations of the battery's degradation modes, i.e., Loss of Lithium Inventory (LLI) and Loss of Active Material (LAM) [18,19] for which some could be "hidden" or "silent" and associated with little or no capacity loss until becoming prominent [9,18]. Therefore, the detection of knees under these pathways cannot be achieved without careful evaluation of cell-level electrochemical signals. This can be done using tools like the Incremental Capacity Analysis (ICA), a technique that has been proven successful in quantitatively estimating degradation modes in most available LIB chemistries [32]. Since ICA relies solely on the voltage response of the cell, it is implementable in deployed systems without the need for additional sensors. It is also highly automatable and is therefore a good candidate for integration with ML and DL approaches and for knee points prediction.

However, the applicability of ICA for DL is limited by the nature of the battery testing data. Typically, the data for ICA analysis is obtained during Reference Performance Tests (RPTs) carried out at different stages of the battery cycle-life [32,33]. Although the data is sequential, there is not a single data point for each RPT as would occur in a typical time series, but a series of points representing the IC curves versus cell voltage evolution, as illustrated in Fig. 1. As a result, the use of DL algorithms for sequential/time series data analysis like RNNs is constrained since they are not designed to support this kind of data. Instead, current ML methods for analyzing the evolution of IC curves rely on independent diagnostics at each RPT (this would be equivalent to diagnosing each data point independently in a regular time series) [34–37], which are combined to obtain the final diagnosis, rather than analyzing the complete degradation trajectory [26,37]. While this approach is valid, independent diagnostics do not rely on previous data which can lead to some inconsistencies. Besides, since the models learn the signature of each IC curve their performance can vary on real data due to cell-to-cell variations as their signature may not match exactly [38–40]. In contrast, analyzing the evolution of the IC curves as a whole offers the opportunity to model the temporal changes of the IC curves, rather than their own signature, which may be less prone to errors. To the best of our knowledge, despite its potential advantages, this specific aspect remains unexplored in current literature.

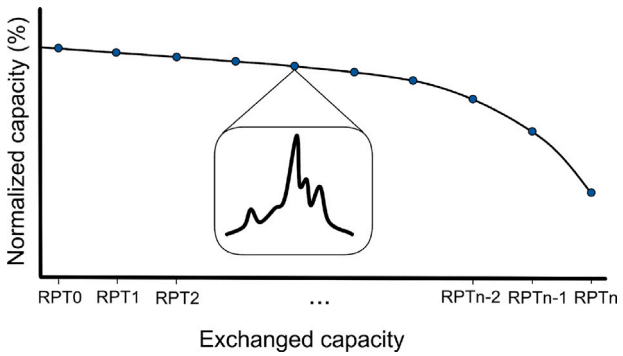

Fig. 1. Normalized battery capacity plotted against the reference performance test (RPT) number, where each RPT provides an individual incremental capacity (IC) curve throughout the battery cycle-life. It is important to note that while a typical time series have a single data point for each variable at each time step, IC analysis yields a complete curve representing the incremental capacity versus cell voltage evolution.

Modeling the evolution of IC data over time could potentially be undertaken by novel models capable of processing multidimensional data effectively, like the Transformer. Transformers [41] are a type of neural network that have emerged as revolutionary in sequential data analysis because of their ability to capture long-term dependencies in the data with the self-attention mechanism that enables the model to selectively focus on different parts of the input sequence and weigh their importance. Also, because unlike RNNs, they can process data in parallel rather than sequentially and handle different types of data, making them broadly applicable to fields such as language modeling, computer vision and audio processing [42,43]. Although Transformers offer significant promise for analyzing battery degradation, they heavily rely on the availability of large, diverse and precise datasets. In this regard, most available data are limited to small training datasets (e.g., Oxford [44], NASA [45] or CALCE [46]) consisting of cells tested under limited unrepresentative conditions and where capacity loss is mostly linear. This is far from the big amounts of data needed to make these approaches accurate in deployed systems. Steps in the good direction were made with the availability of different larger databases [47,48] and with the emergence of synthetic data [35,49] computed from high-fidelity models. Therefore, given the current data availability, it is most likely that the best chance for training DL models is to rely on synthetic training datasets [49].

Herein, we introduce ICFormer, a novel approach that integrates ICA with a Deep Learning model for informed battery diagnosis and early detection of knees. Contrary to previous works on ICA that analyze each IC curve independently we propose a Transformer architecture that learns to focus on the most relevant changes in the IC

curves throughout the complete degradation trajectory. The resulting model addresses diagnosis through degradation modes quantification and prognosis by forecasting the occurrence of future knee points. The model was trained on the HNEI synthetic prognosis dataset [49], and was validated on both synthetic and experimental data. Additionally, the model's performance was compared with state-of-the-art ML approaches, demonstrating its improved ability to quantify battery degradation. Overall, the approach presented here shall provide a promising solution for improving the performance and longevity of battery systems in real-world applications.

2. Relationship between degradation modes and knee pathways in lithium-ion batteries

Degradation mechanisms in LIBs are numerous and complex, with a multitude of different mechanisms induced by the complex interactions between physical and chemical processes [50,51]. Because of path dependence [7–9], it is well known that each deployed LIBs will experience a unique mix of these degradation mechanisms, of which some could create a pathway where a knee could occur [9,52,53]. While some knee pathways are still hypothetical [9], some are well understood and are known to showcase identifiable changes on the cell's electrochemical response [18]. Among those are the two induced by the loss of active material on the negative and positive electrodes (LAM_{NE} and LAM_{PE}, respectively). These degradation modes are, for graphite-based cells, referred to as "silent" or "hidden" because they do not induce capacity loss in their initial stages, even if cell degradation is taking place [18,19,32]. Nonetheless, when either LAMs is high enough to change which electrode is limiting, either at the end of charge or end of discharge [32], it will eventually play a role in capacity loss, leading to knee points in the cell's trajectory either directly in the case of LAM_{PE} or indirectly via irreversible plating for LAM_{NE} [9,32].

Fig. 2 illustrates the relationship between the degradation modes, their silent nature, and their impact on capacity and cell electrochemical response for three case figures with (a–c) LLI > LAMs, (d–f) LAM_{PE} > LLI, and (g–i) LAM_{NE} > LLI. The figure presents the simulations for C/25 charges of a graphite/lithium iron phosphate (LFP) cell performed using the 'alawa toolbox mechanistic model [18]. The top row (a,d,g) showcases individual electrode evolutions, the middle row (b, e, h) the associated capacity loss and the bottom row (c,f,i) the evolution of the IC signature. When pristine, commercial LIBs exhibit an excess of relative capacity for each electrode outside the full cell voltage window, identified as the initial offset OFS_{ini} and the NE excess (NE_{ex}) [32]. OFS_{ini} is induced by the lithium consumption during formation of the solid electrolyte interphase (SEI) upon the first cycles [54,55]. In addition, since the electrode matching in the mechanistic model is done after SEI formation, this approach accommodates to diverse cell manufacturing processes, including novel prelithiation strategies [56, 57]. As for the NE excess, it is the result of standard manufacturing practices consisting in oversizing the NE to ensure its counterpart with the whole PE surface to avoid lithium plating [58,59]. These excesses of relative capacity can be seen as a "buffer" that allows the electrodes to sustain a certain amount of active material loss before reaching the voltage window of the full cell and causing direct capacity fade [18,32] initially and a knee later on. The rate at which this buffer is consumed is related to the composition of the degradation. For example, when LLI is the main degradation mode (i.e., LLI > LAM_{PE} and LLI > LAM_{NE}, see Fig. 2a), the overall slippage of the NE is larger than the PE's shrinkage caused by LAM_{PE}, so its effect will never catch up on the full cell. As the NE's shrinkage rate is smaller than the slippage (i.e., LAM_{NE} < LLI) the NE will be shifted towards higher relative states of charge without affecting the full cell voltage window, and the capacity loss under these conditions will be solely caused by the effect of LLI, not producing any knee in the cell trajectory (Fig. 2b). However, when LAMs are larger than LLI (i.e., LAM_{PE} > LLI and LAM_{NE} > LLI), see Fig. 2(d,g), the rate at which the electrodes shrink is larger than the NE's slippage rate. Initially, full cell capacity fade is only caused by LLI because of the excess of relative capacities. Eventually, there is a change in the limiting electrode and the LAMs become active within the cell voltage window and they start to play a role in capacity fade [32]. This is illustrated by a star in Fig. 2(d,e,g,h). From this perspective, it can be stated that knees generally occur when LAM > LLI, whereas if LLI is the main degradation mode, LAMs remain silent in terms of direct capacity loss on the full cell regardless of whether LLI's and LAM's evolution are linear or nonlinear. For clarity, while LAM_{PE} could induce direct capacity loss, LAM_{NE} cannot if plating is possible in the considered potential window. When LAM_{NE} > LLI, plating will eventually occur and only its irreversible part will contribute to capacity loss as LLI [32,53,60]. Additionally, we should note that if LLI exhibits an exponential evolution, its rate of capacity loss will be directly reflected in the capacity evolution.

Although LAMs do not cause direct capacity loss at the beginning, they do have a traceable impact on the cell's electrochemical response. To track these changes, the IC technique can be used, as shown in Fig. 2(c,f,i). The arrows indicate the main features of interest (FOIs) that are used to evaluate each of the dominant degradation modes (i.e., LLI, LAM_{PE} and LAM_{NE}). LAM_{PE} induced capacity loss will always be accompanied by the gradual disappearance of peak 5 and the reversible part of the plating will infer a new peak 0 at high voltage. Those interested in the IC technique coupled with the expert-designed FOIs in graphite/LFP systems to track degradation modes can refer to previous works [32,61–64] for detailed insights. These indicators allow for knee prediction [60] and this emphasizes the potential of using the IC curves to link silent degradation modes and sensitive features (i.e., FOIs) to identify knee trajectories.

The aim of this work is to develop a model that analyzes temporally the changes in the IC curves and associates them with the current or future presence of knees. To do so, it is necessary to first identify and label knee occurrences within the data. While some methods for knee detection were proposed in the literature [65–67] we elected to a custom one to take full advantage of the information available within the synthetic datasets with the capacity loss and the evolution of the different degradation mode per cycle. The algorithm first checks on the linearity for the capacity loss. For the non-linear duty cycles, a second check is done on the degradation modes to decipher if a knee was present or not, and if that knee was anticipable (i.e., resulting from the interaction between modes and not from a superlinear behavior). The latter is a unique aspect of our work that cannot be obtained from any other methods to the best of our knowledge. More details on the procedure and the algorithm itself are provided in Appendix A.

It is important to note that this represents an initial approach to addressing the problem, and further exploration is required in future work to refine the methodology and address any potential bugs or limitations.

3. Methods

Transformers are responsible for great advances in AI in recent years such as GPT [68], GATO [69], ALPHAFOLD [70], or DALL-E [71], however, they have just started to be explored in the battery field. This can be due to data scarcity, as explained in the introduction, since Transformer models require training with significant amounts of data in order to generalize well to a wide range of task and/or scenarios, yet data availability in the battery field is still a limitation. Additionally, the focus on natural language processing (NLP) and image recognition applications may have led researchers to overlook the potential of Transformers in battery technology. Nevertheless, there have been recent work in which Transformer encoders were trained for SOH prediction and RUL estimation (RUL) [72,73], and where researchers explored the use of an encoder–decoder Transformer for aging-aware battery discharge prediction [74] and for SOC/SOC estimation [75,76], however, these works make use of variables different that IC curves

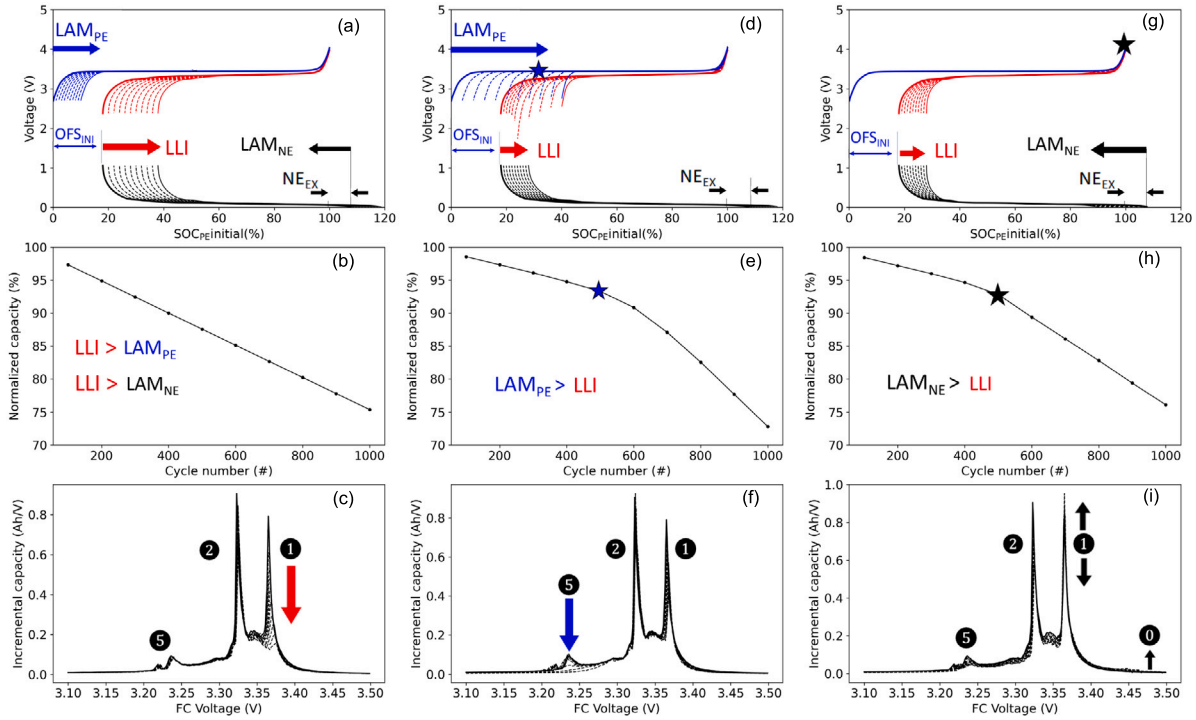

Fig. 2. Examples of the relationship between battery degradation modes on half-cell perspective (panels a, d, g), their impact on capacity loss (panels b, e, h) and the resulting incremental capacity curves (panels c, f, i). The half-cell representation illustrates the silent nature of the degradation modes, where the star represents the knee onset. Knees in capacity trajectory occur when LAM > LLI (panels d–i), while LAMs remain silent if LLI is the dominant degradation mode (panels a–c).

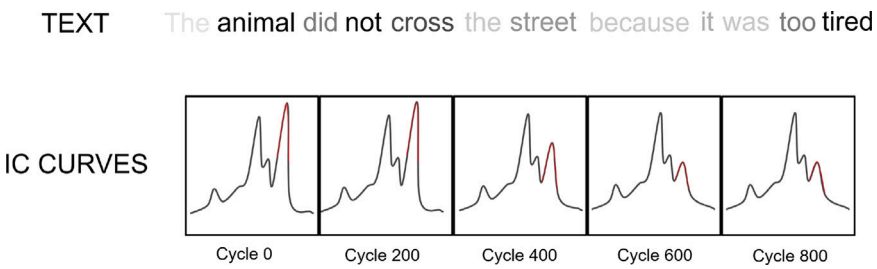

Fig. 3. Examples of the application of attention for text and IC curve analysis, in which the resulting scores will allow paying attention to the most important parts of the sequence to be analyzed.

such as current, voltage and temperature and are not intended for knee detection or degradation modes estimation. In any case, as more data becomes available, it is likely that we will see more applications of Transformer-based models in battery technology in the future.

Transformers originally emerged as an alternative to RNNs for text processing to solve the problems of memory collapse and sequential training, for which attention and positional encoding were proposed as components of the architecture [41]. The attention mechanism allows the model to focus on the most important parts of the input sequence regardless of their position, which enables processing very long data sequences without losing information. For text, the key idea behind attention is that each word in a sentence can be related to all other words and that the strength of these relationships can be learned by the model. This allows the model to capture complex relationships between words that may be far apart in a sentence and to effectively encode long-term dependencies. The upper part of Fig. 3 shows a simple example of the attention mechanism for text. Each word has an associated attention score (see Appendix B for more details on how attention is calculated), which is represented by the intensity of the color, with white being the lowest score and black the highest. Thus, attention is placed on words that are relevant to the context like "animal", or "tired" and not so much on other less important such as "the" or "was".

Although Transformers were initially oriented to text processing, they have been adapted to deal with different types of data and multiple dimensionalities. This have been made possible by the positional encoding layer, which allows to process all the information at once. The layer receives a sequence of data and provides positional information for the model to know the order of the data. In the case of text, different sentences are fed to the model, they are given a positional encoding and the resulting information is fed to a Transformer block, which contains the attention mechanism. In [77], the authors proposed to decompose images into multiple parts and treat them as a sequence. This sequence of sub-images, also called "patches", was then passed to the positional encoding layer to give the model information about the order between the different parts and then to a standard Transformer. Recently, this approach has also been investigated in time series processing. In [78], the authors divided time series data into different parts or patches and fed that to the positional encoding layer and to a standard Transformer. In layman's term, a good analogy could be to see the encoding process as the creation of a stop-motion film from events that occurred at different times.

As discussed in the previous section, the analysis of evolution of the IC curves along cycles is a substantially more complicated task than the analysis of common variables such as current or temperature

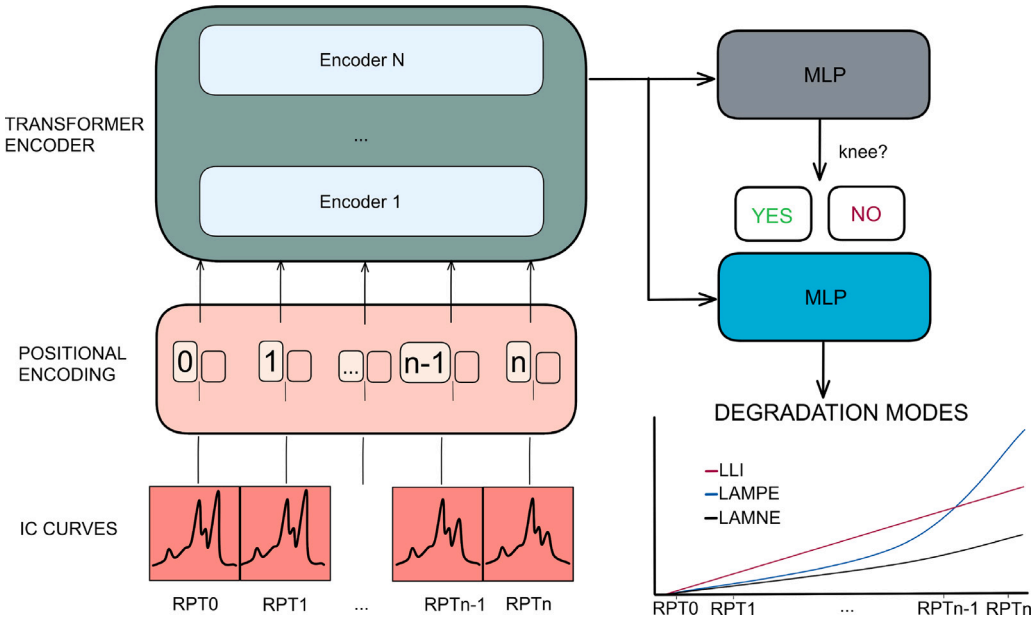

Fig. 4. Pipeline of the proposed solution. First, the IC curves are processed through a positional encoding layer for giving a sequential order to the input data. Then, the processed data is fed to N encoder layers whose output is passed to two independent MLP layers, one for the identification of knees and the other for quantification of degradation modes.

because they do not form a typical time series but a sequence of curves (Fig. 1). Analogously to the aforementioned works where images or text were divided into different parts/patches, in our problem the battery degradation histories are already divided into different patches, the RPTs, where each patch has a corresponding IC curve. As a result, the full set of curves obtained during RPTs could be fed to a positional encoding layer and subsequently to a Transformer encoder, which contains the attention layers. The bottom of Fig. 3 shows how this could be adapted to IC curves. The peak marked in red loses height while the others remain intact, so the model will pay more attention to that part of the curve to identify where the changes have occurred. That is, the model will place more focus on peaks that change during battery cycling to model the degradation state of the battery.

3.1. Model

The main components of the model are illustrated in Fig. 4. First, IC curves obtained from a set of cycling conditions are fed to a positional encoding layer. The resulting representation is then processed by a traditional Transformer encoder, which is actually made up of several stacked encoders containing the attention layers (the reader can refer to Appendix B for specific information on the encoder implementation). The encoder output contains information that eases the resolution of other subsequent tasks, which in this case are the identification of knees in capacity loss and the quantification of degradation modes. For this, the encoder output can be processed by simpler models such as multi-layer perceptrons (MLPs). An MLP with a sigmoid activation function is used to predict whether the battery is undergoing a knee or a silent degradation that eventually leads to a knee (gray block) and, equally, another MLP with a linear activation is used to predict the three degradation modes: LLI, LAM_{PE}, and LAM_{NE} (blue block).

It is noteworthy mentioning that not only the IC curves can be fed to the model but also other available variables such as current, temperature, or resistance. The methodology for incorporating new variables is direct, instead of feeding the complete sequences, they are divided into as many parts/patches as IC curves have been fed to the model, which we will explore in future work.

3.2. Experimental design

The data used for this work is the LFP HNEI synthetic publicly available prognosis dataset, which is composed of more than 125,000 duty cycles at C/25 charges [49]. It includes capacity and voltage at most every 200 cycles for each duty cycle, from which the IC curves were calculated. The synthetic thermodynamic C/25 data provided at cyclical intervals effectively replicates the information that would be available from RPT datasets. An added benefit of synthetic data is its ability to represent any conceivable battery degradation scenario that could take during the duty cycling scheme, including different linear degradation rates, linear degradation combined with exponential variations, or delayed exponential increases. The resolution of the data is originally 1001 points over the voltage window. In order to avoid processing very long sequences, the resolution was downscaled to a point every 2.3 mV per IC curve using a 1-D monotonic cubic interpolation with the Scipy Pchip Interpolator [79], resulting in curves of 128 points. This is similar to what was used in our previous work where this data reduction was shown to have no impact on peak position, intensity, nor shape [37] and it thus kept the main FOIs intact while limiting the curve size.

Since each duty cycle undergoes different degradations, the end of life of each sample can be reached at different stages, which means that the size of the overall samples is variable. Although Transformers can deal with sequences of varying size, it is known that they work best if all samples have the same length, as this allows for efficient matrix multiplication and parallelization [41]. To mitigate this issue, a fixed sequence size was used with a sliding window approach to cover the entire length of the duty cycles. For each duty cycle within the dataset, the curves were sampled every 800 cycles, i.e., data are taken first for cycles 0 to 800, then 200 to 1000 and so on until the end of the life, as shown in Fig. 5. The 800 cycles width was chosen based on this particular dataset (200 cycles resolution and little failed cells at cycle 800) and this will need to be adjusted for other datasets. Windows that contained exactly the same IC curves were discarded in order not to repeat data. The resulting samples constitute the training set, which was composed of 293,286 samples from the original 125,000+ different duty cycles. This data augmentation via the sliding window approach also reinforces the analysis of degradation patterns that can help the algorithm to better model the battery behavior.

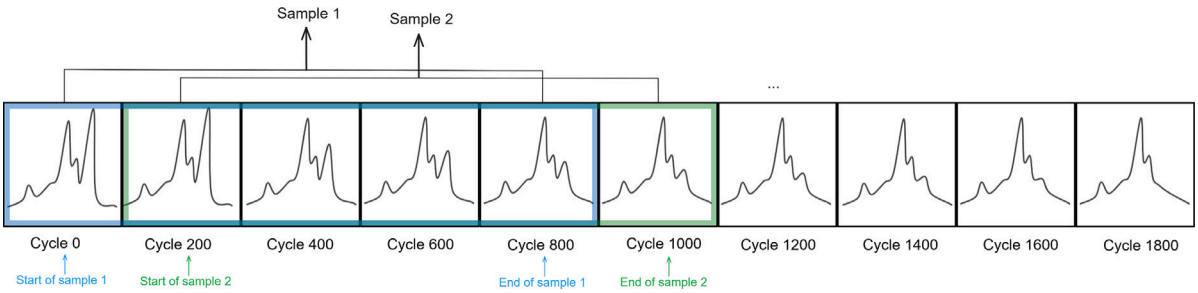

Fig. 5. Sliding window on the IC cycling history for one duty cycle. A window size of 800 cycles is applied, in which the first sample extracted would be the history from cycle 0 to 800, the second from 200 to 1000, and so on.

A standard training/validation/test split protocol was followed to evaluate the performance of the model. 70% of the 293,286 training samples were used to train the model and optimize its parameters while the remaining 30% were used as the validation set. The test set comprised the three virtual cells introduced in our previous work [37]. These virtual cells were specifically designed to differ from the original dataset's cell configurations, thus mitigating the risk of overfitting. In total, 1000 synthetic duty cycles were generated for three distinct cell configurations denoted as C1, C2, and C3 covering a wide range of degradation levels. The parameters used to generate the data in the 'alawa toolbox can be seen in Table 1.

For each sample in the training set, the model is trained to identify whether it has or will have a knee and to quantify the 3 modes of degradation (LLI, LAM_{PE} and LAM_{NE}). The information on the degradation modes is already included in the original dataset while the information on the knees is obtained as a result of applying our custom algorithm, as explained in Section 2. The metrics chosen to optimize the model are the binary cross entropy loss for knee identification, as it is a classification task, and the Mean Squared Error loss for quantification of degradation modes, as it is a regression task. The model was trained to satisfy both objectives therefore the error function to be minimized is the equal sum of both:

$$\mathcal{L} = \underbrace{-\frac{1}{N}\sum_{i=1}^{N} y_i \log(\hat{y}_i) + (1-y_i)\log(1-\hat{y}_i)}_{\text{binary cross-entropy}} + \underbrace{\frac{1}{N}\sum_{i=1}^{N}(y_i-\hat{y}_i)^2}_{\text{Mean Squared Error}}$$

The validation set was used to fine-tune the hyperparameters, which mainly affect the number of encoders and the number of neurons in the MLP layers (see Appendix B for more details). A grid search approach was used for hyperparameters search, where we defined a range of values for each hyperparameter and then trained and evaluated the model for each combination of hyperparameter values, from which the best combination was selected to determine the final model. In addition, to achieve the optimal models' performance, an early stopping procedure was used to relegate the training stop condition to the validation error instead of the number of epochs. Once the hyperparameters were optimized, we evaluated the performance of the model with the best set of hyperparameters on the test set. The WandB framework [80] was used to track all the process.

The test set enabled to thoroughly assess the diagnostic capabilities of our model in accurately quantifying degradation and providing prognostic insights regarding knee anticipation. This approach enabled us to evaluate the model's performance on unseen data and ensuring the generalizability of our findings.

4. Results and discussion

To comprehensively evaluate the performance of the proposed model, this section is structured in three parts. The model's workflow to predict knees and estimate degradation modes will be discussed first. Second, the numerical results for knee identification and degradation modes estimation will be presented with a comparison with state-of-the-art approaches for the latter. Finally, a detailed case study using experimental data from commercial batteries to provide a real-world validation of our approach will be provided.

Table 1
Details about virtual test sets. Three cells, labeled as C1, C2, C3, were generated using the 'alawa toolbox for each chemistry. The values of LR (Loading Ratio) and OFS (offset) with which they were generated are included. Parameters used for the training data are also added to highlight the differences with respect to the test sets.

	Training data		C1		C2		C3	
	LR	OFS	LR	OFS	LR	OFS	LR	OFS
LFP	0.95	12.5	0.96	11.5	0.94	12.5	0.95	11.5
NCA	1.05	1.5	1.06	0.5	1.04	2.5	1.05	0.5
NMC	0.90	10	0.91	9	0.89	11	0.90	9

4.1. Validation of model workflow

To validate the model workflow, Fig. 6 provides an illustration of knee prediction and degradation modes quantification using four representative scenarios selected from the 125,000 synthetic duty cycles dataset used to train the model. The scenarios correspond to different predominant compositions for the degradation: LLI > LAM, LAM_{PE} > LLI, LAM_{NE} > LLI, and LLI~LAM_{NE}. Using the entire IC curve history (Fig. 6, left column), the model analyzed every 800-cycle window and predicted the presence of knees with an assertive, simple label (i.e., "YES" or "NO"), Fig. 6, central column. It is important to remember that capacity loss was not used as an input for the model nor was it part of the predictions. It is shown in Fig. 6 center column for illustrative purposes only. Finally, the model provides a quantitative estimation of the evolution of the degradation modes Fig. 6, right column.

The first degradation scenario (LLI > LAM) is shown in the first row of Fig. 6. By analyzing the IC curves, the model did not predict any knee appearance throughout cycling, as indicated by the label, "NO". The primary degradation mode predicted by the model is linear LLI, with a degradation rate that matches that of the normalized capacity. This indicates that both LAMs remain silent in terms of capacity fade, and all capacity loss is directly caused by LLI. The second degradation scenario (LAM_{PE} > LLI) is shown in the second row of Fig. 6. Based on the evolution of the IC curves in the first 800 cycles, the model accurately predicts that a knee will appear at some point during cycling, which is the case at cycle 1200. The predicted degradation modes reveal a higher rate of LAM_{PE}, which initially remains silent before eventually catching up to create a knee in the capacity trajectory. The third degradation scenario (LAM_{NE} > LLI) is shown in the third row of Fig. 6. While the algorithm does not immediately identify the knee, it is detected in the window corresponding to cycles (400–1200), before the knee takes place. In this scenario, LAM_{NE} initially remains silent before catching up to cause direct capacity fade. The last degradation scenario (LLI~LAM_{NE}) is shown in the fourth row. Contrary to the

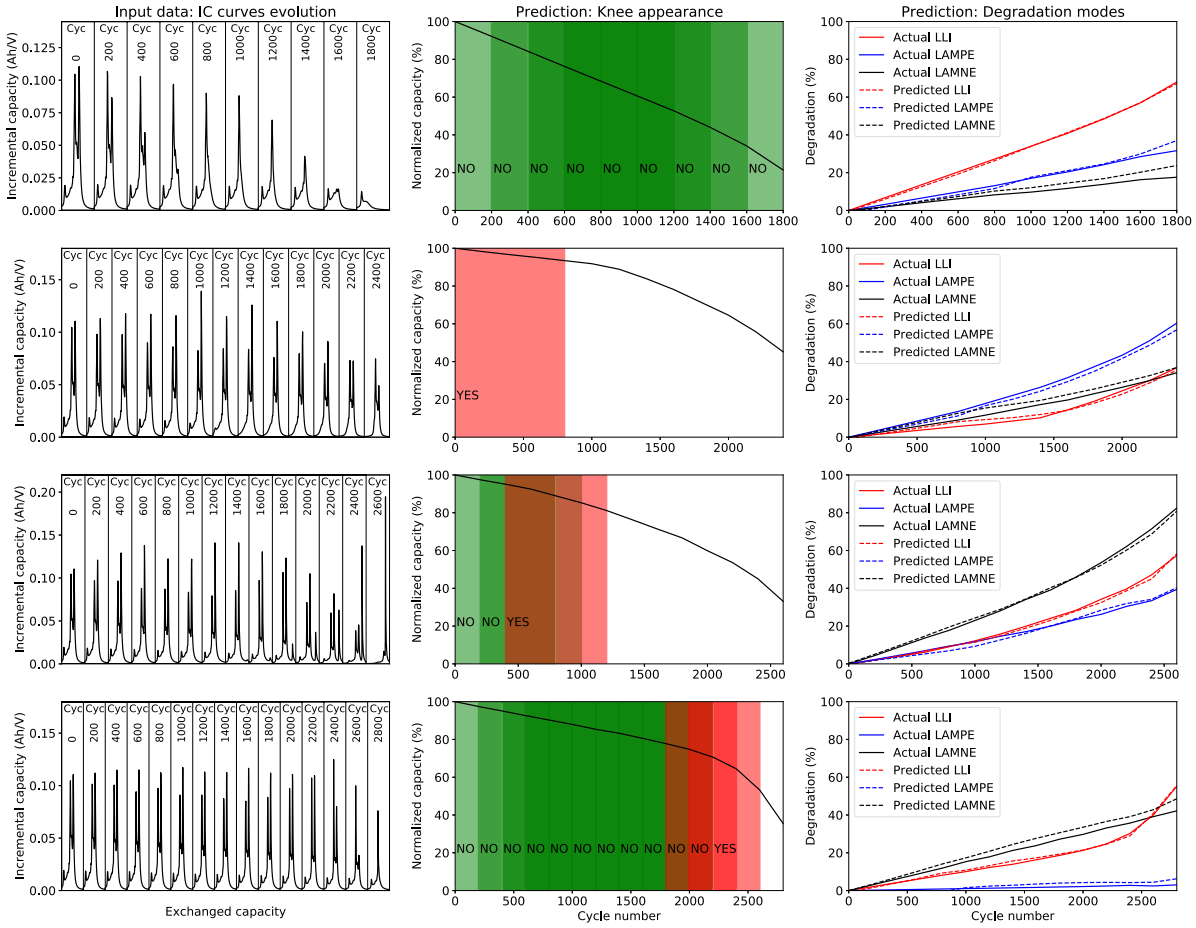

Fig. 6. Model performance in four representative degradation scenarios. The left column shows the IC curves to be analyzed and the two adjacent columns show the model predictions: knee identification and estimated degradation modes. The model predictions are made for each 800-cycle window every 200 cycles, so the intensity of the red/green colors is only due to the overlap of each prediction window, for example from cycles 200–600 of the first window from 0 to 800 and the second from 200 to 1000. In the case of degradation modes, the final predictions are shown, which are taken as the average of predictions in the overlapping windows. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

previous scenarios, in this example, the algorithm does not predict the knee appearance until it already takes place. The model predicts a similar degradation rate for LAM_{NE} and LLI during the first 2000 cycles. Since both degradation rates are similar and there is no knee in capacity evolution, this is an indication that capacity fade is caused by LLI while LAM_{NE} and LAM_{PE} remain silent. However, at cycle 2200 a sudden increase of LLI takes place, rapidly surpassing the degradation rate of LAM_{NE}. As LLI is producing direct capacity fade and does not induce any silent incubation, the model cannot predict this sudden change of LLI rate, yet still provides an indicator that a knee is taking place. Additional examples on the training data illustrating the impact of degradation modes at different cycling stages on the emergence of knees and their relation to degradation modes can be found in Fig. D.8, and the model's performance using real experimental data will be presented in Section 4.3.

4.2. Numerical results

In order to test the performance of our approach, the diagnosis accuracy was tested versus five state-of-the-art methods. The methods tested in this study were the same that were used in our previous work [37,81] with a Random Forest Regressor [82], a 1D Convolutional Neural Network (1DConv) [35], a Feed-forward Neural Network (FNN) [36], and our Dynamic Time Warping Convolutional Neural Network (DTW-CNN) [37]. More details on the methods can be found in [37], which are publicly available. These five methods are only intended to predict degradation modes and do not directly provide any information on the existence of knees. To the best of our knowledge, there is no other model that enables both simultaneously.

Table 2 presents the diagnosis accuracy in terms of RMSE for the quantification of the three degradation modes at five different cycles (200, 400, 600, 800, and 1000) for the 1000 duty cycles of the three test cells C1, C2 and C3. The best performing method for each degradation mode is highlighted in bold. Our method exhibits a more evenly distributed dispersion among the different degradation modes, which is also noticeable in the standard deviation (see Table 3). Notably, our model demonstrates superior diagnosis performance with an average error of 2.33% versus errors from 3.12 to 6.85% for the other approaches. This highlights the effectiveness of our model in accurately predicting degradation modes by analyzing the evolution of the IC curves. Specifically, it is worth noting that predicting LAM_{PE} in graphite/LFP systems with the other methods presents unique challenges because the flat voltage response of LFP results in almost no discernible peaks on the IC curves [32], making the detection of LAM_{PE} more difficult. Additionally, the intrinsic stability of LFP [83] further contributes to this challenges, as any potential degradation in such stable systems tends to be relatively small and therefore, difficult to determine. However, this is not the case for the ICFormer. This phenomenon is directly attributed to the difference between analyzing the history of IC curve patterns and analyzing their individual signatures. When curves are analyzed independently there is a risk of prediction errors, which may penalize one mode and favor others (in this case LAM_{PE} over LLI and LAM_{NE}), because different degradation modes may

Table 2
RMSE results for each degradation mode and cycle for the tested cells.

		LLI					LAM_{PE}					LAM_{NE}				
RF	C1	3.62	3.23	5.28	7.02	9.21	3.15	5.16	7.06	8.26	9.13	3.82	6.24	9.36	10.39	11.83
	C2	3.64	3.14	5.22	7.03	9.22	3.17	4.97	7.30	8.42	9.79	3.82	6.38	8.99	10.31	11.55
	C3	3.62	3.20	5.19	7.10	9.13	3.15	5.07	7.22	8.27	9.37	3.82	6.18	9.18	10.30	11.66
FNN	C1	1.82	1.67	2.31	3.26	3.91	3.28	3.58	5.51	8.09	11.11	2.10	2.15	2.80	4.34	6.31
	C2	1.81	1.55	2.25	3.08	3.67	2.73	3.46	5.54	8.30	11.32	2.77	2.35	2.85	4.31	6.19
	C3	1.68	1.73	2.39	3.35	4.02	3.52	3.78	5.65	8.27	11.31	2.64	2.44	2.93	4.38	6.37
1DConv	C1	1.06	1.68	2.39	2.87	3.21	2.80	3.38	5.10	7.49	10.73	**1.71**	2.83	4.01	5.04	6.60
	C2	1.11	1.62	2.25	2.68	3.15	2.62	3.50	5.21	7.72	10.85	2.36	2.86	3.85	4.88	6.58
	C3	0.96	1.75	2.40	2.91	3.35	2.95	3.44	5.17	7.60	10.86	2.07	2.93	3.92	4.96	6.61
DTW-CNN	C1	1.16	1.31	1.59	2.28	2.47	2.67	3.59	4.90	6.40	8.64	1.98	**1.93**	**1.87**	**2.51**	**3.86**
	C2	1.18	1.32	1.55	2.00	**2.15**	3.22	3.92	5.19	6.63	8.89	**1.11**	**1.41**	**1.82**	**2.62**	**4.01**
	C3	**0.95**	1.12	1.45	2.20	2.58	2.72	3.67	5.10	6.45	8.63	1.43	**1.64**	**1.73**	**2.41**	**3.94**
ICFormer	C1	**1.05**	**1.11**	**1.29**	**1.61**	**2.27**	**1.40**	**1.70**	**2.42**	**3.25**	**4.18**	1.74	2.13	2.65	3.29	4.14
	C2	**1.00**	**1.06**	**1.23**	**1.54**	2.19	**1.35**	**1.72**	**2.54**	**2.48**	**4.49**	1.83	2.32	2.94	3.70	4.70
	C3	1.03	**1.08**	**1.25**	**1.58**	**2.27**	**1.52**	**1.77**	**2.46**	**3.29**	**4.22**	1.87	2.32	2.91	3.65	4.61
		200	400	600	800	1000	200	400	600	800	1000	200	400	600	800	1000

Table 3
RMSE results summary calculated as the average and the standard deviation of predictions in all cycles for all cells.

	RF	FNN	1DConv	DTW-CNN	ICFormer
Mean ± std	6.85 ± 1.31	4.20 ± 1.95	4.0 ± 1.86	3.12 ± 1.97	**2.33 ± 0.8**

result in similar IC curves or because the same degradation modes may have different effects depending on the cell configuration. An example of this limitation is provided in Appendix D.

It is worth noting that while these methods aimed to quantify the different degradation modes, they did not directly address knee predictions. While some methods attempting to predict knees are available in the literature, they are not directly comparable to our approach. For instance, the methods proposed in [7,66,67] necessitate the entire capacity loss trajectory, which differs from our approach of predicting them online. Other recent studies, such as [84,85] predict the capacity trajectory, which may be associated with the occurrence of knees. Others [10] combine knee identification with complete degradation curve estimation to achieve similar results. However, these approaches do not encompass any classification task nor address the online anticipation of knees. The data associated with the 1000 cycles of the three test cells included in this work could be used as a benchmark dataset for future methods to be compared against. For the approach proposed herein, precision, recall and F1-score are provided in Table 4 (see Appendix C for a detailed description of the metrics). The test cell data was divided into two classes with on one side the knee-less samples and on the other the samples with knees. The precision for the knee-less class indicates that the model 95% of the time was accurately predicting its absence. For the knee class, the precision was around 0.90. Similarly, the model achieved high recall scores for which it correctly identified around 97% of the knee-less class samples and 85% of the knee samples. The F1-scores, which provide a combined measure of precision and recall, were also high, with a score of 0.96 for the knee-less class and around 0.88 for the knee class.

Overall, these results indicate that our model is effective in predicting whether the degradation path of a battery has a knee in capacity loss. Furthermore, the 0.94 accuracy indicates that it correctly classified 94% of the samples. However, this is only a baseline and there is still work to be done in terms of improving the model's ability to identify the knee class. The lower metrics in this class are mainly due to the fact that many of the samples do not have knees but are labeled as such because their degradation conditions indicate that there would be if the battery were still cycled for a longer period. This type of situation can be easily confused with other trajectories that do not have a knee at any point of their cycle-life because the evolution of the IC curves is very similar, therefore the anticipation in these cases is complicated. As future work we intend to improve the knee labeling algorithm or propose a new metric to quantify these situations. In any case, it is important to note that the predictions are not erroneous in terms of knee identification because once a knee appears the model correctly identifies it, but it does fail sometimes in terms of anticipating the knee. While this is not of concern as the knee has not yet occurred, anticipating those knees have significant implications for battery management and some further improvements are needed.

4.3. A case of study: validation on real battery data

In this section, the proposed model's performance is evaluated on experimental battery data to demonstrate its practical applicability. The challenge of this task is compounded by the scarcity of datasets featuring real batteries exhibiting knees. To the best of our knowledge, only one dataset, as presented in [86], is available. However, this dataset lacks the necessary RPTs at C/25, rendering it incompatible with the data we used for training our model. Instead, two commercial high-power graphite/LFP cells manufactured by A123 Systems: CReal#1 (ANR26650M1a, 2.3 Ah) and CReal#2 (ANR26650M1b, 2.5 Ah) were used. Both cells were subjected to an electric vehicle driving scheme test, i.e., dynamic stress testing (DST) during their duty cycling schemes, and exhibited knees in capacity loss. To provide the thermodynamic conditions from which the aging modes can be estimated, sequential RPTs with included C/25 cycles were carried out. While CReal#1 was previously studied in detail in [60], this paper presents data on CReal#2 for the first time. The testing procedures for both cells were identical, except for the DST schedule, which was scaled from 400 W/kg for CReal#1 to 800 W/kg for CReal#2. The testing equipment, data sampling, testing temperatures, and RPTs that include C/25 charge/discharge cycles were the same for both cells, and were described in [60]. The raw experimental data is shown in Fig. D.10.

The results for CReal#1 during C/25 charging are presented in the first row of Fig. 7. The algorithm analyzed the C/25 IC curve history from the processed data (left) and provided knee predictions (center) and degradation mode estimations (right). Our model accurately predicted the onset of the knee point in the battery aging trajectory in advance when evaluating the first 800 cycles (see "YES" label in Fig. 7b). In terms of degradation modes evaluation, the model predicted an overall LAM_{NE} rate larger than that of LLI throughout cycling without LAM_{PE}. Furthermore, the LLI exhibited two distinct evolution rates, with an inflection point at the knee cycle (i.e., within cycle 800–1000) where the degradation rate increases likely because of irreversible plating. From a qualitative perspective, our model provided support for the conceptual interpretations made in Section 2 that a knee will eventually occur when $LAM_{NE} > LLI$, as the model output

Table 4
Classification results for knee identification for the tested cells.

	C1				C2				C3			
	precision	recall	F1-score	Accuracy	precision	recall	F1-score	Accuracy	precision	recall	F1-score	Accuracy
knee-less	0.95	0.97	0.96		0.95	0.98	0.96		0.95	0.97	0.96	
knee	0.90	0.85	0.87		0.93	0.85	0.89		0.90	0.83	0.87	
				0.94				0.95				0.94

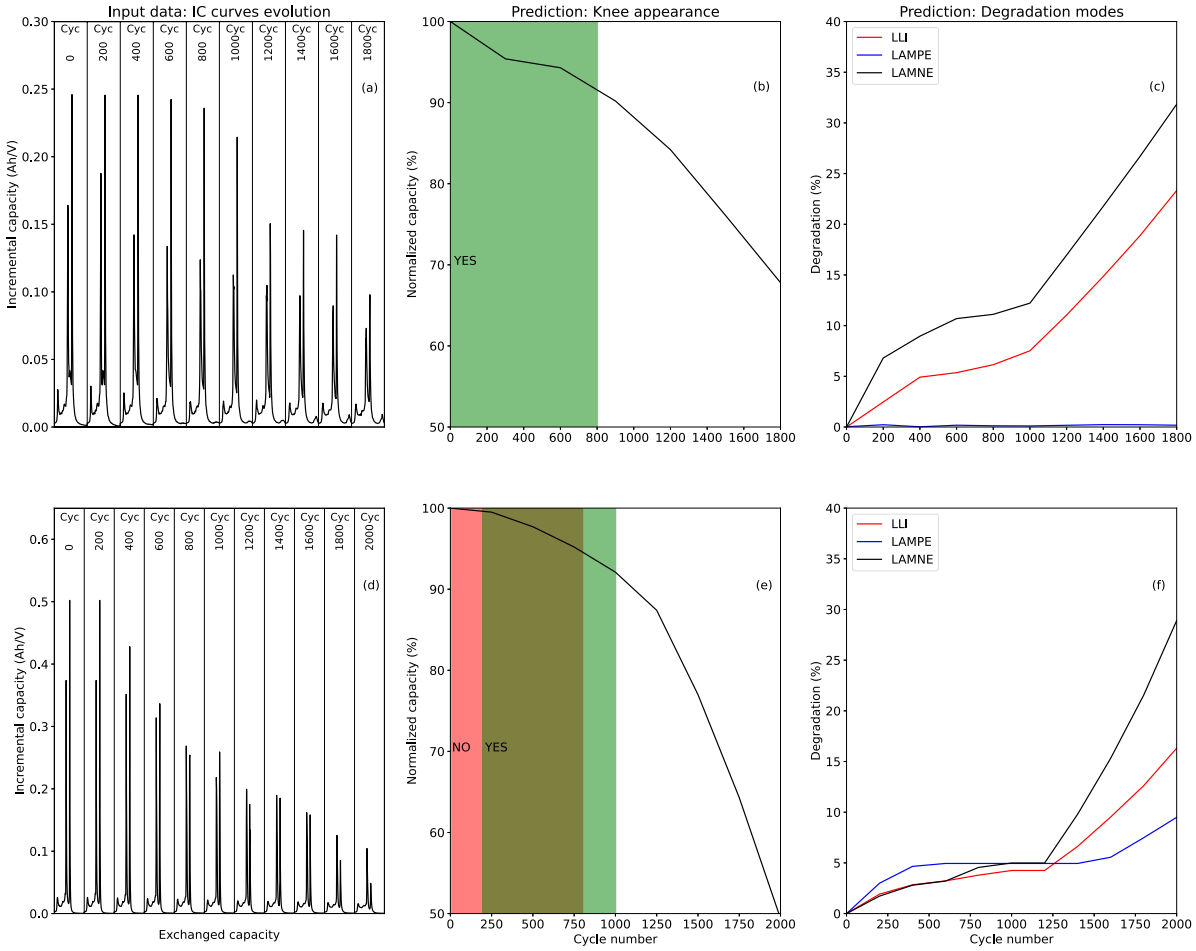

Fig. 7. Model performance validation for the experimental testing results from CReal#1 and CReal#2. The left column displays the IC curves derived from the experimental testing data at C/25, while the two adjacent columns show the model's predictions, including knee identification and estimated degradation modes. Note that predictions for LAM_{PE} for CReal#1 are hardly visible given the model predicted values near zero for all data points.

predicted. Additionally, the experimental IC curves exhibited a relative increase in Peak 1 and the appearance of Peak 0 (see Fig. D.10, left) which is a documented sign of LAM_{NE} and plating [32,60,87,88].

For a quantitative analysis, the model's predicted degradation modes were compared with the diagnosis reported in our previous study [60], Fig. D.11 and their evolutions were similar. The model predicted a final degradation for LAM_{NE} of 32% and LLI of 22% which is within a 3% difference from previous study's diagnosis. Similarly, both analyses reveal that the degradation trajectory for LLI is divided into two well-defined stages, with the second stage exhibiting a steeper rate. This changing rate of LLI occurs at the knee point and is caused by the occurrence of both irreversible lithium plating [60]. The degradation rate of LLI before and after the knee point is about a factor of four increase which also matches the results in previous studies. The model therefore provided a useful quantitative estimation of the detrimental effect of lithium plating. For LAM_{NE}, the overall evolution trend also matched quite well, with some deviations during the first cycles, but ultimately resulting in the same final degradation percentage.

The results for CReal#2 during C/25 charging are presented in the second row of Fig. 7. The model again predicted the onset of knee appearance when analyzing the data between cycles 200 and 1000, which is confirmed by a knee occurring at cycle 1250 (see Fig. 7e). Regarding the degradation modes prediction (cf. Fig. 7f), the model clearly distinguishes two stages with the degradation rates, i.e., before and after the onset of knee appearance. The degradation modes before knee appearance are within the same rate, although LAM_{PE} appears more prominent during the first 500 cycles and plateaus thereafter, while LAM_{NE} rate slightly increases above LLI at around cycle 750. At the knee occurrence point (i.e., cycle 1250), both LAM_{NE} and LLI rates increase substantially, while LAM_{PE} only slightly increases towards the end of cycling. The comprehensive analysis of CReal#2 is under preparation and will be published elsewhere. In the meantime, interested readers can access the data in the following public repository https://github.com/NahuelCostaCortez/ICFormer. A qualitative analysis of the degradation modes reveals that during the first 500 cycles, capacity fade on the full cell is solely caused by LLI, despite LAM_{PE} > LLI (see Fig. 7f). Although the degradation rate for LAM_{PE} almost doubles those of LLI, the initial offset on the PE (OFS_{ini}) acts as a buffer, which typically accounts for 10%–15% of total capacity for this

type of cell [89]. Therefore, degradation on the PE does not directly impact the full cell, and no knee in the capacity trajectory is shown or detected by the algorithm. At cycle 750, a change in the trajectory of the degradation modes is observed as LAM_{NE} becoming higher than LLI (see Fig. 7f). Once again, degradation on LAM_{NE} remains silent on cell capacity, in this case thanks to the initial excess acting as a buffer. However, since the NE excess is minor in this technology, typically NE_{EX} < 5% [89,90], even though both degradation modes are similar, NE shrinkage eventually catches up and yields to thermodynamic lithium plating, leading to a knee in capacity evolution and an acceleration rate of both degradation modes LAM_{NE} and LLI, as observed in Figs. 7d,f. Again, the inner degradation and its complexity remains transparent to the final user, as the model accurately predicted and indicated with a simple "YES" the appearance of a knee well before it takes place.

Compared to CRreal#1, CReal#2 does not display the appearance of Peak 0, and Peak 2 height remains higher than that of Peak 1 in the last cycles, while the opposite is true for CReal#1 (see Fig. D.10). The difference in the degradation modes observed in both cells can be attributed to the slight variation in cell configuration and, more importantly, to the difference in cycling scheme power, where CReal#2 was cycled at twice the power rate (i.e., 800 W/kg) with maximum currents reaching 8C-rate during discharge and 4C during charge. This could be the main reason why CReal#2 exhibits LAM_{PE} while CReal#1 shows no degradation on the LFP material. It is essential to note that online quantifying LFP degradation in this type of technology is challenging [50,61,90] and its true extent is usually only obtainable in post-mortem studies [89,91,92]. Therefore, the detection of LAM_{PE} with our model is a promising development for accurate online diagnosis and prognosis even under most challenging scenarios. Moreover, it is noteworthy that the algorithm can detect the appearance of a knee caused by thermodynamic lithium plating without the appearance of Peak 0, as shown in CReal#2 (and contrary to CReal#1, where Peak 0 appears on the IC curves). This result also indicates that lithium plating under this scheme appears to be almost 100% irreversible leaving no visible traces on Peak 0 of the IC curves during charge. Finally, our testing and simulation results suggest that the new version of the tested cell, CReal#2 (M1b), is indeed an improved new version. Despite doubling the power testing scheme, the capacity fade at 80% of CReal#2 remains nearly as that of the older version, CReal#1 (M1a), while the knee appearance is delayed by at least 250 cycles.

It should also be acknowledged that despite the notable advantages of our approach in modeling changes in IC curves when applying it directly to real data some preprocessing steps are required to ensure optimal performance. The data was normalized in the same way as was done with the training dataset and an interpolation was performed in those cycles where data was not available to meet the algorithm's requirement of having an IC curve every 200 cycles. Thus, we successfully aligned the real data with the expected format, enabling the application of our model. In future work, we intend to further refine our methodology and explore additional approaches to narrow the gap between real and simulated curves.

5. Concluding remarks and future work

Accurate diagnosis and knee detection of LIB degradation trajectories are critical for ensuring safe and reliable use in real-world applications. However, accurate LIB degradation modes estimation and prediction of knees pose notable challenges, highlighting the need for a dedicated and comprehensive model framework. Degradation modes tracking can provide a direct link to knee formation, and experimental techniques such as ICA have proven effective for quantifying them. In this paper, we highlighted some limitations of the application of Deep Learning algorithms to experimental data where current methods focus on independent diagnostics of the battery at different cycling stages, rather than analyzing the complete degradation trajectory, which could be an advantage in the automation of ICA. To this end, we presented a novel Transformer-based Deep Learning model that leverages self-attention to identify relevant changes in the evolution of the IC curves through cycling. The model was trained with synthetic data from the HNEI prognosis dataset and shows that it was able not only to identify knees but also to anticipate them while also outperforming state-of-the-art approaches in the estimation of degradation modes with RMSE errors around 2% on average for 1000 duty cycles for 3 different cells. We also included a report on the application of the model to experimental data from commercial batteries to illustrate its performance in a real-world scenario.

All in all, the model demonstrates that it meets the proposed expectations. The information it provides is of critical importance in many real-world applications, as it enables proactive maintenance, helps prevent unexpected failures and downtime, and allows for potential second-life use of the battery in less demanding applications. The model offers an informed prediction of the battery health from which different types of users can benefit. On the one hand, knee identification can satisfy those who simply want to know whether the battery should be replaced or proactive maintenance should be taken on it. On the other, the model provides an accurate estimation of degradation modes that offers detailed insight into the electrochemical performance of the battery, which provides valuable information that can be used both by battery specialists and researchers.

Our work can be extended in multiple ways. First, the data used still has room for improvement. The motivation for the use of IC curves is driven by the current gap in terms of degradation history analysis, but this opens new opportunities to monitor other types of variables that may also have an influence on the discovery of knees. It is known that LAM failures can lead to knees, but as discussed in the paper, not all knees can be identified in this way. In this sense, including other variables such as impedance or coulombic efficiency may provide more information regarding this type of defect. In addition, the dataset used assumes a C/25 charge rate, whereas it would be of interest to analyze the impact on the predictions in variable rates, and also with noise in the measurements and with a narrower characterization voltage window, which can be doable with synthetic data. For these reasons, we believe that the choice of a Transformer for our architecture is promising since it will allow us to scale easily in the future as it is designed to support more variables as inputs for analysis. Moreover, the potential of the model can be extrapolated to other tasks beyond knee identification and prediction of degradation modes such as estimation of RUL or prediction of future capacity.

Also, an ablation study on the attention weights may provide valuable insights into the state of degradation of the batteries. Finally, there are still some limitations on applying models trained on synthetic data to real data, as they have significant differences that may limit their integration in real environments. We leave the exploration of these open directions to future work.

CRediT authorship contribution statement

N. Costa: Conceptualization, Methodology, Software, Validation, Formal analysis, Data curation, Writing – original draft, Writing – review & editing. **D. Anseán:** Validation, Investigation, Resources, Writing – original draft, Writing – review & editing, Project administration, Funding acquisition. **M. Dubarry:** Writing – review & editing. **L. Sánchez:** Resources, Supervision, Project administration, Funding acquisition.

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.

Data availability

The manuscript provides the link to the data/code.

Acknowledgments

This work has been partially supported by the Ministry of Economy, Industry and Competitiveness (“Ministerio de Economía, Industria y Competitividad”) from Spain/FEDER under grants PID2020-112726-RB-I00 and PID2022-141792OB-I00, and by Principado de Asturias, grant SV-PA-21-AYUD/2021/50994. M.D. was funded by the Office of Naval Research (ONR), grant number N00014-19-1-2159.

Code availability

All models and experiments were implemented in TensorFlow [93]. Further details regarding the experimental setup, the model and the source code to reproduce the experimental results are available in the following public git repository: https://github.com/NahuelCostaCortez/ICFormer.

Appendix A. Knee labeling

The knee labeling strategy proposed in this work is as follows: Assuming knowledge of the capacity loss curve profile and the associated degradation modes for each duty cycle, as is the case with the dataset utilized in this study [49], first, the algorithm checks on the linearity for the capacity loss. For the non-linear duty cycles, each cycle is examined to determine if the primary degradation mode was LLI and if the knee occurs within the history up to that cycle (through another linear fit). For those cases, it is implied that the knee was caused by exponential LLI degradation, making knee anticipation unfeasible due to the immediate impact on capacity loss. Conversely, if the primary degradation mode was LAM and a knee occurs in a subsequent cycle, it indicates the presence of a latent effect that will eventually lead to a knee. Consequently, knee anticipation is possible before the knee fully manifests. The corresponding cycle is labeled accordingly, enabling the algorithm to identify knee occurrences and anticipate them.

Algorithm 1 Knee Labeling

```
kneeLabels ← []
for sample in dataset do
    kneeLabelsSample ← []    ▷ in each cycle is marked whether or not a knee is present
    if knee in sample then    ▷ if there is knee, check from which cycle to label it
        for cycle in sample do
            flag_knee ← False
            if mainDegradationMode == LLI and knee in QLossUpToCycle then
                flag_knee ← True    ▷ there is a knee caused by exponential LLI
            else
                if knee in FutureQLoss then
                    flag_knee ← True    ▷ there is going to be a knee caused by a silent mode (LAM)
                end if
            end if
            kneeLabelsSample[cycle] ← flag_knee
        end for
    end if
    kneeLabels[sample] ← kneeLabelsSample
end for
```

Appendix B. Model details and hyperparameters

The model has three main components: the positional encoding layer, the Transformer encoder and the MLP layers. The positional encoding layer is used to introduce the notion of order of the input sequence to the model. The original Transformer used fixed positional embeddings, combining sine and cosine functions and did not change during training. However, more recent Transformer implementations use learnable positional embeddings, which are learned during training along with other model parameters. This allows the model to adapt the embeddings to the specific task and dataset, rather than using fixed embeddings that may not be optimal [94]. In this work, a 5×128 shape tensor containing IC curves representing the degradation history of cells was fed into a learnable positional embedding layer with dimension h, resulting in a $5 \times$h-dimensional tensor. 5 was the number of IC curves chosen for each sample (that correspond to a 800-cycle window) and 128 was the size of the IC curves. This tensor was fed to a standard Transformer encoder, which in turn is made up of several encoders that contain the attention layers.

The formula for attention is as follows:

$$\text{Attention}(Q, K, V) = \text{softmax}(\frac{QK^T}{\sqrt{d_k}})V$$

Q (Query), K (Key) and V (Value) are three vectors derived from the positional embeddings using fully connected layers that are modified during training, and each serves a different purpose in the attention calculation. The attention mechanism starts by computing the similarity scores between the Q and K vectors. This is done by taking the dot product of the Q vector with each K vector. To mitigate the effects of large values in the dot product, the result is divided by the square root of the dimension of the Key vector (d_k). This scaling helps stabilize the gradients during training. The scaled similarity scores are then passed through a softmax function, which normalizes the scores and converts them into a probability distribution. The softmax operation ensures that the attention weights sum up to 1, allowing the model to allocate proper importance to different positions in the input sequence. Finally, the softmax probabilities are used to compute a weighted sum of the V vectors. The weights obtained from the softmax represent the importance assigned to each V vector. The weighted sum generates the context vector, which captures the relevant information from the V vectors based on the attention scores.

Each encoder consists of an attention layer, whose output is a tensor with shape 5×128 representing the attended parts of the IC curves, and a fully connected layer, whose output results in a h-dimensional tensor. The output of the last encoder layer was fed to two MLP heads, one that returns a 5 x 3-dimensional tensor containing the prediction of the degradation modes (LLI, LAM_{PE} and LAM_{NE}) for the five ICs curves and to another that returns the output of a sigmoid that predicts whether or not a knee exists.

The hyperparameters taken into account and their final values were:

- batch_size: 64
- learning_rate: 0.001
- num_encoders: 4
- ff_dim: 28
- mlp_units: 128
- mlp_dropout: 0.2
- dropout: 0.2

Appendix C. Classification metrics

Metrics used for the binary classification problem of knee identification. The classes are “knee-less” (positive class) and “knee” (negative class).

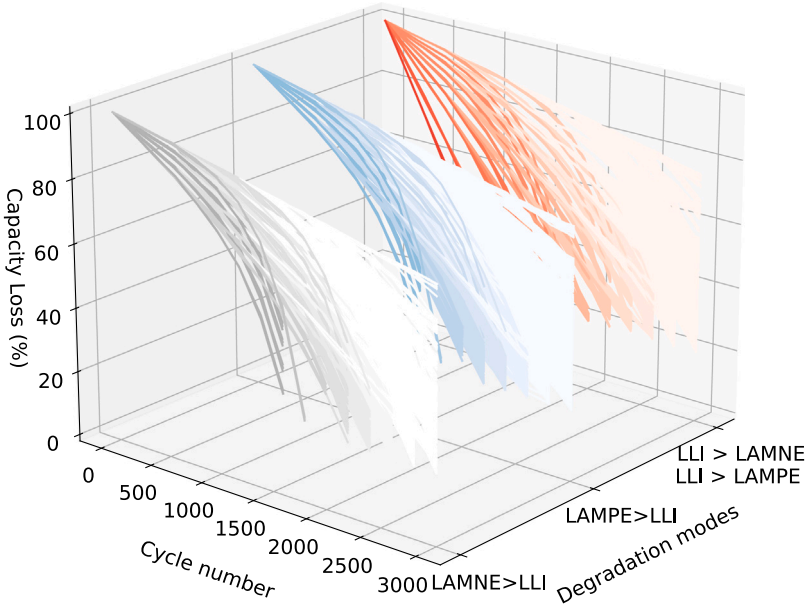

Fig. D.8. Effect of the degradation modes at different cycling stages on the appearance of knees. It is observed that predominant LAM modes (i.e., LAMs > LLI) lead to knees, and these appear earlier the higher the intensity of degradation, while LLI generally does not produce knees unless its evolution is highly exponential.

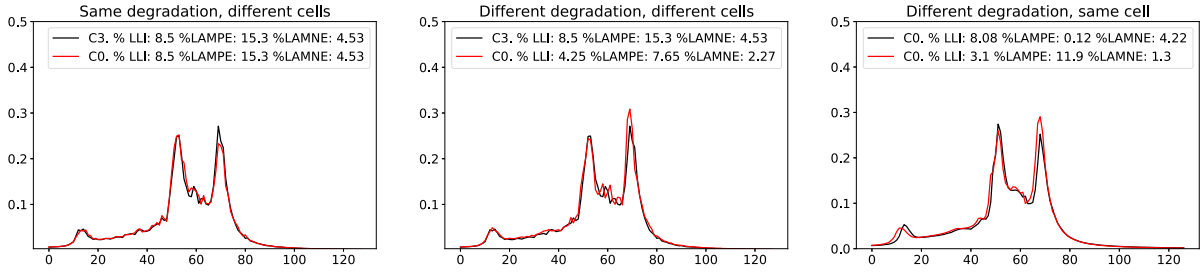

Fig. D.9. Similarities between IC curves for different cell configurations and degradation modes.

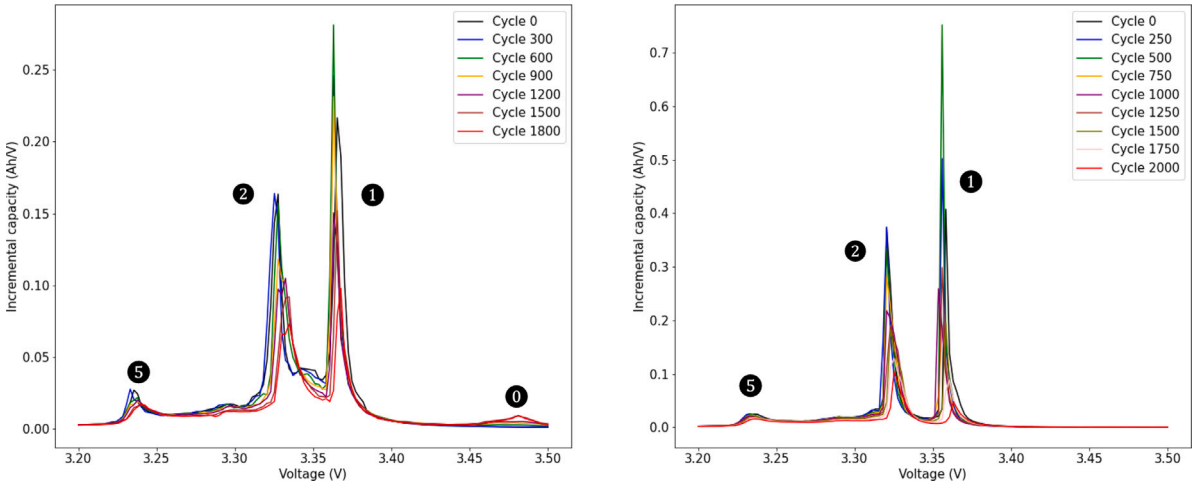

Fig. D.10. Experimental data for CReal #1 (left) and CReal #2 (right).

Precision: measures the proportion of correctly predicted positive instances out of the total predicted positive instances. It provides an indication of the model's ability to avoid false positives.

$$\text{Precision} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}}$$

Recall: also known as sensitivity or true positive rate, measures the proportion of correctly predicted positive instances out of the total actual positive instances. It provides an indication of the model's ability to avoid false negatives.

$$\text{Recall} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}}$$

The F1 score: combines precision and recall into a single value. It is the harmonic mean of precision and recall, providing a balanced

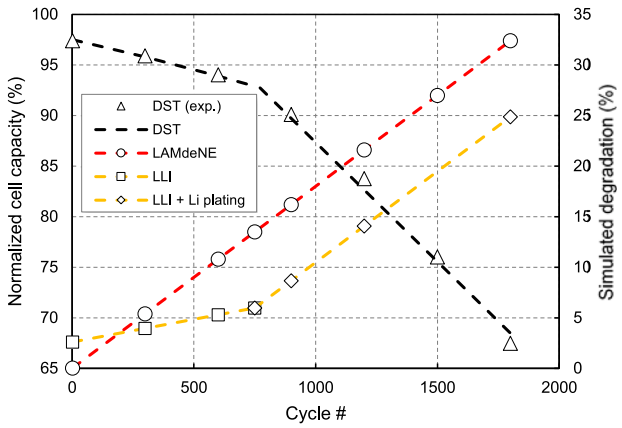

Fig. D.11. Analysis of cell degradation evolution for CReal #1 as estimated in [60]. This figure serves as a comparison to the results presented in Fig. 7.

evaluation of a classifier's performance.

$$\text{F1 Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

Accuracy: measures the proportion of correctly predicted instances (both positive and negative) out of the total instances. It provides an overall assessment of the model's correctness.

$$\text{Accuracy} = \frac{\text{True Positives} + \text{True Negatives}}{\text{True Positives} + \text{True Negatives} + \text{False Positives} + \text{False Negatives}}$$

Appendix D. Supplementary figures

Fig. D.9 illustrates a limitation of analyzing each IC curve independently. For instance, consider the curves on the left side. Due to variations in cell configurations (C0 referring to the training configuration and C3 referring to test cell 3), the resulting IC curves exhibit slight differences and this could prove problematic as these differences could be interpreted as a degradation. Similarly, the center and right figures showcase extremely similar IC curves for different combinations of the degradation modes. As a result, methods that diagnose each IC curve independently are prone to prediction errors since they were trained with signatures of IC curves under different degradation modes combinations. Consequently, when presented with a single IC curve, these methods may predict a plausible degradation mode that, within the context of the cell's degradation history, is not entirely accurate or relevant. This explains the high errors in LAM_{PE} in Table 2 and the more distributed errors in our approach, given that it takes into account the complete degradation history.

References

[1] M. Arbabzadeh, R. Sioshansi, J.X. Johnson, G.A. Keoleian, The role of energy storage in deep decarbonization of electricity production, Nature Commun. 10 (1) (2019) URL www.scopus.com.

[2] J. Deng, C. Bae, A. Denlinger, T. Miller, Electric vehicles batteries: requirements and challenges, Joule 4 (3) (2020) 511–515.

[3] G. Zubi, R. Dufo-López, M. Carvalho, G. Pasaoglu, The lithium-ion battery: State of the art and future perspectives, Renew. Sustain. Energy Rev. 89 (2018) 292–308.

[4] M. Dubarry, G. Baure, D. Anseán, Perspective on state-of-health determination in lithium-ion batteries, J. Electrochem. Energy Convers. Storage 17 (4) (2020) URL www.scopus.com.

[5] K.L. Gering, S.V. Sazhin, D.K. Jamison, C.J. Michelbacher, B.Y. Liaw, M. Dubarry, M. Cugnet, Investigation of path dependence in commercial lithium-ion cells chosen for plug-in hybrid vehicle duty cycle protocols, J. Power Sources 196 (7) (2011) 3395–3403.

[6] T. Raj, A.A. Wang, C.W. Monroe, D.A. Howey, Investigation of path-dependent degradation in lithium-ion batteries, Batter. Supercaps 3 (12) (2020) 1377–1385.

[7] W. Diao, S. Saxena, B. Han, M. Pecht, Algorithm to determine the knee point on capacity fade curves of lithium-ion cells, Energies 12 (15) (2019) http://dx.doi.org/10.3390/en12152910, URL https://www.mdpi.com/1996-1073/12/15/2910.

[8] P. Fermín-Cueto, E. McTurk, M. Allerhand, E. Medina-Lopez, M.F. Anjos, J. Sylvester, G. dos Reis, Identification and machine learning prediction of knee-point and knee-onset in capacity degradation curves of lithium-ion cells, Energy AI 1 (2020) 100006, http://dx.doi.org/10.1016/j.egyai.2020.100006, URL https://www.sciencedirect.com/science/article/pii/S2666546820300069.

[9] P.M. Attia, A. Bills, F.B. Planella, P. Dechent, G. Dos Reis, M. Dubarry, P. Gasper, R. Gilchrist, S. Greenbank, D. Howey, et al., "Knees" in lithium-ion battery aging trajectories, J. Electrochem. Soc. 169 (6) (2022) 060517.

[10] M. Haris, M.N. Hasan, S. Qin, Degradation curve prediction of lithium-ion batteries based on knee point detection algorithm and convolutional neural network, IEEE Trans. Instrum. Meas. 71 (2022) 1–10.

[11] Y. Ma, M. Yao, H. Liu, Z. Tang, State of health estimation and remaining useful life prediction for lithium-ion batteries by improved particle swarm optimization-back propagation neural network, J. Energy Storage 52 (2022) 104750.

[12] Z. Tong, J. Miao, S. Tong, Y. Lu, Early prediction of remaining useful life for Lithium-ion batteries based on a hybrid machine learning method, J. Clean. Prod. 317 (2021) 128265.

[13] Y. Yang, A machine-learning prediction method of lithium-ion battery life based on charge process for different applications, Appl. Energy 292 (2021) 116897.

[14] G. Ma, Y. Zhang, C. Cheng, B. Zhou, P. Hu, Y. Yuan, Remaining useful life prediction of lithium-ion batteries based on false nearest neighbors and a hybrid neural network, Appl. Energy 253 (2019) 113626.

[15] J. Liu, Z. Chen, Remaining useful life prediction of lithium-ion batteries based on health indicator and Gaussian process regression model, IEEE Access 7 (2019) 39474–39484.

[16] T.F. Fuller, M. Doyle, J. Newman, Simulation and optimization of the dual lithium ion insertion cell, J. Electrochem. Soc. 141 (1) (1994) 1.

[17] F.B. Planella, W. Ai, A. Boyce, A. Ghosh, I. Korotkin, S. Sahu, V. Sulzer, R. Timms, T. Tranter, M. Zyskin, et al., A continuum of physics-based lithium-ion battery models reviewed, Prog. Energy (2022).

[18] M. Dubarry, C. Truchot, B.Y. Liaw, Synthesize battery degradation modes via a diagnostic and prognostic model, J. Power Sources 219 (2012) 204–216.

[19] C.R. Birkl, M.R. Roberts, E. McTurk, P.G. Bruce, D.A. Howey, Degradation diagnostics for lithium ion cells, J. Power Sources 341 (2017) 373–386, http://dx.doi.org/10.1016/j.jpowsour.2016.12.011, URL https://www.sciencedirect.com/science/article/pii/S0378775316316998.

[20] Y. Li, K. Liu, A.M. Foley, A. Zülke, M. Berecibar, E. Nanini-Maury, J. Van Mierlo, H.E. Hoster, Data-driven health estimation and lifetime prediction of lithium-ion batteries: A review, Renew. Sustain. Energy Rev. 113 (2019) 109254.

[21] H. Meng, Y.-F. Li, A review on prognostics and health management (PHM) methods of lithium-ion batteries, Renew. Sustain. Energy Rev. 116 (2019) 109405.

[22] K. Luo, X. Chen, H. Zheng, Z. Shi, A review of deep learning approach to predicting the state of health and state of charge of lithium-ion batteries, J. Energy Chem. (2022).

[23] Y. Zhang, Y.-F. Li, Prognostics and health management of Lithium-ion battery using deep learning methods: A review, Renew. Sustain. Energy Rev. 161 (2022) 112282.

[24] M.H. Lipu, M. Hannan, A. Hussain, M. Hoque, P.J. Ker, M.M. Saad, A. Ayob, A review of state of health and remaining useful life estimation methods for lithium-ion battery in electric vehicles: Challenges and recommendations, J. Clean. Prod. 205 (2018) 115–133.

[25] S. Wang, S. Jin, D. Bai, Y. Fan, H. Shi, C. Fernandez, A critical review of improved deep learning methods for the remaining useful life prediction of lithium-ion batteries, Energy Rep. 7 (2021) 5562–5574.

[26] A.G. Li, A.C. West, M. Preindl, Towards unified machine learning characterization of lithium-ion battery degradation across multiple levels: A critical review, Appl. Energy 316 (2022) 119030.

[27] X. Sui, S. He, S.B. Vilsen, J. Meng, R. Teodorescu, D.-I. Stroe, A review of non-probabilistic machine learning-based state of health estimation techniques for Lithium-ion battery, Appl. Energy 300 (2021) 117346.

[28] A. Shah, K. Shah, C. Shah, M. Shah, State of charge, remaining useful life and knee point estimation based on artificial intelligence and machine learning in lithium-ion EV batteries: A comprehensive review, Renew. Energy Focus 42 (2022) 146–164.

[29] W. Li, N. Sengupta, P. Dechent, D. Howey, A. Annaswamy, D.U. Sauer, One-shot battery degradation trajectory prediction with deep learning, J. Power Sources 506 (2021) 230024.

[30] M.L. Baptista, E.M. Henriques, K. Goebel, More effective prognostics with elbow point detection and deep learning, Mech. Syst. Signal Process. 146 (2021) 106987.

[31] S. Sohn, H.-E. Byun, J.H. Lee, Two-stage deep learning for online prediction of knee-point in Li-ion battery capacity degradation, Appl. Energy 328 (2022) 120204.

[32] M. Dubarry, D. Anseán, Best practices for incremental capacity analysis, Front. Energy Res. 10 (2022) http://dx.doi.org/10.3389/fenrg.2022.1023555, URL https://www.frontiersin.org/articles/10.3389/fenrg.2022.1023555.

[33] M. Dubarry, G. Baure, Perspective on commercial Li-ion battery testing, best practices for simple and effective protocols, Electronics 9 (1) (2020) 152.

[34] K.S. Mayilvahanan, K.J. Takeuchi, E.S. Takeuchi, A.C. Marschilok, A.C. West, Supervised learning of synthetic big data for Li-ion battery degradation diagnosis, Batter. Supercaps (2021).

[35] S. Kim, Z. Yi, B.-R. Chen, T.R. Tanim, E.J. Dufek, Rapid failure mode classification and quantification in batteries: A deep learning modeling framework, Energy Storage Mater. 45 (2022) 1002–1011.

[36] S. Lee, Y. Kim, Li-ion battery electrode health diagnostics using machine learning, in: 2020 American Control Conference, ACC, IEEE, 2020, pp. 1137–1142.

[37] N. Costa, L. Sanchez, D. Ansean, M. Dubarry, Li-ion battery degradation modes diagnosis via convolutional neural networks, J. Energy Storage 55 (2022) 105558.

[38] L. Wildfeuer, M. Lienkamp, Quantifiability of inherent cell-to-cell variations of commercial lithium-ion batteries, eTransportation 9 (2021) 100129.

[39] P. Dechent, S. Greenbank, F. Hildenbrand, S. Jbabdi, D.U. Sauer, D.A. Howey, Estimation of Li-ion degradation test sample sizes required to understand cell-to-cell variability, Batter. Supercaps 4 (12) (2021) 1821–1829.

[40] M. Ank, T. Kröger, M. Schreiber, M. Lienkamp, Experimental analysis of lithium-ion cell procurement: Quality differences, correlations, and importance of cell characterization, J. Energy Storage 66 (2023) 107430.

[41] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, Adv. Neural Inf. Process. Syst. 30 (2017).

[42] T. Lin, Y. Wang, X. Liu, X. Qiu, A survey of transformers, AI Open (2022).

[43] Y. Tay, M. Dehghani, D. Bahri, D. Metzler, Efficient transformers: A survey, ACM Comput. Surv. 55 (6) (2022) 1–28.

[44] C. Birkl, Oxford Battery Degradation Dataset 1, University of Oxford, 2017.

[45] B. Bole, C. Kulkarni, M. Daigle, Randomized battery usage data set, in: NASA AMES Prognostics Data Repository, Vol. 70, 2014.

[46] W. He, N. Williard, M. Osterman, M. Pecht, Prognostics of lithium-ion batteries based on Dempster–Shafer theory and the Bayesian Monte Carlo method, J. Power Sources 196 (23) (2011) 10314–10321.

[47] V.D. Angelis, Battery archive, 2022, URL www.batteryarchive.org.

[48] L. Ward, S. Babinec, E.J. Dufek, D.A. Howey, V. Viswanathan, M. Aykol, D.A. Beck, B. Blaiszik, B.-R. Chen, G. Crabtree, et al., Principles of the battery data genome, 2021, arXiv preprint arXiv:2109.07278.

[49] M. Dubarry, D. Beck, Big data training data for artificial intelligence-based Li-ion diagnosis and prognosis, J. Power Sources 479 (2020) 228806.

[50] X. Han, L. Lu, Y. Zheng, X. Feng, Z. Li, J. Li, M. Ouyang, A review on the key issues of the lithium ion battery degradation among the whole life cycle, eTransportation 1 (2019) 100005.

[51] J.S. Edge, S. O'Kane, R. Prosser, N.D. Kirkaldy, A.N. Patel, A. Hales, A. Ghosh, W. Ai, J. Chen, J. Yang, et al., Lithium ion battery degradation: what you need to know, Phys. Chem. Chem. Phys. 23 (14) (2021) 8200–8221.

[52] X.-G. Yang, Y. Leng, G. Zhang, S. Ge, C.-Y. Wang, Modeling of lithium plating induced aging of lithium-ion batteries: Transition from linear to nonlinear aging, J. Power Sources 360 (2017) 28–40.

[53] T. Waldmann, B.-I. Hogg, M. Wohlfahrt-Mehrens, Li plating as unwanted side reaction in commercial Li-ion cells – A review, J. Power Sources 384 (2018) 107–124, http://dx.doi.org/10.1016/j.jpowsour.2018.02.063, URL https://www.sciencedirect.com/science/article/pii/S0378775318301848.

[54] S.J. An, J. Li, C. Daniel, D. Mohanty, S. Nagpure, D.L. Wood, The state of understanding of the lithium-ion-battery graphite solid electrolyte interphase (SEI) and its relationship to formation cycling, Carbon 105 (2016) 52–76, http://dx.doi.org/10.1016/j.carbon.2016.04.008, URL https://www.sciencedirect.com/science/article/pii/S0008622316302676.

[55] D.L. Wood, J. Li, S.J. An, Formation challenges of Lithium-ion battery manufacturing, Joule 3 (12) (2019) 2884–2888, http://dx.doi.org/10.1016/j.joule.2019.11.002.

[56] Y. Yang, J. Wang, S.C. Kim, W. Zhang, Y. Peng, P. Zhang, R.A. Vilá, Y. Ma, Y.K. Jeong, Y. Cui, In situ prelithiation by direct integration of lithium Mesh into battery cells, Nano Lett. (2023).

[57] Q.A. Nguyen, A.K. Haridas, T. Terlier, S.L. Biswal, Prelithiation effects in enhancing silicon-based anodes for full-cell Lithium-ion batteries using stabilized lithium metal particles, ACS Appl. Energy Mater. 6 (10) (2023) 5567–5579.

[58] M. Tang, P. Albertus, J. Newman, Two-dimensional modeling of lithium deposition during cell charging, J. Electrochem. Soc. 156 (5) (2009) A390, http://dx.doi.org/10.1149/1.3095513.

[59] I. Zilberman, J. Sturm, A. Jossen, Reversible self-discharge and calendar aging of 18650 Nickel-Rich, silicon-graphite lithium-ion cells, J. Power Sources 425 (2019) 217–226, http://dx.doi.org/10.1016/j.jpowsour.2019.03.109.

[60] D. Anseán, M. Dubarry, A. Devie, B. Liaw, V. García, J. Viera, M. González, Operando lithium plating quantification and early detection of a commercial LiFePO4 cell cycled under dynamic driving schedule, J. Power Sources 356 (2017) 36–46.

[61] X. Han, M. Ouyang, L. Lu, J. Li, Y. Zheng, Z. Li, A comparative study of commercial lithium ion battery cycle life in electrical vehicle: Aging mechanism identification, J. Power Sources 251 (2014) 38–54.

[62] M. Dubarry, B.Y. Liaw, Identify capacity fading mechanism in a commercial LiFePO4 cell, J. Power Sources 194 (1) (2009) 541–549.

[63] M. Dubarry, C. Truchot, B.Y. Liaw, Cell degradation in commercial LiFePO4 cells with high-power and high-energy designs, J. Power Sources 258 (2014) 408–419.

[64] M. Dubarry, M. Berecibar, A. Devie, D. Anseán, N. Omar, I. Villarreal, State of health battery estimator enabling degradation diagnosis: Model and algorithm description, J. Power Sources 360 (2017) 59–69.

[65] W. Diao, S. Saxena, B. Han, M. Pecht, Algorithm to determine the knee point on capacity fade curves of lithium-ion cells, Energies 12 (15) (2019) 2910.

[66] P. Fermín-Cueto, E. McTurk, M. Allerhand, E. Medina-Lopez, M.F. Anjos, J. Sylvester, G. dos Reis, Identification and machine learning prediction of knee-point and knee-onset in capacity degradation curves of lithium-ion cells, Energy AI 1 (2020) 100006.

[67] C. Zhang, Y. Wang, Y. Gao, F. Wang, B. Mu, W. Zhang, Accelerated fading recognition for lithium-ion batteries with Nickel–Cobalt–Manganese cathode using quantile regression method, Appl. Energy 256 (2019) 113841.

[68] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever, et al., Improving Language Understanding by Generative Pre-Training, OpenAI, 2018.

[69] S. Reed, K. Zolna, E. Parisotto, S.G. Colmenarejo, A. Novikov, G. Barth-Maron, M. Gimenez, Y. Sulsky, J. Kay, J.T. Springenberg, et al., A generalist agent, 2022, arXiv preprint arXiv:2205.06175.

[70] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, et al., Highly accurate protein structure prediction with AlphaFold, Nature 596 (7873) (2021) 583–589.

[71] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, I. Sutskever, Zero-shot text-to-image generation, in: International Conference on Machine Learning, PMLR, 2021, pp. 8821–8831.

[72] M.A. Hannan, D.N. How, M.H. Lipu, M. Mansor, P.J. Ker, Z. Dong, K. Sahari, S.K. Tiong, K.M. Muttaqi, T.I. Mahlia, et al., Deep learning approach towards accurate state of charge estimation for lithium-ion batteries using self-supervised transformer model, Sci. Rep. 11 (1) (2021) 19541.

[73] D. Chen, W. Hong, X. Zhou, Transformer network for remaining useful life prediction of lithium-ion batteries, IEEE Access 10 (2022) 19621–19628.

[74] L. Biggio, T. Bendinelli, C. Kulkarni, O. Fink, Dynaformer: A deep learning model for ageing-aware battery discharge prediction, 2022, arXiv preprint arXiv:2206.02555.

[75] X. Gu, K. See, P. Li, K. Shan, Y. Wang, L. Zhao, K.C. Lim, N. Zhang, A novel state-of-health estimation for the lithium-ion battery using a convolutional neural network and transformer model, Energy 262 (2023) 125501.

[76] H. Shen, X. Zhou, Z. Wang, J. Wang, State of charge estimation for lithium-ion battery using transformer with immersion and invariance adaptive observer, J. Energy Storage 45 (2022) 103768.

[77] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al., An image is worth 16x16 words: Transformers for image recognition at scale, 2020, arXiv preprint arXiv:2010.11929.

[78] Y. Nie, N.H. Nguyen, P. Sinthong, J. Kalagnanam, A time series is worth 64 words: Long-term forecasting with transformers, 2022, arXiv preprint arXiv:2211.14730.

[79] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S.J. van der Walt, M. Brett, J. Wilson, K.J. Millman, N. Mayorov, A.R.J. Nelson, E. Jones, R. Kern, E. Larson, C.J. Carey, İ. Polat, Y. Feng, E.W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E.A. Quintero, C.R. Harris, A.M. Archibald, A.H. Ribeiro, F. Pedregosa, P. van Mulbregt, SciPy 1.0 Contributors, SciPy 1.0: Fundamental algorithms for scientific computing in Python, Nature Methods 17 (2020) 261–272, http://dx.doi.org/10.1038/s41592-019-0686-2.

[80] L. Biewald, Experiment tracking with weights and biases, 2020, Software available from wandb.com, URL https://www.wandb.com/.

[81] M. Dubarry, N. Costa, D. Matthews, Data-driven direct diagnosis of Li-ion batteries connected to photovoltaics, Nature Commun. 14 (1) (2023) 3138.

[82] K.S. Mayilvahanan, K.J. Takeuchi, E.S. Takeuchi, A.C. Marschilok, A.C. West, Supervised learning of synthetic big data for Li-ion battery degradation diagnosis, Batter. Supercaps 5 (1) (2022) e202100166.

[83] W.-J. Zhang, Structure and performance of LiFePO4 cathode materials: A review, J. Power Sources 196 (6) (2011) 2962–2970, http://dx.doi.org/10.1016/j.jpowsour.2010.11.113, URL https://www.sciencedirect.com/science/article/pii/S037877531002104X.

[84] Z. Chen, L. Chen, Z. Ma, K. Xu, Y. Zhou, W. Shen, Joint modeling for early predictions of Li-ion battery cycle life and degradation trajectory, Energy 277 (2023) 127633.

[85] K. Luo, H. Zheng, Z. Shi, A simple feature extraction method for estimating the whole life cycle state of health of lithium-ion batteries using transformer-based neural network, J. Power Sources 576 (2023) 233139.

[86] K.A. Severson, P.M. Attia, N. Jin, N. Perkins, B. Jiang, Z. Yang, M.H. Chen, M. Aykol, P.K. Herring, D. Fraggedakis, et al., Data-driven prediction of battery cycle life before capacity degradation, Nat. Energy 4 (5) (2019) 383–391.

[87] R.V. Bugga, M.C. Smart, Lithium plating behavior in lithium-ion cells, ECS Trans. 25 (36) (2010) 241.

[88] M. Petzl, M.A. Danzer, Nondestructive detection, characterization, and quantification of lithium plating in commercial lithium-ion batteries, J. Power Sources 254 (2014) 80–87.

[89] M. Klett, R. Eriksson, J. Groot, P. Svens, K.C. Högström, R.W. Lindström, H. Berg, T. Gustafson, G. Lindbergh, K. Edström, Non-uniform aging of cycled commercial LiFePO4//graphite cylindrical cells revealed by post-mortem analysis, J. Power Sources 257 (2014) 126–137.

[90] D. Anseán, M. Dubarry, A. Devie, B. Liaw, V. García, J. Viera, M. González, Fast charging technique for high power LiFePO4 batteries: A mechanistic analysis of aging, J. Power Sources 321 (2016) 201–209, http://dx.doi.org/10.1016/j.jpowsour.2016.04.140, URL https://www.sciencedirect.com/science/article/pii/S0378775316305249.

[91] M. Kassem, C. Delacourt, Postmortem analysis of calendar-aged graphite/LiFePO4 cells, J. Power Sources 235 (2013) 159–171.

[92] E. Sarasketa-Zabala, F. Aguesse, I. Villarreal, L. Rodriguez-Martinez, C.M. López, P. Kubiak, Understanding lithium inventory loss and sudden performance fade in cylindrical cells during cycling with deep-discharge steps, J. Phys. Chem. C 119 (2) (2015) 896–906.

[93] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G.S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, X. Zheng, TensorFlow: Large-scale machine learning on heterogeneous systems, 2015, Software available from tensorflow.org, URL https://www.tensorflow.org/.

[94] Q. Wen, T. Zhou, C. Zhang, W. Chen, Z. Ma, J. Yan, L. Sun, Transformers in time series: A survey, 2022, arXiv preprint arXiv:2202.07125.