Your Battery Is a Blast!
Safeguarding Against Counterfeit Batteries with Authentication

Francesco Marchiori
University of Padova
Padua, Italy
francesco.marchiori@math.unipd.it

Mauro Conti
University of Padova
Padua, Italy
mauro.conti@unipd.it

ABSTRACT

Lithium-ion (Li-ion) batteries are the primary power source in various applications due to their high energy and power density. Their market was estimated to be up to 48 billion U.S. dollars in 2022. However, the widespread adoption of Li-ion batteries has resulted in counterfeit cell production, which can pose safety hazards to users. Counterfeit cells can cause explosions or fires, and their prevalence in the market makes it difficult for users to detect fake cells. Indeed, current battery authentication methods can be susceptible to advanced counterfeiting techniques and are often not adaptable to various cells and systems.

In this paper, we improve the state of the art on battery authentication by proposing two novel methodologies, **DCAuth** and **EISthentication**, which leverage the internal characteristics of each cell through Machine Learning models. Our methods automatically authenticate lithium-ion battery models and architectures using data from their regular usage without the need for any external device. They are also resilient to the most common and critical counterfeit practices and can scale to several batteries and devices. To evaluate the effectiveness of our proposed methodologies, we analyze time-series data from a total of 20 datasets that we have processed to extract meaningful features for our analysis. Our methods achieve high accuracy in battery authentication for both architectures (up to 0.99) and models (up to 0.96). Moreover, our methods offer comparable identification performances. By using our proposed methodologies, manufacturers can ensure that devices only use legitimate batteries, guaranteeing the operational state of any system and safety measures for the users.

CCS CONCEPTS

• **Hardware** → **Batteries**; • **Security and privacy** → **Authentication**; • **Computing methodologies** → *Machine learning*.

KEYWORDS

Lithium-ion Batteries; Authentication; Identification; Machine Learning

ACM Reference Format:
Francesco Marchiori and Mauro Conti. 2023. Your Battery Is a Blast! Safeguarding Against Counterfeit Batteries with Authentication. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS '23), November 26–30, 2023, Copenhagen, Denmark.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3576915.3623179

1 INTRODUCTION

Lithium-ion (Li-ion) batteries [45] are currently dominating the market due to their higher energy efficiency and low memory effects [56]. An active line of research backs up this growth, with many works on the improvement of their materials and architectures [60], studies on their aging behaviour [8] and devices for their management [69]. Their rapid development over the years has made them particularly effective solutions in many applications, ranging from small electronic devices to electric vehicles (EVs) [43]. Indeed, the global market for Li-ion batteries in 2022 was estimated to be up to 48 billion U.S. dollars, with an expected compound annual growth rate of 18.1% from 2022 to 2030 [23].

This widespread demand for Li-ion cells, however, is causing a rise in their counterfeiting, i.e., the production of fake or unauthorized replicas of legitimate batteries. Indeed, counterfeit batteries are flooding the market and the supply chain from the Original Equipment Manufacturer (OEM) to the vendors [33]. A significant number of batteries are seized worldwide each year for an estimated value of several million dollars [16, 46]. These cells are branded similarly to their legitimate counterpart and advertised to work in the same way. However, they are often made of lower quality materials and have no safety certification. For these reasons, the risk of fire hazards is significantly higher with respect to the original battery manufacturer cells. Indeed, the operative temperature for cells can reach critical values, and illegally manufactured batteries or recycled ones might not correctly enforce safety measures. However, detecting these cells is particularly difficult since their physical condition is often as good as a legitimate one. By rewrapping the cells, it is possible to fake their rating or capacity, remark them as a different model or even disguise them as a different cell [59]. In Figure 1, we show three examples of counterfeit rewrapped batteries. In those images, smaller capacity cells are connected through a step-up circuit to the positive and negative terminals of the outer cell. While cells such as the one in Figure 1a might behave similarly to their legitimate counterpart, their lifespan are dramatically decreased. However, weighting those cells can detect these counterfeiting attempts since they will be much lighter. Instead, in Figure 1b and 1c, the size gap is filled with powders or other materials to compensate for the weight difference, making the detection of those fakes much harder. Users generally report many issues or power interruptions

CCS '23, November 26–30, 2023, Copenhagen, Denmark

ACM ISBN 979-8-4007-0050-7/23/11…$15.00
https://doi.org/10.1145/3576915.3623179

during their usage.[1] Others that use those counterfeited cells in on-road vehicles instead are particularly in danger and can also pose a threat to nearby vehicles.[2]

Contribution. In this paper, we present **DCAuth** and **EISthentication**, two novel solutions for the automatic authentication and identification of Li-ion battery cells. While other methods for the authentication of Li-ion batteries are present in the literature, to the best of our knowledge, ours are the first to consider only physical and chemical features and thus do not rely on external devices or Challenge-Response (CR) authentication protocols. This is significant since other methods have been proven weak to several attacks or not scalable to many different battery models. Table 1 shows an overview of the strengths of DCAuth and EISthentication compared to other popular authentication methods. Our methodologies take advantage of data retrieved through, respectively, Differential Capacity Analysis (DCA) and Electrochemical Impedance Spectroscopy (EIS). The former is a technique used in electrochemical measurements to study the behavior of the electrode-electrolyte interface. The latter is a non-destructive method for characterizing Li-ion batteries [44]. By considering different cell models and performing DCA or EIS on them at different States Of Charge (SOC) and States Of Health (SOH), we extract features to train several Machine Learning (ML) models. Of the different models considered, the best results are reached by the Random Forest (RF) classifier, which obtains scores up to 0.96 in the authentication of different battery samples. The main contributions of our work can be summarized as follows:

- We improve the state of the art on battery authentication by leveraging the internal characteristics of each cell and Machine Learning models. We propose two methodologies (**DCAuth** and **EISthentication**) for the authentication and identification of batteries in any setting or environment.
- We define and publish a common procedure to process battery data coming from different datasets. Our procedure considers several types of equipment and their characteristics to facilitate future research in this field.
- We evaluate our methodologies on 20 datasets we collected and processed. Our evaluation is differentiated into various steps and considers both the authentication and identification tasks.
- We make available the code and implementation for all of our methodologies, including both the Machine Learning models and their experimental evaluation. Our repositories can be accessed at https://github.com/Mhackiori/DCAuth and https://github.com/Mhackiori/EISthentication.

Organization. The paper is organized as follows. Section 2 reviews related works on different battery authentication methods and the techniques used to perform DCA and EIS. An overview of the system model and examples of practical deployment are given in Section 3. In Section 4, we give an overview of the data collection process and the characteristics of the different datasets used for each methodology. In Section 5, we explain in detail our methodology, presenting an experimental evaluation in Section 6. Section 7 presents a discussion of the obtained results and possible limitations of our methodologies. Finally, Section 8 concludes this work.

Table 1: Comparison of different methodologies for battery authentication. In each cell, we put a checkmark if the considered methodology is resilient against a particular attack or weakness.

Method	Cloning	Replay Attacks	Unscalability	Rewrapping
Markings		✓	✓	
External Features		✓	✓	
Form Factor		✓	✓	
Resistor		✓	✓	
Chip	✓			
CR (in clear)	✓			
CR (encrypted)	✓	✓		
DCAuth	✓	✓	✓	✓
EISthentication	✓	✓	✓	✓

2 RELATED WORKS

We now overview several techniques that are currently used for battery authentication and their flaws that led to the need for a device-independent authentication method (Section 2.1). Furthermore, we summarize the core concepts behind the techniques that we use in our methodologies: Differential Capacity Analysis (Section 2.2) and Electrochemical Impedance Spectroscopy (Section 2.3).

2.1 Battery Authentication

Battery authentication refers to the process of verifying the authenticity of a battery, typically to ensure that it is a genuine product and not a counterfeit or a lower-quality substitute. Several methods can be used to authenticate batteries, including visual inspection, chemical analysis, and various non-destructive testing techniques. In this review, we will only analyze the methods that authenticate the battery before a system is allowed to function from it. Indeed, the system should refuse power from a battery that failed the authentication process [68].

- **Visual Inspection** – One method of battery authentication is visual inspection, which involves looking for physical characteristics that are unique to the manufacturer or product. For example, many batteries have specific markings or labels that can be used to identify them. In addition, the appearance of the battery itself, such as its color, shape, and size, can also be used to verify its authenticity. This method, however, is vulnerable to rewrapping, and in general, it is really easy to replicate the markings on the wrapping or other external characteristics.

[1] http://e-motion.lt/2016/04/13/ultrafire-5800mah/
[2] https://endless-sphere.com/forums/viewtopic.php?t=80451

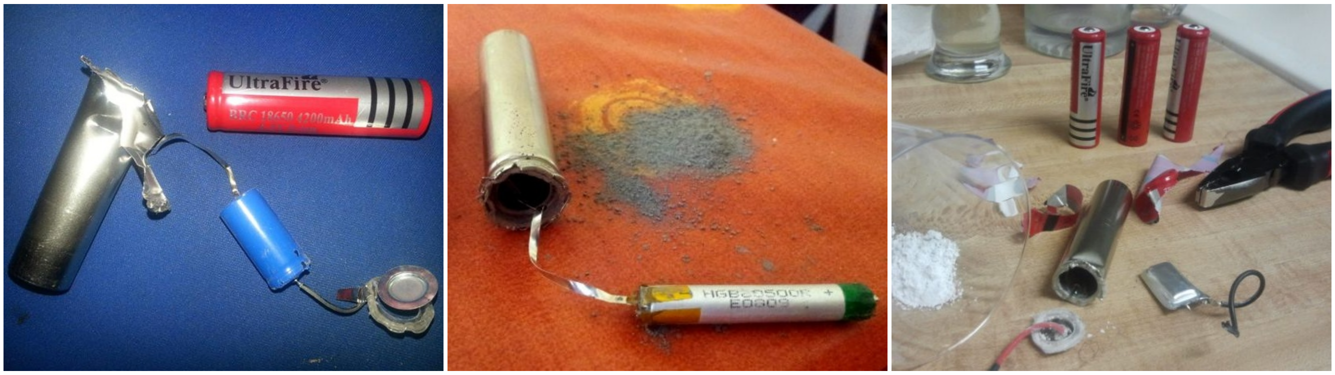

(a): A 18350 cell posing as a 18650 cell. **(b): Smaller cell rewrapped with sand.** **(c): Counterfeit with pouch battery.**

Figure 1: Examples of counterfeit 18650 batteries.

- **Form Factor** – A similar authentication method involves leveraging the form factor of the battery casing and connectors. However, third parties can easily reproduce the physical properties of the cells to create their counterfeit samples.
- **Resistor or Chip** – Another basic approach is to place a resistor in the battery pack to identify its type. Alternatively, instead of a resistor, it is possible to place a memory chip containing several data on the battery, including type, ID, and manufacturing date. However, these resistors or chips can be extracted and replaced on counterfeit samples, which will pose as legitimate to a prover.
- **Challenge-Response Protocols** – It is possible to use a Challenge-Response (CR) protocol between the chip and the prover to authenticate the battery. In EVs, this is usually handled by the fuel gauges, which manage data coming from the battery to monitor its State Of Function (SOF) [54]. The most basic implementation of this protocol is using an unchanging stream of bits that, however, can be sniffed by an attacker and thus is vulnerable to replay attacks. More advanced gauges instead include some cryptographic hash function in the protocol, making it impossible for an attacker to steal the authentication codes [1]. While this method is particularly effective in preventing cloning, it requires the creation of the keys during the battery's manufacturing process. This makes the legitimate battery sample compatible only with a specific family of fuel gauges.

We thus propose two techniques for battery authentication that do not depend on external devices and instead leverage physical and chemical characteristics to determine the legitimacy of a battery. Moreover, updating the labels corresponding to the legitimate samples makes it possible to contemplate the legitimacy of new battery packs or architectures that did not exist at the time of the implementation.

2.2 Differential Capacity Analysis

Differential Capacity Analysis is a common method for examining the interfaces between electrodes and electrolytes in electrochemical systems. Differential Capacity can track the increase in capacity when charged or a decrease when discharged of an electrochemical system. By creating a plot of differential capacity versus voltage, a "fingerprint" of the system can be obtained and monitored over time. This provides insight into the system's thermodynamics and kinetics as the characteristic features of the curve change.

While initially explored as a quality control measure, differential capacity analysis can aid in foreseeing malfunctions by observing variations in the electrochemical characteristics of a system [76]. This capacity for prediction would aid in preventing safety hazards and ensuring that a battery or cell is retired only when it can no longer function effectively.

The use of DCA to estimate State Of Charge and State Of Health has also been widely researched [35, 65]. Battery management and monitoring rely on SOH estimation to gain insights into a battery's condition over time. Since DCA works by analyzing the difference in capacity between the initial discharge and subsequent cyclic discharge of a battery, by measuring this difference, it is possible to estimate the amount of irreversible capacity loss. Using DCA, battery manufacturers and designers can enhance their understanding of the battery's performance throughout its lifespan, improving battery longevity and reducing overall system costs.

With DCAuth, we leverage the relationship between DCA and the degradation of battery samples to authenticate and identify different models and architectures. Indeed, some reaction processes consume electrolyte species and/or active lithium in Li-ion cells. These are often referred to as parasitic reactions. These electrochemical reactions occur at the interfaces between the high-voltage positive electrode or low-voltage negative electrode and the electrolyte. Since parasitic reactions have been shown to alter peaks in the differential capacity plot [34], our model can extract meaningful features and use them to correctly identify and authenticate the model and architecture of battery samples. Furthermore, since both voltage and capacity measurements are fairly accessible in most systems, DCA estimations can be performed on many different devices [12]. This makes DCAuth particularly affordable and inexpensive in most authentication scenarios.

2.3 Electrochemical Impedance Spectroscopy

Electrochemical Impedance Spectroscopy (EIS) is a powerful analytical technique that allows for the characterization of electrochemical systems by measuring their electrical impedance [13]. This technique is widely used in various fields, including materials science, corrosion science, electrochemistry, and biomedical engineering, as it provides valuable insights into the behavior and properties of a system [36]. In an EIS measurement, a small AC voltage is applied to an electrochemical cell, and the resulting current is measured. The system's impedance is then calculated as the ratio of the applied voltage to the measured current. It is possible to obtain a detailed picture of the system's electrical properties by measuring the impedance over a range of frequencies. In addition to its analytical capabilities, EIS has several other advantages. It is a non-invasive technique, meaning it does not require the destruction of the sample or the addition of chemical probes. It is also relatively simple to set up and perform and can be carried out at a wide range of temperatures [11].

EIS can be used to study the electrical and electrochemical properties of batteries and supercapacitors, such as the charge/discharge behavior, the rate capability, and the cycling stability. It can provide valuable information about the performance and durability of these energy storage devices. Especially with lithium-ion batteries, several studies have been performed on the usage of EIS for the estimation of the State Of Charge (SOC) [52, 70], State Of Health (SOH) [31, 39] and other relevant data for the diagnosis of the cells [79]. In particular, SOC, SOH, and temperature are important parameters that can influence materials and devices' electrical and electrochemical properties. They can therefore affect the Nyquist plot of the impedance. The Nyquist plot is retrieved by plotting the imaginary part of the impedance extracted through EIS versus its real part [40]. Figure 2 gives a graphical representation of this influence.

Given the unique trend of the impedance with respect to the electrochemical components of the battery, with EISthentication we leverage the real and imaginary parts of the plot to extract features that will be used to train Machine Learning classifiers. These features allow us to profile each battery sample and thus authenticate different models of Li-ion batteries.

3 SYSTEM MODEL

In this section, we outline our conceptual framework and provide possible implementations of our authentication system. Figure 3 shows an overview of the system model. While the majority of the detailed steps are shared between DCAuth and EISthentication, their difference lies primarily in the data collection process, which will be treated separately in Section 3.1 and Section 3.2.

3.1 DCAuth System Model

As described in Section 2.2, Differential Capacity Analysis is a technique that requires only voltage and capacity measurements in order to be performed. These data types are accessible in most battery-powered systems since they are often already used for SOC or SOH estimation. Thus, DCAuth can be performed by the device without external equipment. To make the process viable for compact, portable devices, authentication is carried out using Machine Learning models of minimal complexity, adding minimal overhead to the system. Essentially, the authentication system can constitute a software module of the system. This approach allows for pre-executed model training, while Over-The-Air (OTA) updates enable parameter adjustments to enhance performance or incorporate new battery types. During the regular usage of the device, voltage and capacity data are collected and processed to estimate the DCA plot over one battery cycle. This cycle can be either the charging or discharging cycle, or both can be considered to increase accuracy. After sufficient data has been gathered, a filtering process is employed to eliminate any artifacts, and features are extracted from the processed plots. Finally, the data samples are tested on the deployed model, and an authentication response is generated. While authentication can be performed once each time a battery swap is detected, it is also possible to provide continuous authentication while concurrently gathering data to enhance performance.

3.2 EISthentication System Model

The main difference of using EIS instead of DCA relies on the data collection process. While with DCAuth it is possible to gather data for authentication while the device is in use, with EISthentication external equipment is needed in order to compute the real and imaginary part of the battery impedance. However, since EIS can be performed at any stage of the battery's life and charge level, data collection takes significantly less time than DCA, depending on the battery implementation. This allows authentication to be performed before even powering the device. In this way, the system model can be fully implemented in an external device comprising both the measurement apparatus and the computational power necessary for data processing and testing on the models. Furthermore, several research works have studied the possibility of retrieving EIS without specialized equipment, which in the future might close the gap between the convenience of the internal system model of DCAuth and the quickness of EISthentication [30, 40, 67].

4 DATASETS

This section details the used datasets and how they have been integrated with our methodologies. We first detail our data collection process (Section 4.1) while focusing separately on the datasets for DCA (Section 4.1.1) and the ones for EIS (Section 4.1.2). After that, we describe the post-processing performed on these datasets to remove artifacts and clean the data (Section 4.2).

4.1 Collection

One of the most crucial aspects of our work is the data collection process. Indeed, by using Machine Learning models to perform our tasks, we need huge amounts of data that fulfill specific requirements.

- **Different battery models** – Since our task can be summarized as a classification task, we need as many classes as possible to represent a real-world application for our tool accurately. However, those classes' types depend on what we want to classify. We can indeed classify battery models, i.e., a specific production line by the manufacturer, or battery architectures, i.e., the physical and chemical design and structure of the battery (e.g., Lithium Iron Phosphate, Nickel

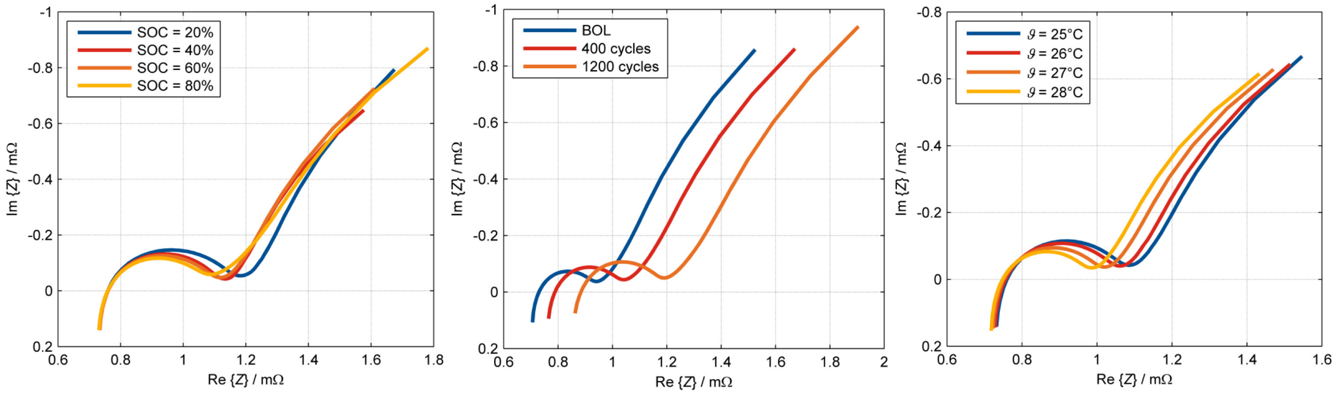

(a): Dependence on SOC.

(b): Dependence on SOH.

(c): Dependence on temperature.

Figure 2: Nyquist plots of EIS for different parameters. Images from [40].

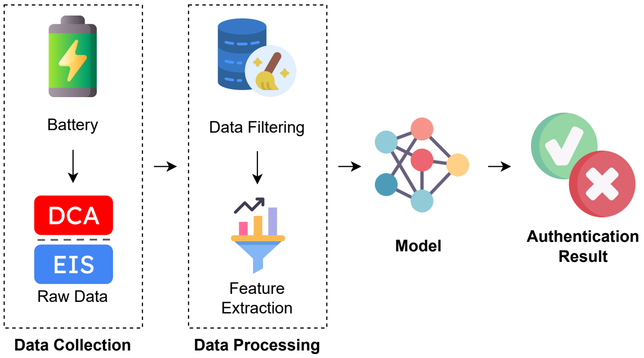

Figure 3: System Model of DCAuth and EISthentication.

Manganese Cobalt). Thus, we need datasets that contain data on many different battery cells, which include different form factors, lithium-ion battery designs, and nominal voltage and capacity values.

- **Varying conditions** – Both our methodologies are designed to be adaptable to many different environments (e.g., authentication of a battery pack in an EV, authentication of a battery in a portable device). Therefore, our datasets should include data on battery cells that have been cycled under different conditions. In particular, we highlighted how SOC, SOH, and temperature can influence DCA and EIS plots. Therefore, our study should consider the different stages of each battery lifecycle.
- **Cycling type** – Data itself is extracted from each battery sample by stressing it through many charge cycles, i.e., by charging and discharging it. However, the way in which each battery is cycled can vary depending on different conditions. For example, a typical method of charging small batteries is Constant Current Constant Voltage (CCCV), where the cell is initially charged at a constant current, but when it is nearly full, it switches to constant voltage. These constant values, however, can change depending on the battery model and its future implementation. Furthermore, many different discharging profiles are publicly available to mimic their implementation on different environments (e.g., driving profiles [63]). Therefore, our datasets should include heterogeneous data gathered from different cycling types and profiles.

These requirements impose some heavy constraints on our data creation process. For this reason, we contacted many different private companies to collaborate with them on the generation of the dataset. To our dismay, no institution agreed to share their data with us or to generate new datasets. This is due to several reasons. First, cycling many different batteries is an expensive procedure requiring constant maintenance. On top of the energy costs in maintaining the infrastructure, cell samples can be quite costly when we consider scenarios like automotive, where an EV battery can cost from $5000 to $20000. Another aspect to consider is the degradation of those cells. Indeed, by constantly cycling them, they would degrade much faster than with regular usage. For these reasons, we instead collect most of the available datasets in the literature and process them to extract the data that we need for each of the methodologies.

The usage of publicly available datasets restricts us to using only original battery sample data in their intended application. However, this aspect does not prevent us from detecting counterfeit batteries for the following reasons.

- As shown in Figure 1, counterfeit batteries are obtained by masquerading a lower-quality battery as a different cell. This implies that recognizing a fake sample is the same as determining whether the utilized battery is appropriate for that specific purpose. Therefore, the following sections will treat different battery models as counterfeit based on their legitimate application.
- While pursuing variety in the datasets regarding condition parameters and type of cycling performed, we obtain data from different cells and different environments. This allows us to gather information on different implementations of the same cell and study similar cycling procedures' effects

on different batteries. Thus, achieving good identification and authentication results will indicate the ability of our methodologies to distinguish between different battery usages and thus detect the legitimacy of a battery in a specific application.

4.1.1 DCA Datasets. As stated in Section 2.2, to perform Differential Capacity Analysis, we need measurements on the voltage and capacity of each battery sample. These measurements are quite common in many different datasets since they are fairly accessible and easy to retrieve. To collect the most amount of data available in the literature, we refer to a survey performed by Dos Reis et al. [18]. In this work, the authors analyze over 30 datasets and their characteristics. This allows us to choose among them the ones that are most in line with our requirements. Indeed, while most of the listed datasets can be eligible for our work, some are now unavailable or do not contain data needed for processing. In Table 2, we overview each dataset we use for DCAuth. Furthermore, we highlight many different aspects of each dataset, such as cycling type, cycling equipment, and other data that can be retrieved from their respective papers. Indeed, for our purposes, having datasets in which batteries have been cycled with different techniques can help us identify and authenticate each cell under different stressing conditions. Moreover, by being differential, DCA is less sensitive to the absolute values of the charge and voltage than other techniques. This is because the differential value is calculated as a ratio of changes in capacity and voltage rather than absolute values. Thus, we ensure that the type of cycling used in the datasets does not introduce bias in our models' performance.

As we notice, data in each dataset has been extracted with different equipment and in many different formats. Thus, as a preprocessing procedure, we manage the different file formats by converting each dataset into many csv files containing only the data that we need (i.e., voltage and capacity). To allow researchers to work in this field more efficiently, we publish both the code for our preprocessing procedure and the processed datasets in our repository.

4.1.2 EIS Datasets. While datasets containing voltage and capacity can be easily found in the literature, studies on their internal impedance or Electrochemical Impedance Spectroscopy are far more rare. Indeed, EIS is a complex technique that requires specialized equipment and trained technicians to perform the experiments. Furthermore, EIS equipment is expensive and requires regular maintenance, which makes it difficult for researchers with limited funding to perform EIS experiments. However, many research works are moving toward estimating Electrochemical Impedance Spectroscopy in a device or a vehicle. While some works still need some additional electronics [67], others leverage data on internal components such as the excitation of the motor controller [30] or real driving data [40]. However, due to the unavailability of such datasets, our focus remains exclusively on EIS data obtained directly from measurements.

The first dataset that we use is the SiCWell dataset [20]. This dataset contains data on automotive-grade lithium-ion pouch bag cells cycled with two popular driving profiles (sWLTP and UDDS). The data collection's main focus is investigating the influence of ripple currents in the EV battery. Still, several checkups on each pouch bag are performed periodically. In particular, the authors performed electrochemical impedance spectroscopy from 0.001 to 50000 Hz using an EIS-meter for each sample. Each procedure has been repeated at every checkup for four different SOCs (20%, 40%, 60%, and 80%). The processing resulted in a dataset containing 37 different battery samples. The dataset is freely available on IEEE DataPort[3].

To train and evaluate the capability of our methodology in distinguishing different cell architectures, we will use another dataset provided by Sandia National Laboratories (SNL) [7]. In this dataset, authors considered several commercial 18650 Li-ion battery models with four different chemistries: $\mathrm{LiCoO_2}$ (LCO), $\mathrm{LiFePO_4}$ (LFP), $\mathrm{LiNi_xCo_yAl_{1-x-y}O_2}$ (NCA), and $\mathrm{LiNi_{0.80}Mn_{0.15}Co_{0.05}}$ (NMC). The cells have been tested by cycling them at different temperatures, and EIS has been performed at the same SOC to guarantee consistent results. The frequency range for EIS is from 0.1 to 100000 Hz with a 0.010 V perturbation and has been performed at five different temperatures (5°C, 15°C, 25°C, 35°C, and 45°C). Furthermore, EIS has also been performed on brand-new cells at pristine conditions. The dataset is freely available in the official Sandia R&D Data Repository[4].

The final dataset that we consider is provided by Zhang et al. in [75]. In their work, the authors cycled 12 LCO graphite Li-ion cells at three different temperatures (25°C, 35°C, and 45°C) and performed EIS at nine different stages of their charging/discharging cycles from 0.02 to 20000 Hz. Each cell is cycled until it reaches its End Of Life (EOL), which is defined as when its SOH drops below 80% its initial value. The dataset is freely available on Zenodo[5].

4.2 Processing

As anticipated in Section 4.1.1, collecting different datasets from the literature often comes at the cost of having different data formats and extraction mechanisms. While the datasets that we use for EISthentication contain the raw spectroscopy values (and thus Nyquist plots could be directly generated), the ones for DCAuth instead contain only the voltage (V) and capacity (Q) values. Thus, differential capacity (i.e., dQ/dV) must be calculated from those values and processed to remove artifacts. Indeed, the raw differential capacity can be computed directly with the following equation.

$$(dQ/dV)_i = \frac{(Q_i - Q_{i-1})}{(V_i - V_{i-1})}. \tag{1}$$

However, values obtained in this way often present some noise and contain many bogus local maxima/minima, as shown in Figure 4a, which can heavily influence the learning process of our models. For this reason, our first step is to clean the data, as shown in Figure 4b. To do so, we remove from the raw data the points in which the voltage is too close to the previous data point. Indeed, as we can notice in Equation 1, when the denominator is too close to zero, we will obtain values that are far too high and will constitute fake peaks in the plot. Finally, to remove irregularities on the plot, we need to smoothen it as shown in Figure 4c. To do so, we use the Savitzky-Golay filter, which is a type of linear filter that uses a moving window of data to estimate the value of a given point in the signal [57]. These processing steps are based on the work of

[3] https://ieee-dataport.org/open-access/sicwell-dataset

[4] https://www.sandia.gov/ess/tools-resources/rd-data-repository

[5] https://zenodo.org/record/3633835

Table 2: Datasets used for DCAuth.

Dataset	Battery Model	Battery Architecture	Equipment	Data
Berkley [25]	Sanyo 18650	LCO/Graphite	N.A.	CCCV, MCC, CP-CV, and Boostcharge cycles at various C-rates.
CALCE_1 [27, 72, 77]	INR 18650-20R	NMC/Graphite	Arbin BT2000	Low Current and Incremental Current OCV tests, Dynamic Test Profiles.
CALCE_2 [27, 72]	ANR26650M1A	LFP	Arbin BT2000	Low Current OCV tests, Dynamic Test Profiles.
CALCE_3 [28, 71, 73]	CS2	LCO	Arbin BT2000, CADEX Tester	CCCV with different discharging protocols.
CALCE_4 [28, 73]	CX2	LCO	Arbin BT2000, CADEX Tester	CCCV with different discharging protocols.
CALCE_5 [58]	PL Samples	LCO/Graphite	Arbin BT2000	CCCV cycles on different SOC ranges.
EVERLASTING_1 [22]	INR18650 MJ1	NMC	Maccor	Aged at different C-rated and temperature within a 10-90% SOC window.
EVERLASTING_2 [66]	INR18650 MJ1	NMC	Maccor	Aged at different C-rated and temperature within a 10-90% SOC window.
HNEI [17]	ICR18650 C2	LCO/NMC	Arbin	Cycled at 1.5C to 100% DOD for more than 1000 cycles at room temperature.
OX [10, 51]	SLPB533459H4	LCO	Maccor 4200	1-C charge, 1-C discharge, pseudo-OCV charge, pseudo-OCV discharge.
OX_1 [53]	SLPB533459H4	LCO	Maccor 4200	CCCV charge and CCCV discharge.
OX_2 [50]	NCR18650BD	NCA	Maccor 4200	Different combined profile groups with reference performance tests.
SNL [49]	• APR18650M1A • NCR18650B • LG 18650HG2	• LFP • NCA • NMC	Arbin LBT21084	Charged at 0.5C, discharged at 3C. Cycled at three different SOC ranges (0-100, 20-80, 40-60) at CC or CCCV.
TRI_1 [61]	APR18650M1A	LFP/Graphite	Arbin LBT 48ch	Batteries charged with a one-step or two-step fast-charging policy depending on SOC.
TRI_2 [5]	APR18650M1A	LFP/Graphite	Arbin LBT 48ch	Cells are cycles with one of 224 six-step 10-minutes fast charging protocols.
UCL [29]	INR18650 MJ1	NMC/Graphite	Maccor 4200	CC charging at 1.5 A until 4.2 V. Discharging at 4.0 A to 2.5 V.
UL-PUR [32]	NCR18650B	NCA	Arbin BT2543	Discharged to 2.7 V (CC), charged to 4.2 V (CCCV).

Thompson et al. [64] and are consolidated as best practices when dealing with discrete sampling of data associated with batteries [19, 47]. After this procedure (which we apply to all the datasets), we ensure that our processed data do not contain any artifacts which might create biases during classification.

5 METHODOLOGIES

In this section, we outline in more detail the techniques that we use for authentication. First, we discuss our feature extraction procedure, which is a crucial part of our work given the heterogeneity of the data we use (Section 5.1). Then, we overview the Machine Learning models we use as classifiers and their optimal hyperparameter search (Section 5.2).

5.1 Feature Extraction

In Section 4.2, we described our preprocessing procedure and how we extract and clean the data from our various datasets. However, since we will use Machine Learning models to authenticate the various battery cells, graph data is not so suited in its raw form to be fed to the classifiers. Moreover, in Section 2.2 and 2.3, we have identified which could potentially be the most important characteristics in our data, i.e., peaks values and location for DCA and dependence on battery usage for EIS. For these reasons, we define a common feature extraction procedure that will consider the many different aspects of the plotted data. We can extract a

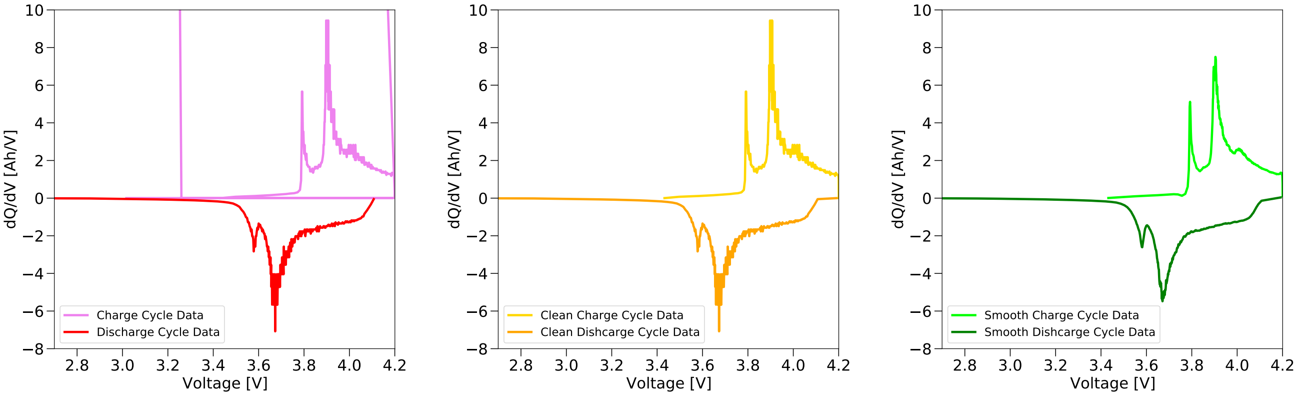

(a): Raw DCA plot. **(b): Clean DCA plot.** **(c): Smooth DCA plot.**

Figure 4: Data processing overview on Differential Capacity Analysis.

constant number of features from any plot by using *tsfresh*, an open-source Python package designed for time series feature extraction and selection [15]. While our data is not time series data per se, tsfresh allows us to extract statistical features, spectral features, and features related to autocorrelation, trend, resampling, and others.

5.2 Models

Once the datasets have been processed and features have been extracted, we define eight different Machine Learning (ML) models to deploy for the authentication task. Machine Learning is the field of Artificial Intelligence (AI) that focuses on developing algorithms that can learn from and make predictions based on data. For this reason, it is often used in authentication [62]. Indeed, by identifying patterns in the DCA and EIS data, it is possible to identify not only the cell chemistry from which it has been extracted but also the specific samples among a pool of battery cells. While DCA and EIS data has been widely used in ML models for SOC and SOH estimation, to the best of our knowledge, it has never been used for the authentication of the battery cells [6, 35].

In Table 3, we show the different models that we use. We selected models commonly used in literature, particularly on tasks involving the same type of data that we use [78]. Additionally, in order to optimize their performance, we perform Grid Search by defining possible values for the hyperparameters and training a model for each combination of values. After cross-validating each model and determining its best combination of hyperparameters, we evaluate the best estimator on a test set.

The reader might notice that, in all our models, we did not include any deep or complex network. Indeed, even the simple neural network that we consider, consists of only one hidden layer with a maximum size of 200. While deep networks have been widely used in the literature with data concerning lithium-ion batteries [42, 48], we opted for more lightweight models both in terms of computational time and needed processing power. Indeed, practical implementations of both DCAuth and EISthentication should be affordable and accessible on many different battery-powered devices, which might not possess the necessary hardware to run complex models, or might not be fast enough for authentication purposes. Nonetheless, as will be shown in Section 6, the best models can still obtain high scores while maintaining their lightweight characteristics.

Table 3: Machine Learning models deployed for classification and hyperparameters subject to Grid Search.

Models	**Hyperparameters**
AdaBoost (AB)	• Number of estimators
Decision Tree (DT)	• Criterion • Maximum Depth
Gaussian Naive Bayes (GNB)	• Variance Smoothing
Nearest Neighbors (KNN)	• Number of neighbors • Weight function
Neural Network (NN)	• Hidden layer sizes • Activation function • Solver
Quadratic Discriminant Analysis (QDA)	• Regularization Parameter
Random Forest (RF)	• Criterion • Number of estimators
Support Vector Machine (SVM)	• Kernel • Regularization parameter • Kernel coefficient

6 EVALUATION

We now give an experimental evaluation of our methodologies on the proposed datasets. First, we will disclose the metrics and testing scenarios (Section 6.1). Given the differences in data and processing, we differentiate the two methodologies and evaluate them separately in Section 6.2 and 6.3.

6.1 Metrics

We use four standard metrics to evaluate our models: accuracy, precision, recall, and F1 score. By using True Positives (TP), True Negatives (TN), False Positives (FP), and False Negatives (FN), those metrics are defined as follows.

$$Accuracy = \frac{TN + TP}{TP + TN + FP + FN}, \tag{2}$$

$$Precision = \frac{TP}{TP + FP}, \tag{3}$$

$$Recall = \frac{TP}{TP + FN}, \tag{4}$$

$$F1\ \ Score = 2\frac{Precision \cdot Recall}{Precision + Recall}. \tag{5}$$

While the main focus of our work is the authentication of battery models and architectures, we also study the ability of our models to perform identification. Although the data and features for these two tasks are the same, some core conceptual differences set them apart.

- **Authentication** – The process of verifying whether a battery is genuine or not. As a Machine Learning classification problem, this translates to a binary classification where the only two labels are *authenticated* (or legitimate) and *not authenticated*.
- **Identification** – The process of identifying the type and specifications of a battery. As a Machine Learning classification problem, this translates to a multiclass classification where each label constitutes a battery model or architecture.

Since our datasets are composed of data retrieved from original battery samples cycled in their intended application, in the provided evaluation, a counterfeit model is equivalent to a different battery model. When dealing with authentication, in particular, we are interested in two more metrics: False Acceptance Rate (FAR) and False Rejection Rate (FRR). FAR is the rate at which the authentication system incorrectly accepts an invalid or unauthorized battery as a valid battery. FRR is the rate at which the authentication system incorrectly rejects a valid battery as an invalid battery. These metrics are defined as follows.

$$FAR = \frac{FP}{FP + TN}, \tag{6}$$

$$FRR = \frac{FN}{FN + TP}. \tag{7}$$

6.2 DCAuth

We start our evaluation by first measuring the F1 score of all our ML models in the identification task for DCAuth. To do that, we split the whole dataset into a training set and a test set with respective percentages of 80% and 20%. Further splits to include a validation set are not needed since we are performing 5-fold cross-validation during our Grid Search. When dealing with identification, we perform multiclass classification where the samples for each label have been collected from various datasets. This leads to an imbalanced distribution of the labels, which can affect our processing and generate bias in our results. For this reason, we first balance the number of labels in the dataset by performing random undersampling [38]. While oversampling techniques, such as SMOTE [14], might help us in pursuing the same objective without removing samples from our dataset, we found out that it negatively affected the performance of the classifier, probably due to the high number of features. Results for both architecture authentication and battery model authentication are shown in Figure 5. While most ML models manage to achieve good results, the Random Forest classifier appears to be the best one, obtaining an F1 score close to 1 in battery architecture identification (0.99) and 0.92 in battery model identification. With the sole exception of the Gaussian Naive Bayes classifier, all ML models obtain better results in battery architecture identification than battery model identification. Indeed, the number of architectures is far less than that of battery models (respectively, 5 against 11). Furthermore, battery cells that share the same architecture usually have similar specifications. Thus, DCA plots between cells are comparable, and the extracted features appear to be correlated with each other. More details on the other metrics of evaluation can be found in our repository.

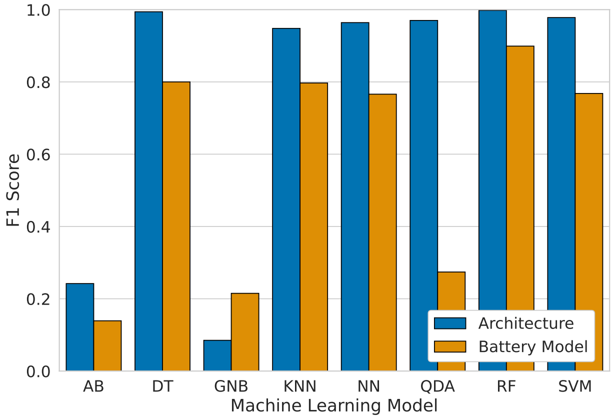

Figure 5: Identification of architectures and battery models in DCAuth.

In the previous evaluation, we considered the dataset composed of data retrieved from both charging and discharging cycles. While also these two types of cycles are equally distributed, their individual effect on the ML model performance might be different. To study this, we divide the dataset into two different parts, one considering charging cycles and the other one considering only discharging cycles. After that, we train different classifiers on each of them and compare the results, which are shown in Figure 6. As we can see, features extracted from the discharging cycles tend to perform slightly better in most of the models (Random Forest included). Indeed, as seen in Table 2, while charging protocols are often similar in different studies (e.g., CCCV charge), discharging techniques instead are more varied and might even depend on dynamic profiles (i.e., driving profiles for EV batteries). Nonetheless, results are still comparable to the ones obtained when considering the whole dataset.

We now move to the authentication task, which effectively translates to binary classification. However, given the number of labels in the dataset, we must consider the effect that an imbalanced distribution can have on the evaluation results. Indeed, despite the undersampling performed in the identification task, when considering

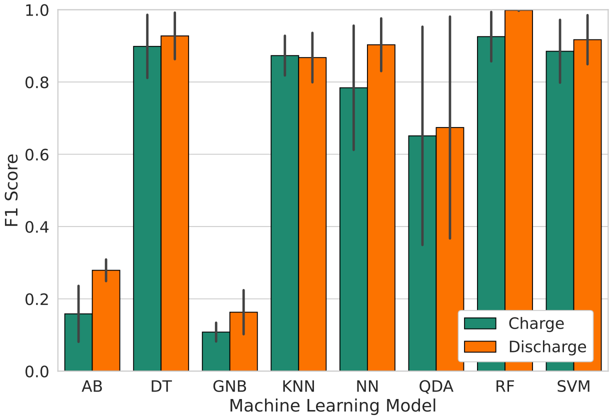

Figure 6: Identification in DCAuth when considering separate charging and discharging cycles. Results are averaged with respect to architecture identification and battery model identification.

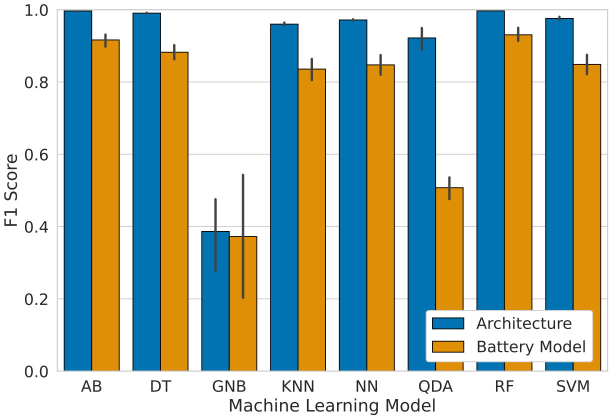

Figure 7: Authentication of architectures and battery models in DCAuth. Results are averaged with respect to the different balance levels for the dataset distribution.

just one label as legitimate and all the others as the not authorized ones, we end up with heavy imbalances in the data distribution. This is expected since many security and privacy scenarios are inherently imbalanced. Two examples are hate speech detection, where datasets are often imbalanced towards the hateful class [24], and intrusion detection systems, where rarer attacks often constitute the minority classes [9]. Thus, we define four different levels of imbalances with respect to legitimate and not authorized classes.

- **50/50** – Samples equally split between the two labels.
- **40/60** – 40% legitimate, 60% not authorized.
- **30/70** – 30% legitimate, 70% not authorized.
- **20/80** – 20% legitimate, 80% not authorized.

As a first type of evaluation for the authentication task, we train all our ML models on all our dataset balances (obtaining a total of $8 \cdot 4 = 32$ models for each task) and compute their F1 score on the test set. While other metrics are published in our repository, we focus on the F1 score in particular since it combines both precision and recall, and we are dealing with imbalanced distribution. Results averaged with respect to the balance levels are shown in Figure 7. These results are acquired by averaging F1 scores from training and testing each model with diverse combinations of genuine and fake battery labels, where a label is treated as legitimate while others as counterfeit. By comparing these results with the ones shown in Figure 5, we notice similar values with the exception of some of the ML models, which are able to obtain better results in a binary classification scenario. Nonetheless, the Random Forest classifier is still the best among the considered ones.

As a second type of evaluation, we study the effect of the different dataset imbalances on the False Acceptance Rate and False Rejection Rate of our models. To do so, we need to average results for architecture authentication and battery model authentication. This is a fair assumption since, for almost all models, the performances of the two tasks are comparable. Results for both FAR and FRR are shown in Figure 8. Unsurprisingly, models performing poorly in authentication have higher FARs and FRRs (e.g., GNB and QDA). It is interesting instead to analyze the trend of those rates across the different distributions. Indeed, while good performers such as RF and AB seem to be almost unaffected by the distribution given their excellent baseline performance, other models seem to present higher or lower performances depending on the balance level and term of evaluation. In particular, FAR is generally higher whenever the distribution is more unbalanced towards the negative class, while FRR has an opposite tendency. This is due to the fact that False Positives are usually more frequent whenever the negative class is the predominant one, and vice versa for False Negatives. Nonetheless, using a balanced dataset seems to be the optimal solution for two main reasons.

1. The effect of the dataset distribution is greater while evaluating FAR with respect to FRR. This means that while the False Rejection Rate might slightly increase with respect to more imbalanced distributions, an imbalanced dataset would otherwise greatly (and negatively) impact the overall performance of the model.
2. The aim of the task is to authenticate one (or more) battery cell in a system. Thus, in order to avoid putting the users at risk, False Acceptance events are more critical than False Rejection events. Nevertheless, having a good variety of samples in the not authorized label can greatly improve the learning process of the models.

A similar analysis to the one shown in Figure 6 (i.e., splitting the datasets in charging and discharging cycles) can also be performed for the authentication task. Given the similarity of the obtained scores, we refrain from including an additional plot to show their values.

6.3 EISthentication

We now move to the evaluation of EISthentication. While the types of evaluations are similar to the ones performed for DCAuth, the characterization of the data is slightly different. Indeed, before

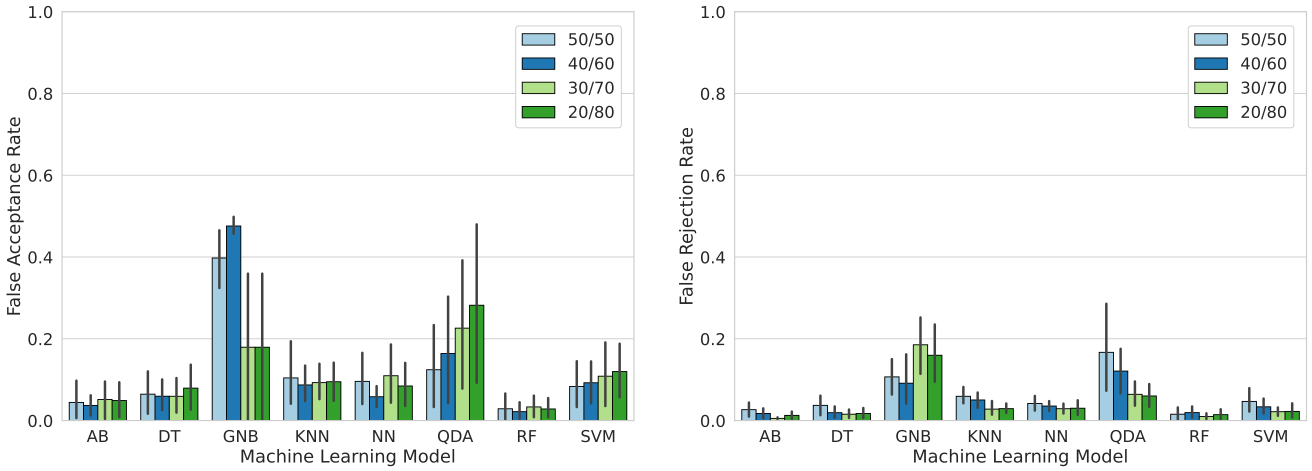

(a): False Acceptance Rate.

(b): False Rejection Rate.

Figure 8: Authentication in DCAuth for each dataset distribution. Results are averaged with respect to architecture authentication and battery model authentication.

splitting the dataset into training set and test set (with the same percentages of DCAuth and using the same 5-fold cross-validation during Grid Search) we filter the features with the *tsfresh* package. Its `feature_selection` package includes a selection method that evaluates the importance of the different extracted features, and thus we keep only the most relevant ones in our datasets. In DCAuth, this procedure does not affect the size of the dataset so much (from 788 features to around 750). In EISthentication, instead, the number of filtered features is almost half that of original features (from 788 to around 400). This could indicate that EIS data might contain some redundancy, and thus results might differ.

As with DCAuth, we start by evaluating the baseline performance of the identification task on both architecture identification and battery model identification. Results are shown in Figure 9. Some differences can be noted while comparing these results with the ones obtained with DCAuth in Figure 5. First, different models produce more consistent results, while in DCAuth, some of them (AB and GNB in particular) performed particularly poorly in both tasks. Secondly, there seems to be a bigger gap in the F1 score between architecture identification and battery model identification. This could mean that the characteristics related to the physical structure and chemistry of the battery extracted through EIS are more directly related to its architecture than its modeling. Lastly, the Random Forest classifier is still the best one among the models considered and reaches F1 scores (0.96 for architecture identification and 0.88 for battery model identification) close to the ones obtained in DCAuth.

Moving to authentication, we notice similar behaviors in the results, which are shown in Figure 10. As with DCAuth, some of the models that had poor performance in identification are now obtaining higher results. Instead, other models that were already performing well in identification now reach approximately the same F1 scores.

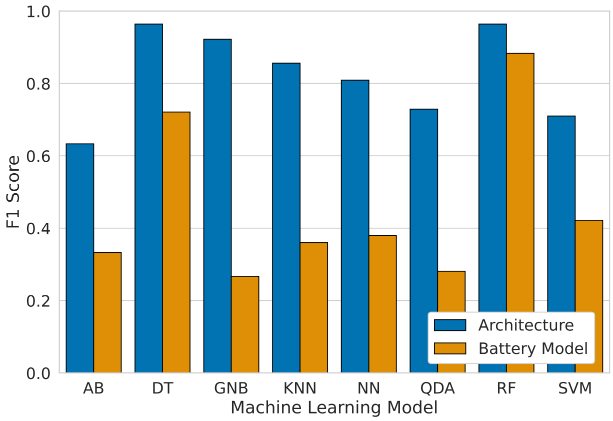

Figure 9: Identification of architectures and battery models in EISthentication.

Even with EISthentication, in the authentication task we are using the same four balance levels for the dataset distribution. Therefore, we can analyze the trend of the False Acceptance Rate and False Rejection Rate at the varying of those balance levels. Results are shown in Figure 11. We can notice the same trends that were present with DCAuth in Figure 8, but they appear to be more uniform across the different models. However, more imbalanced distributions appear to affect both the FAR and the FRR of the models more with respect to DCAuth. Thus, all the more so, in this case, the optimal distribution would be the balanced one for the reasons discussed above.

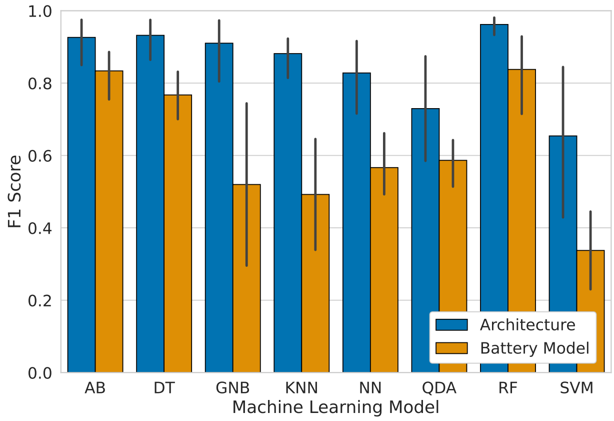

Figure 10: Authentication of architectures and battery models in EISthentication. Results are averaged with respect to the different balance levels for the dataset distribution.

7 DISCUSSION

In this section, we will analyze the obtained results and discuss the possible criticalities of our methodologies. We first analyze the models' performances and use explainable machine learning techniques to have a more detailed insight into the importance of the features (Section 7.1). A comparison of the two presented methodologies is detailed in Section 7.2, while potential limitations are highlighted in Section 7.3. Finally, we discuss the possibility of evasion attacks on the Machine Learning models in Section 7.4.

7.1 Feature Importance

In Section 2.2 and Section 2.3, we detailed the techniques used to retrieve data from the battery cells and made some hypotheses on their goodness for classification. In particular, we speculate that the moving peaks in the DCA plots might be used for authentication and for EIS instead the relationship between different environmental conditions. To confirm our hypothesis, we use Explainable Machine Learning (XAI) techniques on the ML models that we previously trained [55]. We use different techniques on the Random Forest classifier since it is the best performing model in all our tests. Specifically, we use two techniques: Mean Decrease in Impurity (MDI) and SHapley Additive exPlanations (SHAP). The former is a model-specific, global technique that computes the average decrease in Gini impurity or entropy that results from splitting a node using the specific features [26]. Being model-specific, this technique should give us an accurate insight into the model and be faster to compute. On the other hand, this measure suffers from so-called feature selection bias, i.e., it may erroneously assign high MDI values to features that are not highly correlated to the output. For this reason, we also use SHAP, which on the contrary, is model-agnostic [41].

For DCAuth, we obtained similar results with both XAI techniques, which means that the 20 most important techniques extracted with each method greatly overlap. The most important features appear to be related to Continuous Wavelet Transform (CWT) coefficients. It is a technique used in signal processing to analyze the frequency content of a signal and can be used to identify patterns and features in the signal. Another one of the most important features is the number of peaks in the time series, followed by the count of observed values in a specified interval. Lastly, there is also autocorrelation, which is somehow expected since we are dealing with DCA plots in which different behaviors at the early stages of charging/discharging might already indicate a different battery cell or architecture. Combining these features might show that peaks have an important role in the classification and thus might confirm our hypothesis. Moreover, the difference between the average impacts that the features have on the model is small enough to consider them equally important (∼ 0.01 in accuracy).

Similar results are obtained with EISthentication. While autocorrelation and the number of peaks were still present in the top 20 most important features, one of the most important ones appears to be the quantile feature. The feature is often present multiple times with different q values in all our models. This could mean that the distribution of the data points is skewed (we performed the analysis in the balanced scenario) and certain percentiles are more meaningful than others. We also find many Continuous Wavelet Transform (CWT) coefficients and Fast Fourier Transform (FFT) coefficients. Thus, it appears that the model is more interested in the particular distribution of the data points than its specific peak values, indicating that variations in the different data points are more homogeneous on the whole space.

7.2 Methodology Comparison

In Section 6.2 and Section 6.3, we show that both our methodologies are able to obtain high scores in both identification and authentication. What sets these methodologies apart is, of course, the type of data that they use for the processing and, thus, their possible scope of application. Indeed, while DCAuth is the technique that manages to obtain the best results, to perform a prediction, it needs a full charge or discharge cycle for the battery. While this assumption can be fair in some implementations, it has the drawback of being time-consuming (depending on the battery capacity) and degrading the State Of Health of the battery. Indeed, each cycle affects the overall Remaining Useful Life (RUL) of a battery. Nevertheless, some batteries are designed to be stressed for such a large number of cycles that a complete charge or discharge has little to no effect on its SOH. Instead, if the capacity of the battery is too big to wait for a full cycle to be completed, EIS is generally faster. The duration of EIS testing for a lithium-ion battery can vary depending on the chosen frequency range and required accuracy level. For example, a typical EIS measurement for a lithium-ion battery might take around 10 to 30 minutes to complete. Moreover, optimal results for EIS are produced when the battery has been at rest for long periods of time, which incentivize its authentication before its initial usage [37]. While performances in authentication might be slightly lower than DCAuth, EISthentication can be a valuable technique given the increasing availability of EIS measurements or estimations in different systems [30, 40, 67]. Both these methodologies have also proven to be computationally lightweight, making their usage affordable on many different low-end devices.

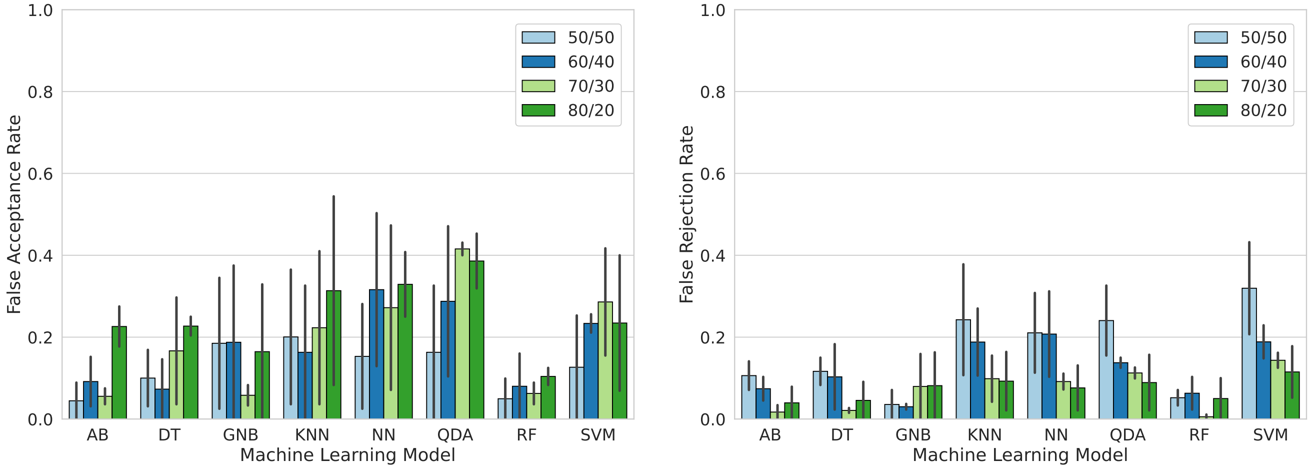

(a): False Acceptance Rate.

(b): False Rejection Rate.

Figure 11: Authentication in EISthentication for each dataset distribution. Results are averaged with respect to architecture authentication and battery model authentication.

7.3 Limitations

In our work, we considered a total of 20 datasets, accounting for 17 different battery models and 5 different battery architectures. Instead, in real-world applications, millions of batteries might potentially be used for a specific scenario, and thus, the number of possible counterfeit samples becomes even larger. In practice, this could increase the number of false positives and false negatives in our approaches, which could affect their accuracy, leading to some issues if applied to critical contexts. To partially address this, selecting a subset of authorized battery models might be possible and then appropriately adjusting the dataset balance by incorporating data labeled as "counterfeit" from different battery models. Moreover, it is unnecessary to consider every battery model for authentication in a single application. For example, though a 18350 battery can be disguised as a 18650 battery through rewrapping, the opposite scenario is physically impossible. Furthermore, adopting such a strategy to conceal a counterfeit battery wouldn't yield economic advantages for an attacker seeking to market an inferior battery at the cost of a genuine one.

Another aspect that could potentially alter performance with respect to the results shown in this paper is the aging of the battery cells. In our work, we collected datasets that contain data retrieved from batteries cycled in many different stages of their lifespan, as described in Table 2 and Section 4.1.2. However, as also shown in Figure 2, several aspects such as SOH and SOC greatly affect the data distribution from which we extract features. Thus, data extracted from batteries cycled in a limited set of conditions might not be authenticated if the dataset used to train the models did not account for that specific combination of parameters. To account for this, we advise the usage of varied datasets for training and collecting data from batteries in many stages of their lifecycle.

7.4 Evasion Attacks

By only leveraging the electrochemical properties of the battery, DCAuth and EISthentication are resilient against most attacks while also providing scalability. However, to achieve these results, our methodologies use Machine Learning models, which have been proven vulnerable to adversarial attacks [21]. In particular, evasion attacks represent a subset of adversarial attacks focused on altering input data to mislead Machine Learning models. These attacks could potentially affect the authentication results of our methodologies since even our best classifier (i.e., Random Forest) has been shown to be vulnerable under specific conditions [4, 74]. However, considering our system model, the implementation of these attacks in most applications is not currently practical due to the following reasons.

- **Attacker Knowledge** – To successfully evade authentication, the attacker needs access to the inner parameters of the target Machine Learning model. Since this information is usually hidden from the attacker, real-world attacks could differ significantly from the scenarios outlined in the literature [3].
- **Model Obfuscation** – To gain access to the model or its training dataset, attackers can use techniques such as model inversion. However, procedures such as model obfuscation can prevent it and protect the model parameters from unauthorized access.
- **Adversarial Transferability** – The knowledge that an attacker can gain to evade detection is not limited to the model architecture. By leveraging the dataset distribution and ground truth balance, malicious actors can craft evasion attacks that are "transferable" among different models. However, it has been shown how surrogate models struggle to evade most models impactfully [2].

8 CONCLUSIONS

In conclusion, the rise of counterfeit Li-ion batteries is a serious issue that demands the attention of the battery industry. Counterfeit batteries can cause significant safety hazards and economic losses, and detecting them is becoming increasingly difficult as counterfeiters become more sophisticated. Furthermore, current techniques for battery authentication can be fooled by various attacks or are not scalable to several battery models and architectures.

Contribution. In this paper, we explored the use of Machine Learning techniques to counter the spread of these counterfeit samples. Through several ML models, we are able to formalize an authentication method that is resilient against several attacks and scalable to many different devices. In particular, we focused on the use of authentication and identification mechanisms that can rapidly and efficiently detect the legitimacy of a battery. We proposed two different methodologies, **DCAuth** and **EISthentication**, which respectively use Differential Capacity Analysis data and Electrochemical Impedance Spectroscopy data to perform authentication on each battery cell. The data is retrievable through the regular usage of the battery, and thus our techniques do not rely on any external devices. Our models are able to detect both battery Li-ion architectures and battery models with high accuracy, obtaining F1 scores that reach up to 0.94 with the best classifier (i.e., Random Forest in battery model authentication).

Future Works. As stated in Section 4.1, all the datasets used in this paper have been found in the literature. While the overall number of battery models and architectures is still relevant, real-world implementations might consider several more. Unfortunately, collecting this type of data is often expensive and requires specific expertise. We believe that our methodologies can have a great impact on the battery industry and can benefit user safety; thus, we invite researchers to integrate our results with their data, when available, and report their findings. Through practical implementations of our methodologies, the counterfeit battery market can be opposed by ensuring the legitimacy of each battery cell and thus guaranteeing the safety of users.

REFERENCES

[1] Ahmad Al Khas and Ihsan Cicek. 2019. SHA-512 based wireless authentication scheme for smart battery management systems. In *2019 8th International Conference on Renewable Energy Research and Applications (ICRERA)*. IEEE, 968–972.

[2] Marco Alecci, Mauro Conti, Francesco Marchiori, Luca Martinelli, and Luca Pajola. 2023. Your Attack Is Too DUMB: Formalizing Attacker Scenarios for Adversarial Transferability. *arXiv preprint arXiv:2306.15363* (2023).

[3] Giovanni Apruzzese, Hyrum S Anderson, Savino Dambra, David Freeman, Fabio Pierazzi, and Kevin A Roundy. 2022. " Real Attackers Don't Compute Gradients": Bridging the Gap Between Adversarial ML Research and Practice. *arXiv preprint arXiv:2212.14315* (2022).

[4] Giovanni Apruzzese and Michele Colajanni. 2018. Evading botnet detectors based on flows and random forest with adversarial samples. In *2018 IEEE 17th International Symposium on Network Computing and Applications (NCA)*. IEEE, 1–8.

[5] Peter M Attia, Aditya Grover, Norman Jin, Kristen A Severson, Todor M Markov, Yang-Hung Liao, Michael H Chen, Bryan Cheong, Nicholas Perkins, Zi Yang, et al. 2020. Closed-loop optimization of fast-charging protocols for batteries with machine learning. *Nature* 578, 7795 (2020), 397–402.

[6] Iman Babaeiyazdi, Afshin Rezaei-Zare, and Shahab Shokrzadeh. 2021. State of charge prediction of EV Li-ion batteries using EIS: A machine learning approach. *Energy* 223 (2021), 120116.

[7] Heather M Barkholtz, Armando Fresquez, Babu R Chalamala, and Summer R Ferreira. 2017. A database for comparative electrochemical performance of commercial 18650-format lithium-ion cells. *Journal of The Electrochemical Society* 164, 12 (2017), A2697.

[8] Anthony Barré, Benjamin Deguilhem, Sébastien Grolleau, Mathias Gérard, Frédéric Suard, and Delphine Riu. 2013. A review on lithium-ion battery ageing mechanisms and estimations for automotive applications. *Journal of Power Sources* 241 (2013), 680–689.

[9] Punam Bedi, Neha Gupta, and Vinita Jindal. 2021. I-SiamIDS: an improved Siam-IDS for handling class imbalance in network-based intrusion detection systems. *Applied Intelligence* 51 (2021), 1133–1151.

[10] C Birkl. 2017. Oxford Battery Degradation Dataset 1.

[11] Lorenzo A Buscaglia, Osvaldo N Oliveira, and João Paulo Carmo. 2021. Roadmap for electrical impedance spectroscopy for sensing: a tutorial. *IEEE Sensors Journal* 21, 20 (2021), 22246–22257.

[12] Lisa Calearo, Charalampos Ziras, Andreas Thingvad, and Mattia Marinelli. 2022. Agnostic Battery Management System Capacity Estimation for Electric Vehicles. *Energies* 15, 24 (2022), 9656.

[13] Byoung-Yong Chang and Su-Moon Park. 2010. Electrochemical impedance spectroscopy. *Annual Review of Analytical Chemistry* 3, 1 (2010), 207.

[14] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. 2002. SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research* 16 (2002), 321–357.

[15] Maximilian Christ, Nils Braun, Julius Neuffer, and Andreas W Kempa-Liehr. 2018. Time series feature extraction on basis of scalable hypothesis tests (tsfresh–a python package). *Neurocomputing* 307 (2018), 72–77.

[16] U.S. Customs and Border Protection. 2016. *CBP Seizes Record Amount of Counterfeit Hoverboards.*

[17] Arnaud Devie, George Baure, and Matthieu Dubarry. 2018. Intrinsic variability in the degradation of a batch of commercial 18650 lithium-ion cells. *Energies* 11, 5 (2018), 1031.

[18] Gonçalo Dos Reis, Calum Strange, Mohit Yadav, and Shawn Li. 2021. Lithium-ion battery data and where to find it. *Energy and AI* 5 (2021), 100081.

[19] Xuning Feng, Yu Merla, Caihao Weng, Minggao Ouyang, Xiangming He, Bor Yann Liaw, Shriram Santhanagopalan, Xuemin Li, Ping Liu, Languang Lu, et al. 2020. A reliable approach of differentiating discrete sampled-data for battery diagnosis. *ETransportation* 3 (2020), 100051.

[20] Erik Goldammer, Marius Gentejohann, Michael Schlüter, Daniel Weber, Wolfgang Wondrak, Sibylle Dieckerhoff, Clemens Gühmann, and Julia Kowal. 2022. The Impact of an Overlaid Ripple Current on Battery Aging: The Development of the SiCWell Dataset. *Batteries* 8, 2 (2022), 11.

[21] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014).

[22] Jaykanth Govindarajan. 2021. Lifecycle ageing tests on commercial 18650 Li ion cell @ 10°C and 0°C. https://doi.org/10.4121/14377295.v1

[23] GrandViewResearch. 2022. *Lithium-ion Battery Market Size, Share & Trends Analysis Report By Product (LCO, LFP, NCA, LMO, LTO, NMC), By Application (Consumer Electronics, Energy Storage Systems, Industrial), By Region, And Segment Forecasts, 2022 - 2030.*

[24] Tommi Gröndahl, Luca Pajola, Mika Juuti, Mauro Conti, and N Asokan. 2018. All you need is" love" evading hate speech detection. In *Proceedings of the 11th ACM workshop on artificial intelligence and security.* 2–12.

[25] Defne Gun, Hector Perez, and Scott Moura. 2015. Fast Charging Tests. https://datadryad.org/stash/dataset/doi:10.6078/D1MS3X. https://doi.org/10.6078/D1MS3X

[26] Hong Han, Xiaoling Guo, and Hua Yu. 2016. Variable selection using mean decrease accuracy and mean decrease gini based on random forest. In *2016 7th ieee international conference on software engineering and service science (icsess)*. IEEE, 219–224.

[27] Wei He, Nicholas Williard, Chaochao Chen, and Michael Pecht. 2014. 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* 62 (2014), 783–791.

[28] Wei He, Nicholas Williard, Michael Osterman, and Michael Pecht. 2011. Prognostics of lithium-ion batteries based on Dempster–Shafer theory and the Bayesian Monte Carlo method. *Journal of Power Sources* 196, 23 (2011), 10314–10321.

[29] Thomas MM Heenan, A Jnawali, MDR Kok, Thomas George Tranter, C Tan, A Dimitrijevic, R Jervis, DJL Brett, and PR Shearing. 2020. An advanced microstructural and electrochemical datasheet on 18650 Li-ion batteries with nickel-rich NMC811 cathodes and graphite-silicon anodes. *Journal of The Electrochemical Society* 167, 14 (2020), 140530.

[30] David A Howey, Paul D Mitcheson, Vladimir Yufit, Gregory J Offer, and Nigel P Brandon. 2013. Online measurement of battery impedance using motor controller excitation. *IEEE transactions on vehicular technology* 63, 6 (2013), 2557–2566.

[31] Bo Jiang, Jiangong Zhu, Xueyuan Wang, Xuezhe Wei, Wenlong Shang, and Haifeng Dai. 2022. A comparative study of different features extracted from electrochemical impedance spectroscopy in state of health estimation for lithium-ion batteries. *Applied Energy* 322 (2022), 119502.

[32] Daniel Juarez-Robles, Judith A Jeevarajan, and Partha P Mukherjee. 2020. Degradation-safety analytics in lithium-ion cells: Part I. Aging under

charge/discharge cycling. *Journal of The Electrochemical Society* 167, 16 (2020), 160510.

[33] Lingxi Kong, Diganta Das, and Michael G Pecht. 2022. The Distribution and Detection Issues of Counterfeit Lithium-Ion Batteries. *Energies* 15, 10 (2022), 3798.

[34] Amelie Krupp, Ernst Ferg, Frank Schuldt, Karen Derendorf, and Carsten Agert. 2020. Incremental capacity analysis as a state of health estimation method for lithium-ion battery modules with series-connected cells. *Batteries* 7, 1 (2020), 2.

[35] Peter Kurzweil, Wolfgang Scheuerpflug, Bernhard Frenzel, Christian Schell, and Josef Schottenbauer. 2022. Differential Capacity as a Tool for SOC and SOH Estimation of Lithium Ion Batteries Using Charge/Discharge Curves, Cyclic Voltammetry, Impedance Spectroscopy, and Heat Events: A Tutorial. *Energies* 15, 13 (2022), 4520.

[36] Andrzej Lasia. 2002. Electrochemical impedance spectroscopy and its applications. In *Modern aspects of electrochemistry*. Springer, 143–248.

[37] Hugo Leduc, Russell Okamura, Eru Kyeyune-Nyombi Jr, Kenny Huynh, and Steven Chung. 2020. Real-Time Under-Load Electrochemical Impedance Spectroscopy (EIS) Analysis and Modeling. In *Electrochemical Society Meeting Abstracts prime2020*. The Electrochemical Society, Inc., 1584–1584.

[38] Guillaume Lemaître, Fernando Nogueira, and Christos K Aridas. 2017. Imbalanced-learn: A python toolbox to tackle the curse of imbalanced datasets in machine learning. *The Journal of Machine Learning Research* 18, 1 (2017), 559–563.

[39] Dezhi Li, Dongfang Yang, Liwei Li, Licheng Wang, and Kai Wang. 2022. Electrochemical impedance spectroscopy based on the state of health estimation for lithium-ion batteries. *Energies* 15, 18 (2022), 6665.

[40] Nils Lohmann, Peter Haussmann, Patrick Wesskamp, Joachim Melbert, and Thomas Musch. 2015. Employing real automotive driving data for electrochemical impedance spectroscopy on lithium-ion cells. *SAE International Journal of Alternative Powertrains* 4, 2 (2015), 308–317.

[41] Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems* 30 (2017).

[42] Bin Ma, Shichun Yang, Lisheng Zhang, Wentao Wang, Siyan Chen, Xianbin Yang, Haicheng Xie, Hanqing Yu, Huizhi Wang, and Xinhua Liu. 2022. Remaining useful life and state of health prediction for lithium batteries based on differential thermal voltammetry and a deep-learning model. *Journal of Power Sources* 548 (2022), 232030.

[43] Mario Marinaro, Dominic Bresser, Engelbert Beyer, Peter Faguy, Kei Hosoi, Hong Li, Julija Sakovica, Khalil Amine, Margret Wohlfahrt-Mehrens, and Stefano Passerini. 2020. Bringing forward the development of battery cells for automotive applications: Perspective of R&D activities in China, Japan, the EU and the USA. *Journal of Power Sources* 459 (2020), 228073.

[44] Nina Meddings, Marco Heinrich, Frédéric Overney, Jong-Sook Lee, Vanesa Ruiz, Emilio Napolitano, Steffen Seitz, Gareth Hinds, Rinaldo Raccichini, Miran Gaberšček, et al. 2020. Application of electrochemical impedance spectroscopy to commercial Li-ion cells: A review. *Journal of Power Sources* 480 (2020), 228742.

[45] Gholam-Abbas Nazri and Gianfranco Pistoia. 2008. *Lithium batteries: science and technology.* Springer Science & Business Media.

[46] Michael O'Brien and Gregg Tatarka. 2008. *Lithium-Ion Batteries: An Emerging Focus of Causation in Consumer Product Fires.* White Paper. Wilson Elser Moskowitz & Dicker LLP.

[47] Jarred Z Olson, Carmen M López, and Edmund JF Dickinson. 2023. Differential Analysis of Galvanostatic Cycle Data from Li-Ion Batteries: Interpretative Insights and Graphical Heuristics. *Chemistry of Materials* 35, 4 (2023), 1487–1513.

[48] Simona Pepe, Jiapeng Liu, Emanuele Quattrocchi, and Francesco Ciucci. 2022. Neural ordinary differential equations and recurrent neural networks for predicting the state of health of batteries. *Journal of Energy Storage* 50 (2022), 104209.

[49] Yuliya Preger, Heather M Barkholtz, Armando Fresquez, Daniel L Campbell, Benjamin W Juba, Jessica Romàn-Kustas, Summer R Ferreira, and Babu Chalamala. 2020. Degradation of commercial lithium-ion cells as a function of chemistry and cycling conditions. *Journal of The Electrochemical Society* 167, 12 (2020), 120532.

[50] T Raj. 2020. Path Dependent Battery Degradation Dataset Part 1.

[51] Trishna Raj, Andrew A Wang, Charles W Monroe, and David A Howey. 2020. Investigation of path-dependent degradation in lithium-ion batteries. *Batteries & Supercaps* 3, 12 (2020), 1377–1385.

[52] Li Ran, Wu Junfeng, Wang Haiying, and Li Gechen. 2010. Prediction of state of charge of lithium-ion rechargeable battery with electrochemical impedance spectroscopy theory. In *2010 5th IEEE Conference on Industrial Electronics and Applications*. IEEE, 684–688.

[53] J M Reniers, G Mulder, and D A Howey. 2020. Oxford energy trading battery degradation dataset.

[54] M Rezal, A Zulaikha, M Sabri, R Yusof, and S Ridzwan. 2014. Orion battery management system (BMS) for lithium-ion battery pack. In *Proceedings of the Colloquium of Education, Engineering & Technology (COLEET 2014) at Universiti Kuala Lumpur Malaysian Spanish Institute, Kulim, Kedah, Malaysia*. 80–86.

[55] Ribana Roscher, Bastian Bohn, Marco F Duarte, and Jochen Garcke. 2020. Explainable machine learning for scientific insights and discoveries. *Ieee Access* 8 (2020), 42200–42216.

[56] Tsuyoshi Sasaki, Yoshio Ukyo, and Petr Novák. 2013. Memory effect in a lithium-ion battery. *Nature materials* 12, 6 (2013), 569–575.

[57] Abraham Savitzky and Marcel JE Golay. 1964. Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry* 36, 8 (1964), 1627–1639.

[58] Saurabh Saxena, Christopher Hendricks, and Michael Pecht. 2016. Cycle life testing and modeling of graphite/LiCoO2 cells under different state of charge ranges. *Journal of Power Sources* 327 (2016), 394–400.

[59] Saurabh Saxena, Lingxi Kong, and Michael G Pecht. 2018. Exploding e-cigarettes: A battery safety issue. *IEEE Access* 6 (2018), 21442–21466.

[60] Bruno Scrosati and Jürgen Garche. 2010. Lithium batteries: Status, prospects and future. *Journal of power sources* 195, 9 (2010), 2419–2430.

[61] Kristen A Severson, Peter M Attia, Norman Jin, Nicholas Perkins, Benben Jiang, Zi Yang, Michael H Chen, Muratahan Aykol, Patrick K Herring, Dimitrios Fraggedakis, et al. 2019. Data-driven prediction of battery cycle life before capacity degradation. *Nature Energy* 4, 5 (2019), 383–391.

[62] Nyle Siddiqui, Laura Pryor, and Rushit Dave. 2021. User authentication schemes using machine learning methods—a review. In *Proceedings of International Conference on Communication and Computational Technologies*. Springer, 703–723.

[63] R Smith, S Shahidinejad, D Blair, and EL Bibeau. 2011. Characterization of urban commuter driving profiles to optimize battery size in light-duty plug-in electric vehicles. *Transportation Research Part D: Transport and Environment* 16, 3 (2011), 218–224.

[64] Nicole L Thompson, Theodore A Cohen, Sarah Alamdari, Chih-Wei Hsu, Grant A Williamson, Vincent C Holmberg, et al. 2020. DiffCapAnalyzer: A Python Package for Quantitative Analysis of Total Differential Capacity Data. *Journal of Open Source Software* 5, 54 (2020), 2624.

[65] Soichiro Torai, Masaru Nakagomi, Satoshi Yoshitake, Shuichiro Yamaguchi, and Noboru Oyama. 2016. State-of-health estimation of LiFePO4/graphite batteries based on a model using differential capacity. *Journal of Power Sources* 306 (2016), 62–69.

[66] Khiem Trad. 2021. Lifecycle ageing tests on commercial 18650 Li ion cell @ 25°C and 45°C. https://doi.org/10.4121/13739296.v1

[67] Uwe Tröltzsch, Olfa Kanoun, and Hans-Rolf Tränkler. 2006. Characterizing aging effects of lithium ion batteries by impedance spectroscopy. *Electrochimica acta* 51, 8-9 (2006), 1664–1672.

[68] Perry Tsao. 2018. *How Secure Fuel Gauges Can Prevent Battery Counterfeiting.* White Paper. Maxim Integrated.

[69] Wladislaw Waag, Christian Fleischer, and Dirk Uwe Sauer. 2014. Critical review of the methods for monitoring of lithium-ion batteries in electric and hybrid vehicles. *Journal of Power Sources* 258 (2014), 321–339.

[70] U Westerhoff, T Kroker, K Kurbach, and M Kurrat. 2016. Electrochemical impedance spectroscopy based estimation of the state of charge of lithium-ion batteries. *Journal of Energy Storage* 8 (2016), 244–256.

[71] Nick Williard, Wei He, Michael Osterman, and Michael Pecht. 2013. Comparative analysis of features for determining state of health in lithium-ion batteries. *International Journal of Prognostics and Health Management* 4, 1 (2013).

[72] Yinjiao Xing, Wei He, Michael Pecht, and Kwok Leung Tsui. 2014. State of charge estimation of lithium-ion batteries using the open-circuit voltage at various ambient temperatures. *Applied Energy* 113 (2014), 106–115.

[73] Yinjiao Xing, Eden WM Ma, Kwok-Leung Tsui, and Michael Pecht. 2013. An ensemble model for predicting the remaining useful performance of lithium-ion batteries. *Microelectronics Reliability* 53, 6 (2013), 811–820.

[74] Chong Zhang, Huan Zhang, and Cho-Jui Hsieh. 2020. An efficient adversarial attack for tree ensembles. *Advances in Neural Information Processing Systems* 33 (2020), 16165–16176.

[75] Yunwei Zhang, Qiaochu Tang, Yao Zhang, Jiabin Wang, Ulrich Stimming, and Alpha A Lee. 2020. Identifying degradation patterns of lithium ion batteries from impedance spectroscopy using machine learning. *Nature communications* 11, 1 (2020), 1–6.

[76] Jingyuan Zhao, Heping Ling, Junbin Wang, Andrew F Burke, and Yubo Lian. 2022. Data-driven prediction of battery failure for electric vehicles. *Iscience* 25, 4 (2022), 104172.

[77] Fangdan Zheng, Yinjiao Xing, Jiuchun Jiang, Bingxiang Sun, Jonghoon Kim, and Michael Pecht. 2016. Influence of different open circuit voltage tests on state of charge online estimation for lithium-ion batteries. *Applied energy* 183 (2016), 513–525.

[78] Shan Zhu, Xinyang Sun, Xiaoyang Gao, Jianrong Wang, Naiqin Zhao, and Junwei Sha. 2019. Equivalent circuit model recognition of electrochemical impedance spectroscopy via machine learning. *Journal of Electroanalytical Chemistry* 855 (2019), 113627.

[79] Quan-Chao Zhuang, Xiang-Yun Qiu, Shou-Dong Xu, Ying-Huai Qiang, and SG Su. 2012. Diagnosis of electrochemical impedance spectroscopy in lithium-ion batteries. *Lithium Ion Batteries—New Developments* 8 (2012), 189–227.