SOCXAI: Leveraging CNN and SHAP Analysis for Battery SOC Estimation and Anomaly Detection

Amel Hidouri[1], Slimane Arbaoui[1], Ahmed Samet[1], Ali Ayadi[1], Tedjani Mesbahi[1], Romuald Boné[1], and François de Bertrand de Beuvron[1]

Université de Strasbourg, 67000, France
Institut National des Sciences Appliquées (INSA Strasbourg), 67000, France
CNRS, ICube Laboratory UMR 7357, Strasbourg, 67000, France
`first_name.last_name@insa-strasbourg.fr`

Abstract. In the domain of battery energy storage systems for Electric Vehicles (EVs) applications and beyond, the adoption of machine learning techniques has surfaced as a notable strategy for battery modeling. Machine learning models are primarily utilized for forecasting the forthcoming state of batteries, with a specific focus on analyzing the State-of-Charge (SOC). Additionally, these models are employed to assess the State-of-Health (SOH) and predict the Remaining Useful Life (RUL) of batteries. Moreover, offering clear explanations for abnormal battery usage behavior is crucial, empowering users with insights needed for informed decision-making, build trust in the system, and ultimately enhance overall satisfaction. This paper presents `SOCXAI`, a novel algorithm designed for precise estimation of batteries's SOC. Our proposed model utilizes a Convolutional Neural Network (CNN) architecture to efficiently estimate the twenty five future values of SOC, rather than a single value. We also incorporate a SHApley Additive exPlanations (SHAP)-based post-hoc explanation method into our method focusing on the current feature values for deeper prediction insights. Furthermore, to detect abnormal battery usage behavior, we employ a 2-dimensional matrix profile-based approach on the time series of current values and their corresponding SHAP values. This methodology facilitates the detection of discords, which indicate irregular patterns in the battery usage. Our extensive empirical evaluation, using diverse real-world benchmarks, demonstrates our approach effectiveness, showcasing its superiority over state-of-the-art algorithms.

Keywords: Machine Learning. Battery. Explainability. SOC estimation. Data Mining.

1 Introduction

The integration of technological advancements across industries has significantly enhanced the accessibility and the generation of industrial time series data, a

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63783-4_14

trend expected to persist with the emergence of Industry 4.0 [10]. This development has led to the inevitable generation of vast datasets, underscores the need for versatile methods to effectively mine this information. This transition involves the utilization of data mining techniques to extract valuable insights from large datasets, particularly within the industrial sector. A key challenge in data mining revolves around identifying significant sub-sequences within time series data, where "interesting" patterns may include both repetitive and singular occurrences or deviations from the norm. Meanwhile, battery management has become a focal point within this context. As batteries play a critical role in various applications, understanding their behavior through time series analysis is essential.

The viability of Electric vehicles (EVs) is predominantly contingent upon the performance, range, lifetime cost-effectiveness, and safety of their batteries. At present, rechargeable lithium-ion (Li-ion) batteries are the preferred choice for EVs due to their favorable energy density and lifespan. The high energy density of Li-ion batteries allows for more energy storage in a relatively compact size, which is crucial for maximizing the driving range of EVs [11].

Therefore, ensuring the efficiency and safety of these advanced batteries is becoming increasingly crucial. Effectively managing batteries within a system demands detailed modeling to accurately predict their condition, with particular focus on metrics such as State of Charge (SOC) and State of Health (SOH). These metrics offer crucial insights into remaining energy, power delivery capacity, and overall cell life. Nevertheless, assessing residual lithium in batteries is a challenging task, necessitating precise algorithms embedded within Battery Management Systems (BMS). These algorithms, often leveraging mathematical models or Machine Learning (ML) techniques, play a pivotal role in estimating the battery's states, including SOC and SOH levels, using data such as terminal voltage, terminal current, and surface temperature. These measures are useful comprehending the remaining driving range of an EV or in designing a battery that will exhibit optimal performance in real-world conditions. Often, in this work we are interested in the SOC of the battery within a single charge/discharge cycle. The SOC of a battery refers to the current level of energy stored in the battery, expressed as a percentage of its total capacity. It indicates how much charge is remaining in the battery relative to its fully charged state and it can be computed using the following formula (1) [7]:

$$\text{SOC (\%)} = \left(\frac{\text{Ongoing capacity}}{\text{Total capacity}} \right) \times 100 \tag{1}$$

Recently, there has been a noticeable shift towards employing data mining tools to facilitate eXplainable Artificial Intelligence (XAI). This involves utilizing data mining techniques to elucidate and interpret black-box models, thereby improving transparency and comprehensibility in the decision-making processes of these models, particularly in safety-critical applications, especially within industries or vehicles. As a result, the development of XAI techniques has become a priority, aiming to provide insights into AI decision-making processes and make

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63783-4_14

their outputs interpretable to end-users. In this context, SHAP, a model-agnostic approach [14] based on Shapely index, has gained significant popularity in recent years for explaining a wide range of ML models.

In the realm of industrial time series data analysis, examining subsequences for similarities or disimilarities provides valuable insights and explanations regarding the state of a product or process. Commonly utilized terms in literature to describe patterns within sequential data are time series motif and time series discord. Time series motifs predominantly emphasize similarities, while discords concentrate on dissimilarities. Mining time series data, particularly through discord identification, is a topic of extensive research. As a result, the field of time series anomaly detection has witnessed a remarkable surge in interest, with hundreds of algorithms proposed over the last two decades [1, 12].

Our contribution is twofold, aiming to address both the prediction task and the requirement for explainability concurrently. More specifically, this paper presents a novel approach leveraging Convolutional Neural Networks (CNN) to estimate the SOC of batteries. Unlike existing methods, the proposed model predicts 25 SOC values rather than a single value, thereby offering more detailed insights into battery behavior for the next minutes. To provide explanations for the SOC predictions, we employ the SHAP model, that analyzes the contributions of current values towards SOC predictions. This is complemented by the application of a two-dimensional anomaly detection model, enabling us to identify the factors influencing SOC predictions.

2 Related Work

Abundant literature has been dedicated to the task of SOC estimation. Indeed, tremendous progress has been made in developing efficient algorithms that can estimate its future state, i.e., SOC. Two branches of works can characterize existing battery models: model-based approaches and data-driven methods. The former consists of the equivalent circuit models (ECMs) which is based on empirical knowledge and experimental data. Batteries are represented by groups of electrical components, such as resistors and capacitors, forming resistor-capacitor networks that are used to monitor the battery's behavior at different time constants associated with the diffusion and charge-transfer processes [8, 9, 17]. Although, this model is used as main battery models that are widely used in the BMS of EVs for online SOC estimations due to their low computational demands, the accuracy is usually limited to the parameterized range of the model. A further improvement on model-based methods is about development of Physics-Based Models (PBMs) [3, 5], with the pseudo two-dimensional (P2D) model being the most notable. The P2D model provides insights into battery internal dynamics. However, managing its equations is complex and demands significant computational resources, making it impractical for real-time applications. Moreover, PBMs often overlook details about material information. The second line of research pertains to Data-Driven Models (DDMs), which have garnered considerable attention for their adaptability, model-free advantages, and the capacity

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63783-4_14

to handle high degrees of nonlinearity. Possessing self-learning capabilities and robust generalization ability, DDMs are particularly effective for estimating SOC within nonlinear systems. Typically, these systems are constructed using various machine learning techniques, including neural networks [2,6,21], support vector machines [22], to predict SOC without the necessity for a prior knowledge.

3 A Glimpse at Times Series

In this section, we introduce the definitions and concepts that we will use throughout this paper.

Definition 1 (Time Series). *A time series $T \in \mathcal{R}^n$ denoted as $T = [t_1, \ldots, t_n]$ is a time-ordered sequence of values.*

Definition 2 (Subsequence). *A subsequence $T_{i,m} \in \mathcal{R}^m$ of T is a continuous subset of values from T of length m, starting from position i. Formally, $T_{i,m} = [t_i, t_{i+1}, \ldots, t_{i+m-1}]$.*

By selecting any subsequence $T_{i,m}$ as a query and computing its distance from all subsequences within the time series T, then sequentially saving the distances in an array, we generate a distance profile.

Definition 3 (Distance Profile). *A distance profile D_i of a time series T is an ordered array of Euclidean distances between the query subsequence $T_{i,m}$ and all subsequences in time series T. Formally, $D_i = [d_{i,1}, d_{i,2}, \ldots, d_{i,n-m+1}]$ where $d_{i,j}$ for $i \geq 1, j \leq n-m+1$ is the Euclidean distance between $T_{i,m}$ and $T_{j,m}$.*

In the distance profile D_i of query $T_{i,m}$, the i^{th} position represents the distance between the query and itself, resulting in a value of 0. Values preceding and following position i are nearly zero, indicating overlapping subsequences with the query. We focus solely on non-self-matches, disregarding these self-matches.

Definition 4 (Non-Self Match). *In a time series T, with a subsequence $T_{p,m}$ of length m beginning at position p and a matching subsequence $T_{q,m}$ starting at q, $T_{p,m}$ is a non-self match to $T_{q,m}$ with distance $d_{p,q}$ if $|p - q| \geq m$.*

Definition 5 (Time Series Discord). *In time series T, with a subsequence $T_{d,m}$ of length m starting at position d is considered a discord of T if the distance between $T_{d,m}$ and its nearest non-self match is the largest among all subsequences. Formally, for every $T_{c,m} \in T$, with the non-self matching sets M_D of $T_{d,m}$, and non-self matching set M_C of $T_{c,m}, min(d_d, M_D) > min(d_c, M_C)$.*

The Matrix Profile (MP) [23] is the most used solution to compute discords within time series data.

Definition 6 (Matrix Profile). *The Matrix Profile P of a time series T is a vector that records the z-normalized Euclidean distance between each subsequence and its nearest non-self match. Formally, $P = [min(D_1), min(D_2), \ldots, min(D_{n-m+1})$ where $D_{1 \leq i \leq n-m+1}$ represents the distance profile of the query subsequence $T_{i,m}$ in time series T.*

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63783-4_14

Definition 7 (Multidimensional Time Series). *A multidimensional time series $\mathcal{T} \in \mathcal{R}^{n \times d}$ is a n-sized set of d co-evolving time series. Formally,* $\mathcal{T} = [T^1, T^2, \dots, T^d]$.

When extending the matrix profile to multidimensional time series, we introduce a new structure called the multidimensional matrix profile. This adaptation facilitates the analysis of pattern similarity and dissimilarity across multiple dimensions within the time series data.

Definition 8 (Multidimensional Matrix Profile). *Given a multidimensional time series $\mathcal{T} = [T^1, T^2, \dots, T^d]$ with d time series, each of length n, the multidimensional matrix profile is constructed by aggregating the matrix profiles MP_i of all d time series. It stores the z-normalized Euclidean distance between each subsequence and its nearest neighbor across all dimensions.*

In this paper, we are interested in multi-dimensional time series that exhibit discords that may be present on a subset of dimensions, we call such anomalies a *K-dimensional anomaly.*

Definition 9 (K-Dimensional Anomaly). *A K-dimensional anomaly appears on at least K of the time series $\mathcal{T} = [T^1, T^2, \dots, T^d]$. When k equals the total number of time series, such a k-dimensional anomaly is referred to as a natural anomaly [20].*

Definition 10 (Natural Anomaly). *Given a multidimensional time series $\mathcal{T} = [T^1, T^2, \dots, T^d]$ consists of d times series and X a K-dimensional-anomaly in T, X is a natural anomaly if k is equal to the total number of dimensions on which the the anomaly is observed.*

Natural anomaly detection is particularly intriguing because simply declaring the presence of an anomaly is not enough. It is more valuable to identify which specific dimensions, such as sensors, are involved, especially when their number is significant. In this work, our focus is on identifying natural anomalies and pinpointing the specific time series associated with them.

The MP technique [23] has emerged as a valuable tool for uncovering various properties of time series data across a wide range of applications, including seismology, medicine, and vocalization analysis. This technique has demonstrated its utility in identifying numerous structural elements within time series datasets, such as repeated behaviors, known as motifs [13], as well as anomalies, referred to as discords [4, 16, 23], shapelets among others. Indeed, the field of time series discord detection has been gaining increasing interest within the domain of data mining [23].

4 Convolutional Neural Network-based Model for Battery SOC Estimation

In this section, we present the methodology for constructing a deep learning model to estimate the SOC values of Li-ion cells, starting from the dataset used to train and test the model until the model architecture.

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63783-4_14

4.1 Dataset

In order to develop a model able to estimate the SOC of real-world driving cycles, the dataset includes 142 cycles from the Massachusetts Institute of Technology (MIT) battery dataset [18] and a further 425 cycles from the National Institute of Applied Sciences (INSA) [7]. From the MIT dataset, we selected exactly two cycles from of the 72 identified charges policies, ensuring a broad representation of charging conditions. In addition, we incorporated the Basytec XCTS system for assessing lithium ferrophosphate (LFP) battery cells, identical to those featured in the MIT dataset. This advanced system enables us to conduct tests employing diverse protocols, including the Worldwide Harmonized Light Vehicle Test Procedure (WLTP) and the Assessment and Reliability of Transport Emission Models and Inventory Systems (ARTEMIS) cycle. The testing regimen comprises two primary phases: charge and discharge as shown in Figure 1. During the charge phase, we employ the classic CC-CV (constant current-constant voltage) method. This involves applying a constant current to the battery cells, followed by a constant voltage, a process that is crucial for accurately simulating the charging behavior of batteries in practical applications. For the discharge phase, we emulate real-world driving conditions by integrating multiple driving cycles. This phase encompasses regenerative braking, a critical feature that recuperates energy dissipated during vehicle deceleration and braking, effectively recharging the battery cells.

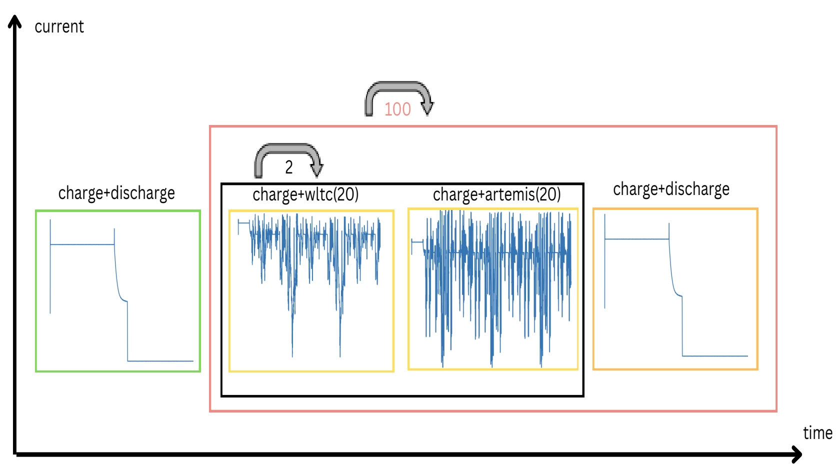

Fig. 1. Structure of the test protocols.

4.2 Data Preprocessing

In the obtained dataset from each cycle, comprising measurements of current, voltage, and temperature recorded during tests, the direct measurement of SOC

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63783-4_14

is not feasible and requires estimation. To initiate a supervisory learning process, for our model, accurately computing the SOC is indispensable. For this purpose, we employ the Coulomb counting technique [19], which involves the cumulative integration of current over time to estimate the SOC. This technique is computed using the following Equation 2:

$$\mathrm{SOC}(t) = \begin{cases} \mathrm{SOC}(t-1) + \frac{1}{\mathrm{cap}} \sum_{i=1}^{n_c} I_c(t_i)\Delta t_i, & \text{if charging} \\ 1 + \frac{1}{\mathrm{cap}} \sum_{i=1}^{n_d} I_d(t_i)\Delta t_i, & \text{if discharging} \end{cases} \tag{2}$$

where:

- $\mathrm{SOC}(t)$ represents the SOC at time t,
- $\mathrm{SOC}(t-1)$ is the previous SOC value,
- $I_c(t_i)$ and $I_d(t_i)$ denote the charging and discharging currents respectively, recorded at time t_i, with negative current values during the discharge.
- Δt_i signifies the time intervals between consecutive measurements,
- n_c and n_d are the number of measurements taken during the charging and discharging phases respectively.
- cap represents the capacity of the battery.

During the charging phase, the SOC is updated by summing the integrated current over the duration of the charge. Conversely, during the discharging, the initial SOC is subtracted from 1 (assuming a full charge), and the integrated current over the discharge duration is added.

Following the calculation of SOC values, we perform the min-max normalization technique to scale our features within a 0 to 1 range, as presented in Equation 3. More precisely, this technique was applied first to the training set, which contains 70% of the total cycles including an equal proportion from both the MIT and INSA datasets. Subsequently, we adopted the same minimum and maximum values obtained from the training set to normalize the test set.

$$\text{normalize value} = \frac{\text{data} - \min(\text{data})}{\max(\text{data}) - \min(\text{data})} \tag{3}$$

As depicted in Figure 2, in our approach, we set the input window size to 100 and the output window size to 25. The choice of these window sizes is strategic; the input window of 100 allows the model to consider a substantial sequence of data points, providing a comprehensive view of the battery's behavior leading up to the current state. This size ensures that the model has enough context to understand the temporal dynamics of SOC changes. The output window of 25, on the other hand, enables the model to predict the SOC for the next 25 time intervals based on the input sequence, offering a detailed forecast that can be invaluable for real-time battery management and planning. More interestingly, in the MIT dataset, the size of 25 corresponds to the SOC values for the next two minutes, whereas in the INSA dataset, it represents the SOC for the following minute.

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63783-4_14

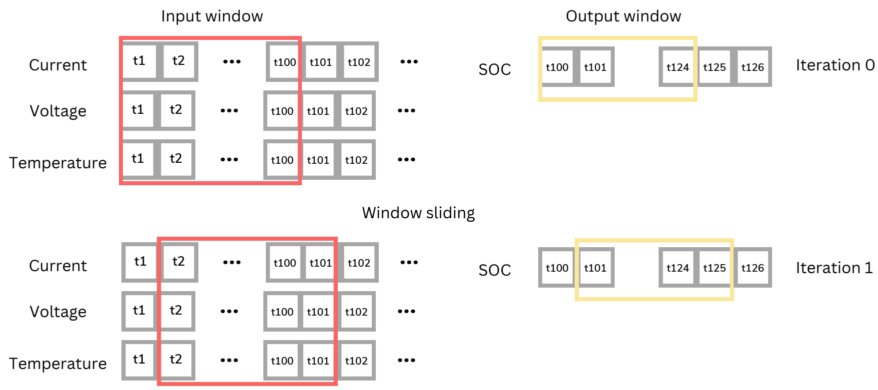

Fig. 2. The sliding window technique.

4.3 Model architecture

In this work, we propose a data-driven model that leverages the power of CNNs to estimate the SOC for lithium-ion battery cells. CNNs are renowned for their efficacy in processing and analyzing structured grid data, making them ideally suited for interpreting time series data, such as the SOC estimation from battery cycles. Here, we explain the fundamental building blocks of our CNN architecture and their roles within the model:

- *Input layer*: the first layer that receives the raw input data.
- *Convolutional layer*: this layer applies convolutional operations to the input data allowing to capture spatial hierarchies and features d
- *Activation layer*: following convolution, an activation function, commonly the Rectified Linear Unit (ReLU), is applied to introduce non-linearity and enhance the model's ability to make predictions on previously unseen data.
- *Pooling layer*: down-sample the spatial dimensions of the input data, reducing its computational complexity. Max pooling and average pooling are common techniques used in this layer.
- *Fully connected layer*: neurons in this layer are connected to all neurons in the previous layer, resembling a traditional neural network. It helps in learning global patterns and their relationships.
- *Flattening layer*: before entering the fully connected layer, the multi-dimensional data is flattened into a one-dimensional vector. This step prepares the data for the fully connected layers.
- *Output layer*: the final layer outputs the SOC estimation.

In our CNN model, we employ the ReLU activation function [1] in all layers except the final layer, where a Sigmoid activation function [2] is utilized. The architecture of this model, as depicted in Figure 3, reflects this design choice with the parameters employed for each layer.

[1] `https://www.tensorflow.org/api_docs/python/tf/keras/activations/relu`
[2] `https://www.tensorflow.org/api_docs/python/tf/keras/activations/sigmoid`

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63783-4_14

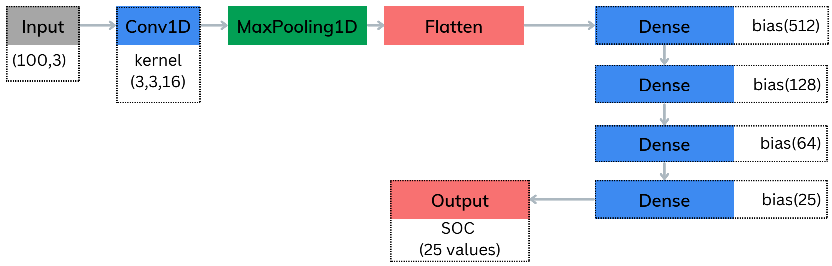

Fig. 3. Architecture of the proposed CNN model.

4.4 Explaining predictions

Efforts to enhance the interpretability and transparency of deep learning models, particularly complex CNNs for SOC estimation, have led to the integration of XAI techniques. These methods aim to elucidate the decision-making process of models, bridging the gap between advanced computational algorithms and human understanding. Among the various XAI methodologies, SHAP [14] stands out for its comprehensive approach to quantifying the influence of each feature on the model's output. Our idea to explaining predictions involves employing a post-hoc SHAP model applied to the output of our model. This allows us to obtain the contributions of the three features: current (I), voltage (V), and temperature (T) toward the SOC estimation. This analysis is visually represented in Figure 4, where each feature's impact on the CNN model's predictions is clearly illustrated.

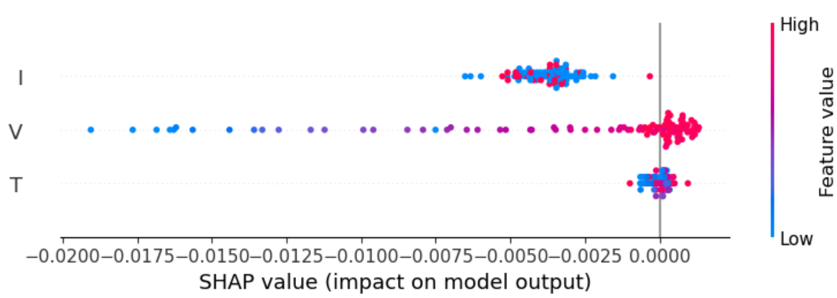

Fig. 4. Feature importance analysis with SHAP.

Upon analysis, we observed that among all the time series provided in the input, only the current exhibits the highest contribution to the predictions. Therefore, we aim to construct a time series using the SHAP values of the current.

Detecting discords Detecting discords within time series data is a crucial aspect of interpreting complex patterns, especially when assessing the impact of

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63783-4_14

different factors on battery SOC estimation. To achieve this, our approach uses a 2-dimensional matrix profile on the time series of SHAP values and current values. This allows us to detect discords, i.e., natural discords which must appear simultaneously in both times series and are identified as the largest values in the MP. We employ an algorithm called Stumpy [3] [24] built on top of the MP and able to detect these discords. Discords represent data points that are most different among all the time series. By setting a window size of $m = 100$, we tailor the algorithm to our specific dataset, allowing for a comprehensive examination of the data over time frame. This window size is chosen to balance the granularity of analysis with computational efficiency, ensuring that the algorithm can effectively detect discords without being hindered by excessive detail or data volume.

5 Experimental Evaluation

This section outlines the experimental setup, the comparative analysis with existing state-of-the-art models, and the metrics employed to assess performance.

5.1 Experimental protocol

Implementation All experiments were conducted on a machine with an Intel Core i5 12th generation CPU, a NVIDIA GeForce RTX 3070 GPU with 6GB of VRAM, 32GB of RAM, and a 512GB SSD. This machine provided the necessary computing power to train and test the models efficiently and effectively.

Competitors To evaluate the `SOCXAI` algorithm, we compared it with the state-of-the-art algorithms: a simple Feed-forward Neural Network (FNN) [6], and a model based on Long Short-Term Memory (LSTM) networks [2].

Error metrics To assess the effectiveness of our model, we utilized three standard machine learning metrics. These metrics take into account the complete set of window values under consideration, which in our case is 25, rather than focusing on individual values. The metrics we employed are defined as follows:

– Mean Squared Error (MSE): As defined in Equation 4, it quantifies the the average of the squares of the errors or deviations, in other words the difference between the estimator and what is estimated. The MSE is calculated as follows (where n represents the total number of samples):

$$MSE = \frac{1}{n}\sum_{i=1}^{n}(predicted\,value_i - observed\,value_i)^2 \tag{4}$$

[3] `https://stumpy.readthedocs.io/en/latest/#`

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63783-4_14

– Mean Absolute Error (MAE): As defined in Equation 5, the MAE measures the average absolute difference between the predicted and actual values. This metric offers insight into the magnitude of errors in the model's predictions:

$$MAE = \frac{\sum_{i=1}^{n} |predicted\ value_i - observed\ value_i|}{n} \tag{5}$$

– Root Mean Squared Error (RMSE): As defined in Equation 6, it is the square root of the MSE and is commonly used to measure the average error between the predicted and actual values in the same units as the original data:

$$RMSE = \sqrt{\frac{\sum_{i=1}^{n} (predicted\ value_i - observed\ value_i)^2}{n}} \tag{6}$$

5.2 Results

SOC estimation In Table 1, we conducted a comparative analysis between `SOCXAI` model and other baseline models, namely FNN and LSTM, using the dataset presented in Section 4.1.

Metric	`SOCXAI`	**FNN**	**LSTM**
MAE	0.0143	0.0510	0.0201
MSE	0.0016	0.0065	0.0021
RMSE	0.040	0.0809	0.0464

Table 1. SOC estimation model performance results in dataset.

We evaluated the models based on error metrics including MSE, MAE, and RMSE. The experimental results demonstrate that our model outperforms the baseline models across all error metrics, marking a threefold improvement in accuracy over the FNN model. This performance is attributed to the utilization of CNN layers, which excel in capturing data dependencies and reducing model complexity through by using fewer parameters compared to LSTM layers.

Figure 5 and Figure 6 depict the SOC estimation values generated by the models alongside the true SOC values for two randomly selected cycles from the MIT and INSA test sets, respectively. The x-axis represents the time, while the y-axis represents the SOC values which range from 0 to 1, representing 0% to 100% charge. These figures clearly demonstrate that the `SOCXAI` and LSTM models provide estimations closely aligned with the true SOC values, significantly outperforming the FNN model. We note here that the proposed model can be generalized to other types of lithium batteries; one simply needs to retrain it.

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63783-4_14

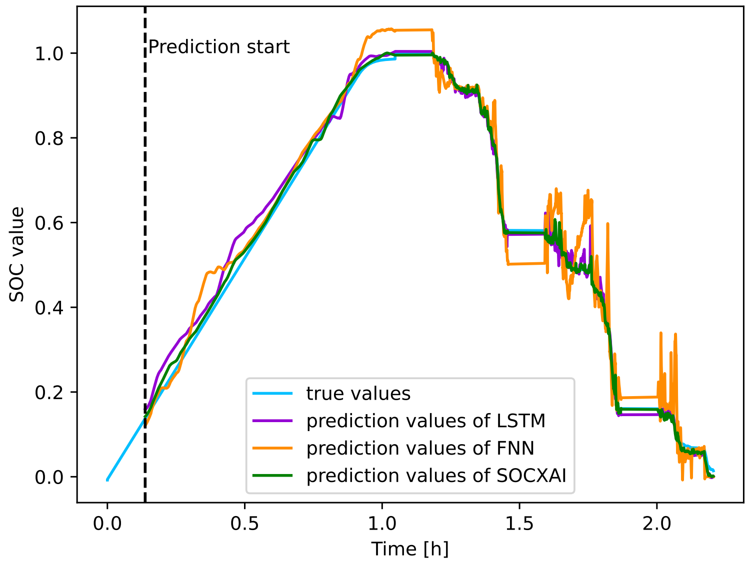

Fig. 5. Comparison of model prediction with LSTM and FNN on INSA driving cycle.

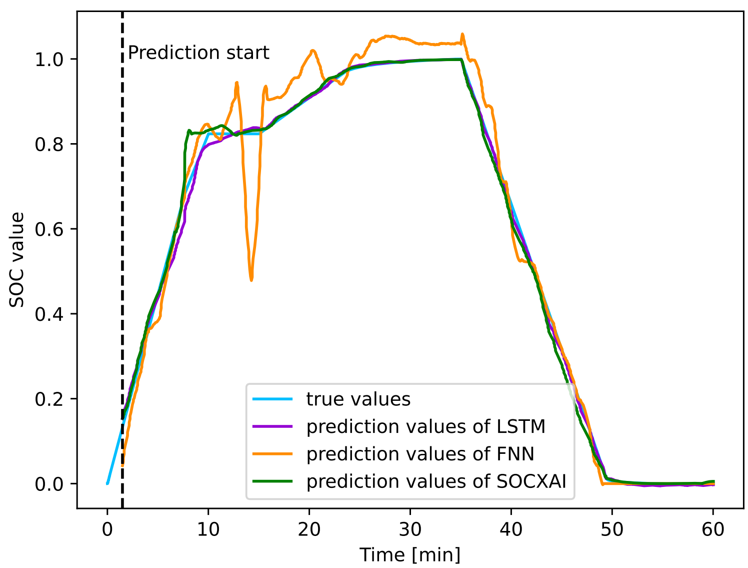

Fig. 6. Comparison of model prediction with LSTM and FNN on MIT driving cycle.

Explaining predictions. As noted in [15], the SHAP model may sometimes provide imprecise or misleading assessments of relative feature importance, par-

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63783-4_14

ticularly failing to capture inter-feature relationships. Our investigation focuses on this aspect, especially in the context of regression problems related to SOC estimation. Consequently, if the SHAP method highlights significant contributions of current values within identified anomalies, we intend to evaluate the consistency between SHAP assessments and the regions within the current values indicative of anomalies. This approach enhances our understanding of battery behavior and facilitates effective anomaly detection. Our findings confirm that the SHAP model effectively explains the presence of anomalies as shown in Figure 7 (top), indicating deviations from expected patterns. Moreover, the absence of conserved behavior in the time series underscores the efficacy of SHAP in elucidating abnormal occurrences. Notably, the contributions of SHAP values slightly increase within the subsequences where anomalies are detected, further affirming the model's ability to capture and explain these irregularities. Practically, one explanation for the abnormal behavior, as illustrated after zooming into the discord in Figure 7 (bottom), of the battery could be attributed to a voltage measurement issue. This issue affects the current values as the driving cycle progresses until it reaches a voltage threshold. However, the rate at which the voltage reaches this threshold can vary, indicating the inconsistent behavior of the current time series values.

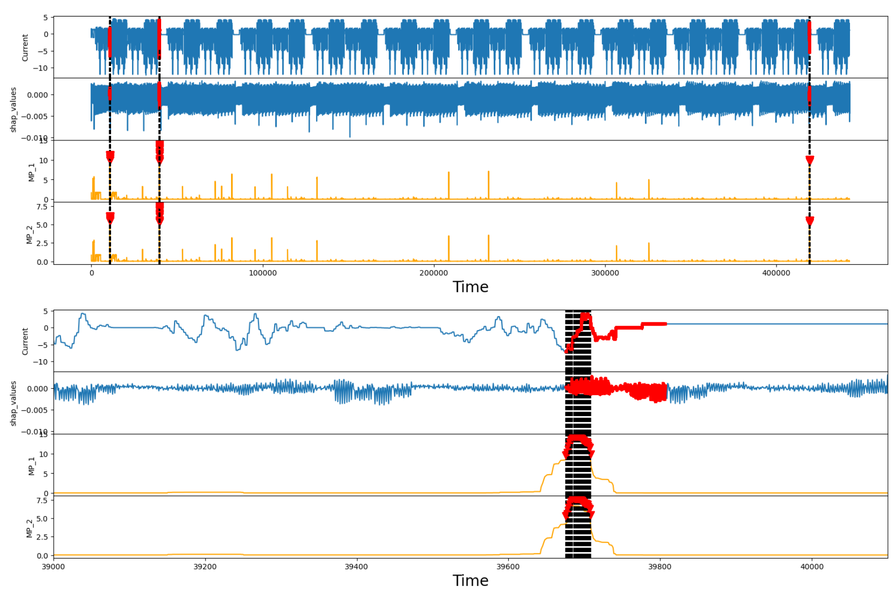

Fig. 7. Illustration of 2-dimensional anomaly detection of current and SHAP values for subsequence m = 100.

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63783-4_14

6 Conclusion and Perspectives

In this paper, we introduced a novel algorithm named `SOCXAI`, aimed at estimating the SOC of batteries. This algorithm distinguishes itself by its capability to predict not just a single future SOC value but 25 future values. Its application extends beyond simple constant discharge scenarios to encompass real-world driving cycles and various charging policies. Furthermore, it provides explanations for these predictions using the SHAP model. Additionally, we proposed an anomaly detection method using the concept of natural anomalies, highlighting abnormal battery usage patterns that deviate from expected behavior.

In future work, we aim to enhance the performance of the proposed method by exploring more advanced techniques such as utilizing a sliding window with a dynamically varying size. Furthermore, we intend to expand this model to deal with a new type of battery known as sodium-ion batteries.

Acknowledgment

This research received partial funding from the French National Research Agency (ANR) under the project 'ANR-22-CE92-0007-02'. Additionally, support was provided by the European Union through the Horizon Europe program and the innovation program under 'GAP-101103667'.

References

1. Boniol, P., Linardi, M., Roncallo, F., Palpanas, T., Meftah, M., Remy, E.: Unsupervised and scalable subsequence anomaly detection in large data series. The VLDB Journal pp. 1–23 (2021)
2. Chemali, E., Kollmeyer, P.J., Preindl, M., Ahmed, R., Emadi, A.: Long short-term memory networks for accurate state-of-charge estimation of li-ion batteries. IEEE Transactions on Industrial Electronics pp. 6730–6739 (2018)
3. Doyle, M., Fuller, T.F., Newman, J.: Modeling of galvanostatic charge and discharge of the lithium/polymer/insertion cell. Journal of the Electrochemical society p. 1526 (1993)
4. El Khansa, H., Gervet, C., Brouillet, A.: Application of matrix profile techniques to detect insightful discords in climate data. International Journal on Soft Computing, Artificial Intelligence and Applications (IJSCAI) (2022)
5. Fuller, T.F., Doyle, M., Newman, J.: Simulation and optimization of the dual lithium ion insertion cell. Journal of the electrochemical society p. 1 (1994)
6. He, W., Williard, N., Chen, C., Pecht, M.: State of charge estimation for li-ion batteries using neural network modeling and unscented kalman filter-based error cancellation. International Journal of Electrical Power & Energy Systems pp. 783–791 (2014)
7. Heitzmann, T., Samet, A., Mesbahi, T., Soufi, C., Jorge, I., Boné, R.: Sochap: A new data driven explainable prediction of battery state of charge. In: Computational Science – ICCS 2023. pp. 463–475 (2023)
8. Huria, T., Ludovici, G., Lutzemberger, G.: State of charge estimation of high power lithium iron phosphate cells. Journal of Power Sources pp. 92–102 (2014)

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63783-4_14

9. Johnson, V.: Battery performance models in advisor. Journal of power sources pp. 321–329 (2002)
10. Kashpruk, N., Piskor-Ignatowicz, C., Baranowski, J.: Time series prediction in industry 4.0: A comprehensive review and prospects for future advancements. Applied Sciences (2023)
11. Lee, J., Sun, H., Liu, Y., Li, X.: A machine learning framework for remaining useful lifetime prediction of li-ion batteries using diverse neural networks. Energy and AI p. 100319 (2024)
12. Li, G., Jung, J.J.: Deep learning for anomaly detection in multivariate time series: Approaches, applications, and challenges. Information Fusion pp. 93–102 (2023)
13. Linardi, M., Zhu, Y., Palpanas, T., Keogh, E.: Matrix profile x: Valmod-scalable discovery of variable-length motifs in data series. In: Proceedings of the 2018 International Conference on Management of Data. pp. 1053–1066 (2018)
14. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions. Advances in neural information processing systems (2017)
15. Marques-Silva, J., Huang, X.: Explainability is not a game. arXiv preprint arXiv:2307.07514 (2023)
16. Nakamura, T., Imamura, M., Mercer, R., Keogh, E.: Merlin: Parameter-free discovery of arbitrary length anomalies in massive time series archives. In: 2020 IEEE international conference on data mining (ICDM). pp. 1190–1195 (2020)
17. Plett, G.L.: Extended kalman filtering for battery management systems of lipb-based hev battery packs: Part 3. state and parameter estimation. Journal of Power sources pp. 277–292 (2004)
18. Severson, K.A., Attia, P.M., Jin, N., Perkins, N., Jiang, B., Yang, Z., Chen, M.H., Aykol, M., Herring, P.K., Fraggedakis, D., Bazant, M.Z., Harris, S.J., Chueh, W.C., Braatz, R.D.: Data-driven prediction of battery cycle life before capacity degradation. Nature Energy p. 383–391 (2019)
19. Stefanopoulou, A., Kim, Y.: System-level management of rechargeable lithium-ion batteries. Rechargeable Lithium Batteries pp. 281–302 (2015)
20. Tafazoli, S., Keogh, E.: Matrix profile xxviii: Discovering multi-dimensional time series anomalies with k of n anomaly detection. In: Proceedings of the 2023 SIAM International Conference on Data Mining (SDM). pp. 685–693 (2023)
21. Tian, J., Chen, C., Shen, W., Sun, F., Xiong, R.: Deep learning framework for lithium-ion battery state of charge estimation: Recent advances and future perspectives. Energy Storage Materials p. 102883 (2023)
22. Yan, Q.: Soc prediction of power battery based on svm. In: 2020 Chinese Control And Decision Conference (CCDC). pp. 2425–2429 (2020)
23. Yeh, C.C.M., Zhu, Y., Ulanova, L., Begum, N., Ding, Y., Dau, H.A., Silva, D.F., Mueen, A., Keogh, E.: Matrix profile i: all pairs similarity joins for time series: a unifying view that includes motifs, discords and shapelets. In: 2016 IEEE 16th international conference on data mining (ICDM). pp. 1317–1322 (2016)
24. Zhu, Y., Zimmerman, Z., Senobari, N.S., Yeh, C.C.M., Funning, G., Mueen, A., Brisk, P., Keogh, E.: Matrix profile ii: Exploiting a novel algorithm and gpus to break the one hundred million barrier for time series motifs and joins. In: 2016 IEEE 16th international conference on data mining (ICDM). pp. 739–748 (2016)

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63783-4_14