





Article

# Aging State Classification of Lithium-Ion Batteries in a Low-Dimensional Latent Space

Limei Jin <sup>1,2,3</sup> , Franz Philipp Bereck <sup>1,3</sup>, Rüdiger-A. Eichel <sup>1,4,5</sup> , Josef Granwehr <sup>1,3</sup>   
and Christoph Scheurer <sup>1,2,\*</sup> 

<sup>1</sup> Institute of Energy Technologies, Fundamental Electrochemistry (IET-1), Forschungszentrum Jülich, 52425 Jülich, Germany

<sup>2</sup> Fritz-Haber-Institut der Max-Planck-Gesellschaft, Faradayweg 4-6, 14195 Berlin, Germany

<sup>3</sup> Institute of Technical and Macromolecular Chemistry, RWTH Aachen University, 52056 Aachen, Germany

<sup>4</sup> Institute of Physical Chemistry, RWTH Aachen University, 52062 Aachen, Germany

<sup>5</sup> Faculty of Mechanical Engineering, RWTH Aachen University, 52062 Aachen, Germany

\* Correspondence: [scheurer@fhi.mpg.de](mailto:scheurer@fhi.mpg.de)

## Abstract

Battery datasets, whether gathered experimentally or through simulation, are typically high-dimensional and complex, which complicates the direct interpretation of degradation behavior or anomaly detection. To overcome these limitations, this study introduces a framework that compresses battery signals into a low-dimensional representation using an autoencoder, enabling the extraction of informative features for state analysis. A central component of this work is the systematic comparison of latent representations obtained from two fundamentally different data sources: frequency-domain impedance data and time-domain voltage-current data. The close agreement of aging trajectories in both representations suggests that information traditionally derived from impedance analysis can also be captured directly from raw time-series signals. To better approximate real operating conditions, synthetic datasets are augmented with stochastic perturbations. In this context, latent spaces learned from idealized periodic inputs are contrasted with those derived from permuted and noise-contaminated signals. The resulting low-dimensional features are subsequently evaluated through a support vector machine with both linear and non-linear kernel functions, allowing the categorization of battery states into *fresh*, *aged* and *damaged* conditions. The results demonstrate that the progression of battery degradation is consistently reflected in the latent space, independent of the input domain or signal quality. This robustness indicates that the proposed approach can effectively capture essential aging characteristics even under non-ideal conditions. Consequently, this framework provides a basis for developing advanced diagnostic strategies, including the design of pseudo-random excitation profiles for improved battery state assessment and optimized operational control.



Academic Editor: Zhenbo Wang

Received: 5 March 2026

Revised: 29 March 2026

Accepted: 2 April 2026

Published: 7 April 2026

**Copyright:** © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article

distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)

[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

**Keywords:** Electrochemical Impedance Spectroscopy; Convolutional Autoencoder; Latent Space; Aging Pathway

## 1. Introduction

The accelerating adoption of electric vehicles as a sustainable transportation solution has significantly increased the demand for high-performance battery systems [1]. However, over repeated charge–discharge cycles, batteries undergo gradual degradation, driven by changes in electrochemical reactions, interfacial properties, and internal structures,

ultimately leading to performance loss. Understanding these multifactorial battery degradation phenomena is essential for optimizing dynamic operation. Various approaches have been developed to monitor and predict battery lifetime.

Physics-based electrochemical models, such as the P2D model [2,3], aim to capture the underlying electrochemical processes and degradation mechanisms in lithium-ion batteries. These models incorporate factors such as Solid Electrolyte Interphase (SEI) growth [4], lithium plating, active material loss, and mechanical degradation. The primary advantage of these models lies in their ability to provide insights into the fundamental aging mechanisms. However, they are often computationally intensive and require detailed knowledge of battery materials and design parameters. Data-driven approaches leverage machine learning and statistical techniques to predict battery lifetime based on historical cycling data. Methods, such as neural networks, support vector machines, and Gaussian process regression [5], capture complex, non-linear relationships in battery aging data without requiring in-depth knowledge of the underlying physics. However, these models require large datasets for training and may not generalize well to new operating conditions. Empirical models, such as Equivalent Circuit Models (ECM) [6], use simplified mathematical expressions to describe capacity fading and resistance growth based on experimental observations. While less detailed than physics-based models, such empirical approaches are computationally efficient and suitable for real-time applications in battery management systems.

Modelling with ECM is typically combined with Electrochemical Impedance Spectroscopy (EIS), which is a powerful diagnostic tool for analyzing battery internal states. EIS-based models typically use frequency-domain measurements to estimate equivalent circuit parameters and monitor changes in electrode and electrolyte properties over time. While EIS is effective for tracking specific degradation mechanisms, it commonly assumes that the system behaves linear and time-invariant (LTI). In this context, a linear system is one where the output is directly proportional to the input, which is rarely satisfied in real battery applications due to non-linear degradation processes. In particular, EIS cannot fully capture transient responses under non-periodic excitation [7], which violate the stationarity assumption required for frequency-domain analysis. EIS also reflects system behavior at a single point in time, often represented as State of Charge (SOC) [8]. It struggles to account for non-periodic dynamics or time-dependent changes such as State of Health (SOH), which are critical in real-world, time-varying conditions. To address such limitations, the analysis should transition from the frequency domain to the time domain.

The analysis of time-domain battery data involves handling large, high-dimensional datasets. Machine learning offers transformative potential by leveraging these datasets to extract meaningful features in a nonlinear manner, thereby enhancing the accuracy and reliability of battery condition assessment [9]. In particular, it enables the model to exploit information embedded in transient behavior, relaxation effects, and measurement noise, which are typically neglected in conventional methods. Dimensionality reduction methods, such as autoencoders, are particularly useful for isolating the key features and focusing on the essential information within the data. In unsupervised learning, autoencoders are widely used to derive compact feature representations and perform dimensionality reduction [10]. The architecture includes two primary components, namely an encoder and a decoder. The encoder maps the normalized input into a lower-dimensional latent representation (latent space, LS), while the decoder reconstructs the original data from this compressed form. The model is trained by minimizing the difference between the input and the reconstructed output. Projecting battery data with different aging states into a low-dimensional latent space enables clearer visualization of the degradation process over time. While previous works such as that by Che et al. [11] relied on features from discharging

capacity-voltage curves, this approach is less effective for battery chemistries such as lithium iron phosphate (LFP), where flat voltage profiles compromise SOH sensitivity. To address this, the present study uses direct current and voltage pairs as inputs for training autoencoders, unlocking time-series data for robust aging diagnostics.

This article explores the shift from linear to non-linear regimes, addressing the limitations of conventional EIS-based analysis. We demonstrate the potential for SOH classification directly within a two-dimensional latent space. This is a significant advantage over many current studies that do not fully utilize the interpretability of a 2D representation, which allows for direct visual inspection of the aging process. Our approach also validates the possibility of transforming from frequency-domain impedance data to time-domain current-voltage data, showing that the aging patterns remain consistent across both data types. This provides confidence that methods traditionally used for laboratory-based EIS can be applied to more readily available time-series data. Furthermore, we demonstrate that the robust aging pattern remains consistent even when using permuted noisy data as an input. This is crucial for implementing robust battery diagnostics that handle the variability and noise inherent in real-world operating conditions. By addressing these points, our research fills the gap of directly visualizing and classifying battery aging states in a low-dimensional, easily interpretable latent space while demonstrating the model's robustness to different data types and noise. This approach demonstrates a pathway toward live tracking of the aging process and convincingly shows that latent feature spaces can capture key aspects such as capacity decay and recovery, bolstering their value for battery diagnostics.

## 2. Methodology

### 2.1. Structure of the Autoencoder

In this work, the decoder is omitted, and only the encoder is utilized, since the objective is to observe battery aging behavior in the latent representation. The Rectified Linear Unit (ReLU) activation function is employed for its simplicity and effectiveness, while Convolutional Neural Network (CNN) layers are utilized to handle the multi-channel input data [12]. A Convolutional Neural Network (CNN) architecture was chosen for this analysis, primarily because of its effectiveness in processing multi-channel time-series data, which can be treated as a 2D object. Unlike models such as Time-Series Foundation Models (Times FM) or Patch Time-Series Transformers (Patch TST) [13], which are designed for handling very long time series or predicting future values, our approach focuses on extracting features from the entire input sequence to create an explicit latent space that represents the overall state of the battery. The convolutional layers are particularly suited for capturing local dependencies and hierarchical features across the multiple input channels (current, voltage, and impedance data), enabling the network to map a non-linear transformation from the high-dimensional input to a low-dimensional latent space. The dimensionality of the latent space is set to two, sufficient for categorization and convenient for visualization. While some applications might require higher-dimensional latent spaces, a two-dimensional representation is sufficient here and offers improved interpretability for tracking battery aging. A one-dimensional latent space was not sufficient as it could not capture the complexity of the data without significant information loss, while a higher-dimensional latent space (e.g., 3D) would not provide the same ease of visualization and interpretability for this specific classification task.

### 2.2. Data Formatting for Autoencoder Input and Output Layers

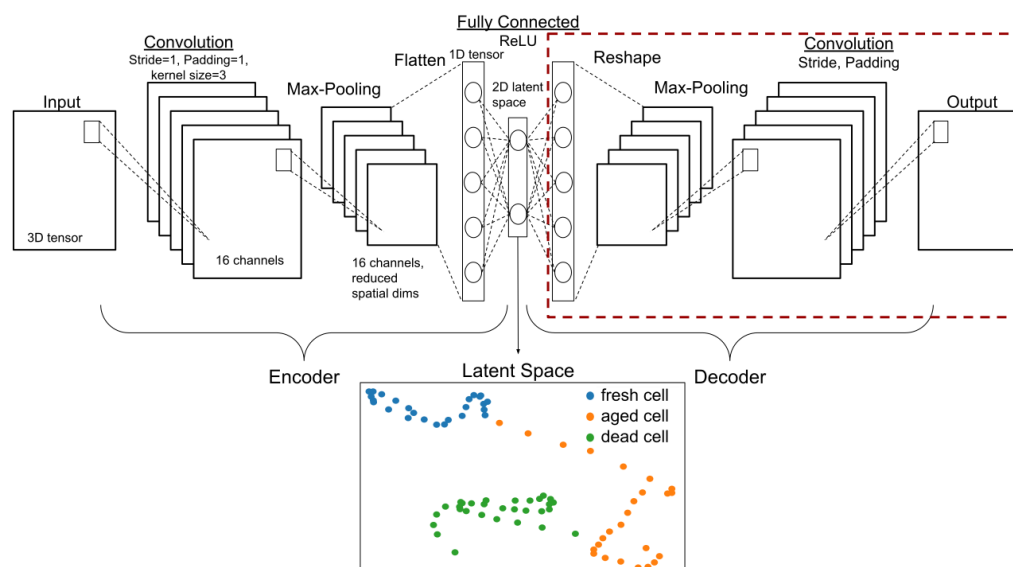
The analysis incorporates both time-series and frequency-domain input data captured across a single discharging cycle at multiple SOC conditions. The frequency-domain data includes real and imaginary parts of the impedance, while the time-based input consists

of a load current and corresponding voltage data. As a result, each sample is represented as a multi-channel tensor combining current, voltage, and SOC information, leading to a three-dimensional input structure for the autoencoder. Due to the sensitivity of neural networks to input scaling, min-max normalization is applied to each sequence [14]. The normalization ensures that all input values are scaled to a uniform range, typically between 0 and 1. It prevents the dominance of any single sequence during training and promotes stable model convergence.

The normalized three-channel input sequence is passed through the encoder, which applies a series of convolutional transformations across multiple layers. Unlike standard 1D convolutions used for single time-series signals, we employ a  $3 \times 3 \times 3$  convolutional kernel to jointly capture correlations across the different signal channels and along the sequence dimension. This design allows the network to learn cross-channel interactions and local coupled patterns between current, voltage, and SOC. Such joint feature extraction is essential for identifying aging-related behavior, as degradation mechanisms manifest through interdependent changes across these signals rather than within any single sequence alone.

The first convolutional layer employs a  $3 \times 3 \times 3$  kernel with a stride of 1 and appropriate padding to preserve proximity, producing an output tensor with 16 channels. Nonlinearity is then introduced through a ReLU activation, allowing the network to model complex patterns. The subsequent layer also uses a  $3 \times 3 \times 3$  kernel but applies a stride of 2, resulting in neighborhood downsampling and a reduction in feature dimensions. This layer retains 16 channels, and the transformation is again followed by a ReLU activation function to ensure non-linear mapping. The final convolutional layer refines the representation further, using another  $3 \times 3 \times 3$  kernel to reduce the number of channels to 2. A ReLU activation function is applied once more to enhance the model's sensitivity to subtle variations in the input. After the convolutional process, the resulting 3D tensor is flattened into a 1D tensor, enabling further dimensionality reduction. The flattened representation is then passed through three fully connected layers, each employing a ReLU activation function to progressively reduce the dimensionality. The final linear layer outputs a compressed representation of the input in the latent space with two dimensions.

This process maps each three-channel input sequence from a single cycle to a single point in the two-dimensional latent space. When applied across multiple cycles, the encoder generates a collection of such points, forming a distribution in the latent space (Figure 1). This distribution effectively represents the unique characteristics of battery aging over time, providing an interpretable and compact visualization of the aging process.



**Figure 1.** Structure of an autoencoder to map battery's aging state onto the latent space.

### 2.3. Classification

In the latent space, each embedding corresponds to a battery sample at a given stage of degradation. The absolute values of these embeddings are not directly interpretable, instead, meaningful information emerges from how the points are arranged relative to one another. Variations in distance and clustering patterns reflect differences in aging behavior, enabling the identification of distinct health conditions. These spatial relationships provide the basis for classifying the SOH of battery cells.

To achieve classification, a Support Vector Machine (SVM) is employed. SVM is a supervised machine learning method that identifies an optimal boundary that maximizes the margin between classes while limiting mis-classification [15]. SVM was chosen for this study because it is well-suited for our example of a limited experimental dataset, as it relies on a subset of critical samples, referred to as support vectors, located near the decision boundary, enabling it to make robust decisions even with small datasets. Its convex optimization framework also guarantees a global optimum, making it insensitive to initial starting points. For larger datasets with a greater number of anchor points, more advanced classifiers or clustering algorithms might be considered.

As a non-parametric model, SVM entails only few explicit assumptions about the underlying data distribution or feature structure compared to other models. Rather than enforcing predefined cluster shapes, it determines a decision boundary directly from the data, allowing for more flexible class separation. This property is particularly advantageous for analyzing aging trajectories in the latent space, where the boundaries between different states may be irregular. Consequently, the method partitions the latent space into regions corresponding to different aging stages, with the capability to accommodate additional classes if required. By observing the distribution of these regions, the progression of cell aging can be understood comprehensively.

SVM classifiers with linear and nonlinear kernels were employed to classify latent representations of battery cells into three aging categories: fresh, aged, and damaged cells. The regularization parameter  $C$  and the kernel parameter  $\gamma$  were tuned empirically to balance classification accuracy and generalization on the limited experimental dataset. This tuning was necessary to prevent overly complex decision boundaries that could overfit noise in the latent space while still capturing the underlying aging trajectories.

### 2.4. Cross-Validation

Cross-validation (CV) was used to select SVM hyperparameters and to assess the generalization performance of the classifier [16]. Specifically, CV was applied to evaluate various combinations of the regularization parameter  $C$  and  $\gamma$  for both linear and nonlinear kernels.

A 3-fold CV scheme was adopted, in which the labeled latent-space samples were randomly partitioned into three non-overlapping subsets. In each fold, two subsets were used for training and the remaining subset for validation. The classification accuracy was averaged across all folds to obtain the CV-score for a given hyperparameter combination.

Final hyperparameters were selected by jointly considering the training accuracy and CV-score. This ensured that the classifier maintained consistent performance across folds while avoiding overfitting, particularly in regions of the latent space where aging trajectories overlap. The selected model was subsequently used for all latent-space classification results reported in this work.

### 3. Materials and Methods

#### 3.1. Experimental Setup

EIS measurements were conducted during charge–discharge cycling on three LiCoO<sub>2</sub> LIR2032 coin cells representing different aging conditions. The investigated samples correspond to a *fresh cell* (37 mAh), an *aged cell* (25 mAh), and a *damaged cell* (16 mAh). Cell capacities were obtained by discharging fully charged cells at their nominal rate until the end-of-discharge voltage of approximately 3.06 V was reached. A total of 18 EIS measurements are performed per SOC point during one cycle using a Biologic SP-200 potentiostat, ensuring minimal SOC variation between measurements. All experiments were carried out at room temperature under uncontrolled ambient conditions. The coin cells are charged under constant current–constant voltage (CC-CV) conditions. The frequency range for galvanostatic EIS measurements spans from 10 mHz to 300 kHz. The applied current amplitude was automatically adjusted to ensure a minimum voltage perturbation of 10 mV. Each impedance spectrum was obtained by averaging over two cycles and consisted of 84 logarithmically spaced frequency points. A resolution of 12 points per decade was used across most of the frequency range, while the lowest-frequency region was sampled more sparsely at 6 points per decade to reduce measurement time. This adjustment minimizes the impact of low-frequency measurement on the total experiment duration.

#### 3.2. Data Processing

The EIS data obtained are parameterized by ECM combined with Distribution of Relaxation Times (DRT) techniques to create virtual models representing the three real battery cells [17]. This process is detailed in [18], which provides a systematic methodology for generating large-scale synthetic datasets from limited experimental measurements. In this approach, ECM parameters are first identified from the measured impedance data. Their dependence on SOC is then approximated using Chebyshev polynomials, yielding a compact functional representation of parameter evolution over a single charging–discharging cycle. The resulting polynomial coefficients serve as a descriptor of cell behavior. To extend the dataset, Quasi-Monte Carlo (QMC) sampling is applied to generate a wide range of coefficient combinations around the experimentally derived states, effectively spanning different regions of battery health. These parameter sets are subsequently used to simulate battery responses under ideal sinusoidal current excitation, producing corresponding voltage signals for training. In this way, the method establishes a clear relationship between the collected EIS data and the simulated time-series data, effectively expanding a small, experimentally constrained dataset into a rich, numerical physics-based data set for robust model training.

The simulated ideal sinusoidal voltage and current data span a frequency range that matched the frequency components used in the original EIS measurements. This alignment ensures comparability between simulated and measured data. However, in practical applications, battery behavior is influenced by factors, including temperature fluctuations [19], electrode degradation [20] and other environmental conditions, rendering ideal sinusoidal waveforms unrealistic. To bridge this gap, Gaussian noise [21] is introduced to the sinusoidal voltage data across all frequencies. While it is acknowledged that Gaussian noise may not fully capture the complete complexity of all real battery operating conditions, it serves as a reasonable and widely accepted approximation to simulate the variability found in these environments. This assumption allows for the systematic testing of the model's robustness against common forms of signal perturbations.

To evaluate the model's generalization capability with limited input data, an investigation is conducted into the number of individual oscillations from the time-series data that are required to reproduce the same latent space topology as full periodic signals. In this context, independently permuting the signal segments without preserving frequency order

is introduced as a strategy to reduce data requirements. It allowed for the exploration of the model's robustness and adaptability with smaller datasets.

In the experimental EIS protocol, impedance spectra are measured at 76 discrete excitation frequencies. To construct permuted time-domain signals with comparable spectral content, the simulated sinusoidal voltage and current response at each individual frequency was first segmented into 76 non-overlapping subsets of equal duration. Each subset corresponds to one oscillation period, such that the total number of subsets matches the number of EIS frequency points. For a single EIS measurement, this procedure results in a set of 76 sinusoidal time-series signals (one per frequency), each divided into 76 subsets. A new permuted signal is then generated by randomly selecting one subset from each frequency-specific signal and concatenating these subsets in the time domain. This process ensures that each reconstructed signal contains exactly one segment from every excitation frequency, thereby preserving full frequency coverage while removing the original monotonic frequency ordering. By repeating this permutation process for all 76 possible subset indices, a total of  $76 \times 76$  scrambled time-domain signals are generated from a single EIS measurement. These permuted signals exhibit pseudo-random temporal structure while maintaining the overall spectral characteristics of the original EIS data. The resulting dataset is used to evaluate the robustness of the latent-space representation and classification performance under non-ideal, mixed-frequency excitation conditions.

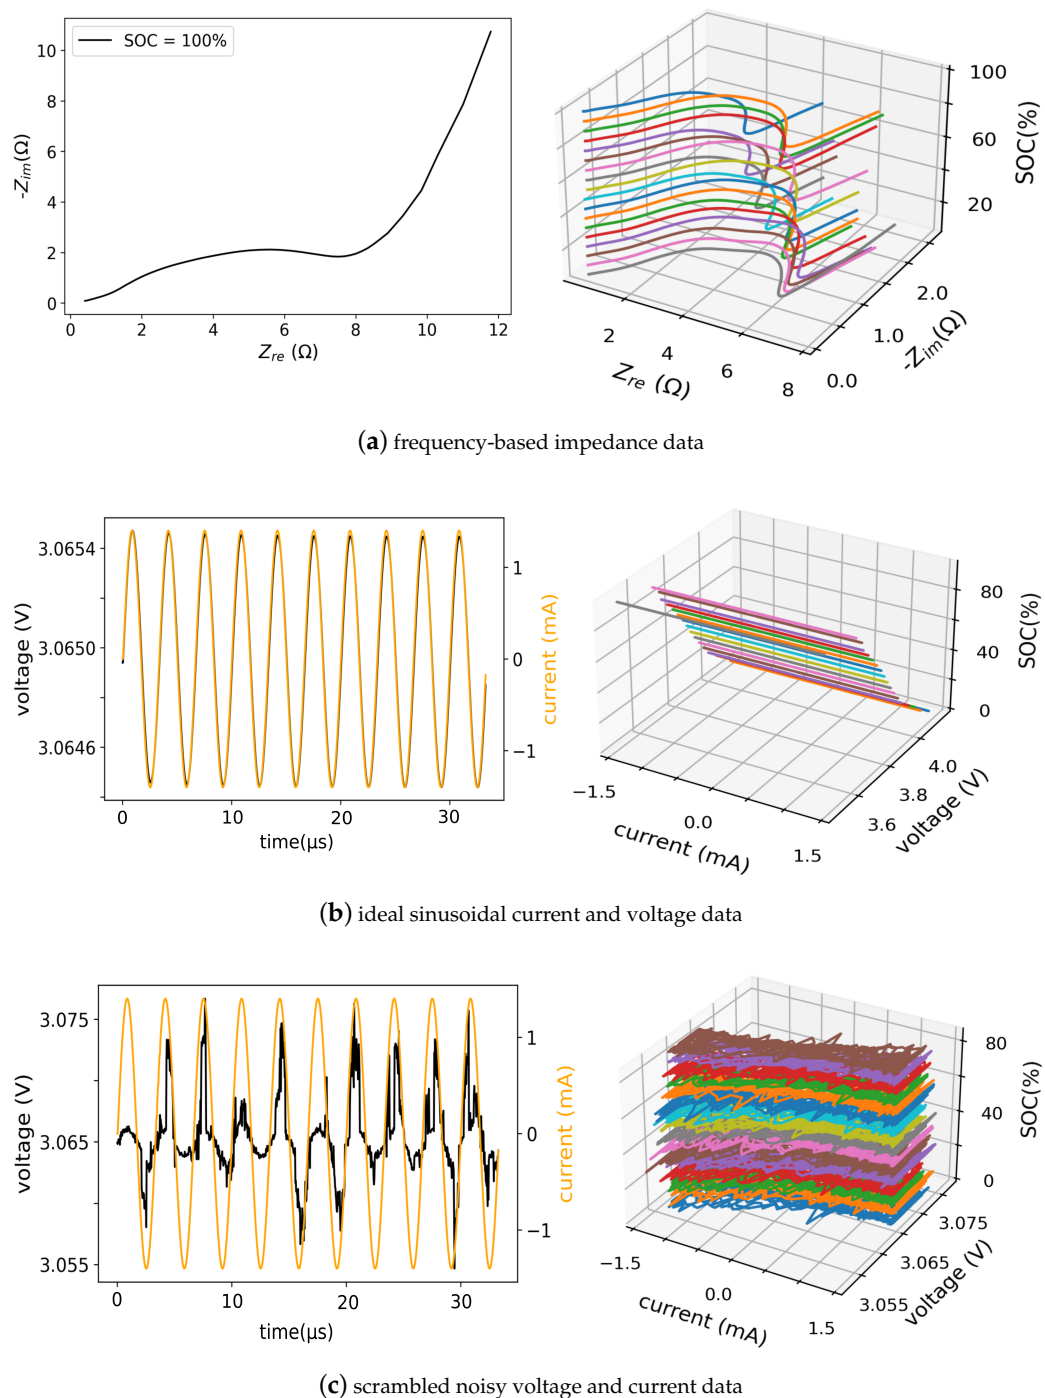
## 4. Results and Discussion

### 4.1. Latent Representation: Time-Domain Data vs. Frequency-Domain Data

The generated synthetic datasets are used to train the autoencoder, and its resulting representations are analyzed in a two-dimensional latent space. Although higher-dimensional embeddings may be advantageous for more complex system behavior, a two-dimensional representation is sufficient here and offers clear interpretability for tracking battery aging. In this space, each point corresponds to a battery instance at a specific cycle, thereby reflecting its SOH.

We make use of three different data set types: traditional, frequency-based EIS data (Figure 2a), time-series data from sinusoidal loads with systematically varied frequency (Figure 2b), and "scrambled" data-sets with randomly permuted sub-sequences from the the time-series data (Figure 2c), which allows us to assess the robustness of our approach to stochastic excitation schemes. Impedance-based analysis remains a standard approach in battery diagnostics and serves as a reliable reference for evaluating alternative methods. Comparing it with time-domain analysis helps assess whether the latter can capture dynamic battery behavior with similar fidelity. Additionally, aligning latent space representations derived from ideal periodic signals with those obtained from impedance data provides a method to verify consistency under linear system assumptions. To extend beyond this regime, latent spaces generated from permuted and noisy signals are also examined. The agreement across these representations indicates that the proposed approach remains valid even under nonlinear and non-stationary conditions.


Figure 3 illustrates the latent space representation obtained from three types of training data. The blue points correspond to synthetic samples generated via QMC to capture variations in aging behavior, while the three star markers denote the experimental reference cells used for validation and comparison.



**Figure 2.** Investigated excitation/response profiles. Left plots: impedance data captured at a single SOC = 100%, and time-dependent data at a single frequency and a specific SOC = 100%. Right plots: Impedance data collected over a full cycle and current-to-voltage relationship across multiple SOCs at a single frequency (reproduced from own work under permissive licensing conditions [22]).

The latent space topologies obtained from all three datasets exhibit a similar overall organization. The close agreement between Figure 3a and Figure 3b suggests that time-domain signals under ideal conditions capture key features comparable to those derived from frequency-domain impedance analysis within the linear regime. Furthermore, the resemblance of Figure 3c to the others indicates that frequency-independent noisy signals can also yield consistent results, enabling reliable dynamic battery performance analysis with impedance-based methods serving as a useful reference.

Notably, the latent representation derived from time-domain data appears rotated relative to that derived from impedance data. This rotation is not problematic, as such transformations are typical in unsupervised dimensionality reduction, where the orientation of the latent space is not uniquely defined. The nonlinear mapping performed by the autoencoder can lead to rotations or rearrangements of the embedding while preserving the underlying relationships within the data.



**Figure 3.** Latent space representation of generated battery samples trained using (a) impedance data, (b) ideal periodic time-series data, and (c) permuted time-series data. The star markers denote experimental reference cells: red stars correspond to damaged cells (16 mAh), green stars to aged cells (25 mAh), and blue stars to fresh cells (37 mAh). (reproduced from own work under permissive licensing conditions [22]).

The placement of the experimental samples within the latent space clearly reflects the progression of battery aging. The transition from blue (*fresh cell*), through green (*aged cell*), to red (*damaged cell*) illustrates the evolution from an initial state to end-of-life. This aging trajectory remains consistently embedded within the synthetic SOH distribution across all three training datasets (Figure 3), indicating stable alignment between experimental reference points and generated data. Even with a small number of experimental samples, the relationship between SOH and the latent space topology is clearly captured. This result highlights the data efficiency of the autoencoder in combination with a data augmentation strategy, demonstrating that essential features and patterns of battery aging can be captured with a very sparse dataset and physics-based synthetic data. The encoding of aging patterns within the latent space establishes a strong link between the observed degradation trajectories and the learned representation, demonstrating the autoencoder's ability to extract meaningful insights from limited data.

#### 4.2. Classification: Permuted Noisy Data vs. Ideal Periodic Data

Due to the observed variations in the mapping of input data to points in the LS by the autoencoder, a robust scheme needs to be devised to analyze the relevant LS topology. As an example, a straightforward task is to classify all LS points as belonging to one of the categories *fresh*, *aged*, and *damaged*. Battery health classification in the latent space is performed by SVM with both linear and RBF kernels. Achieving reliable performance requires appropriate selection of the kernel type and careful tuning of the associated hyperparameters.

The motivation for introducing permuted time-series data is the reduction of required input data and, consequently, measurement time. In contrast to full periodic signals, which require continuous excitation over complete cycles at each frequency, the permutation strategy relaxes the need for strict temporal ordering. This enables the autoencoder to extract informative latent representations from shorter or partially observed signals while preserving the essential frequency content. As a result, fewer time-domain samples are sufficient, which is particularly advantageous for practical battery monitoring where long measurement sequences are undesirable.

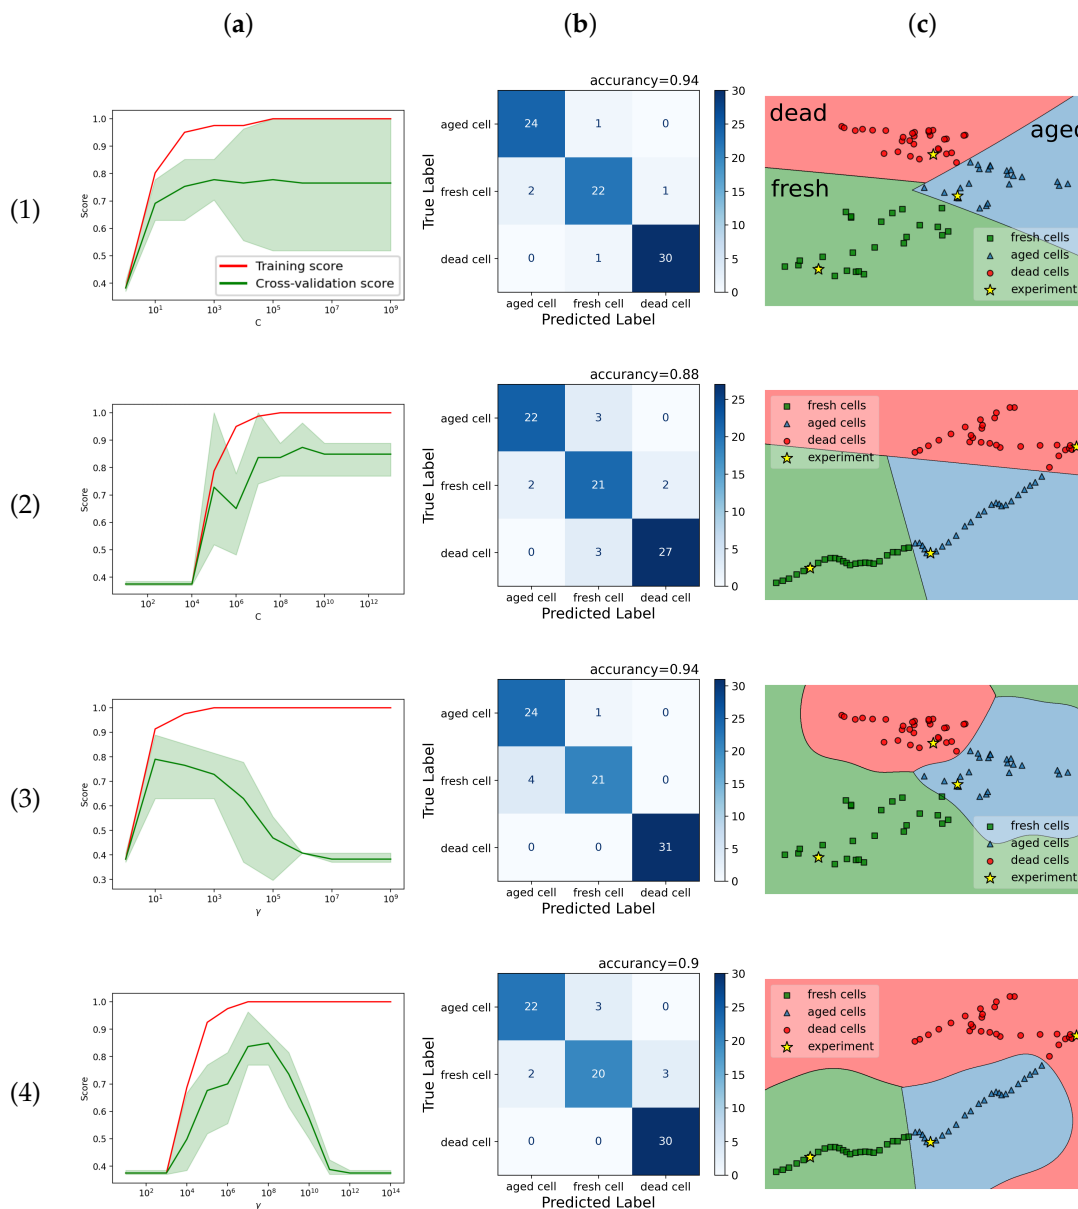
For the linear kernel, the validation curves in Figure 4 column a, rows 1 and 2, illustrate the dependence of training scores and cross-validation scores on the regularization parameter  $C$  for periodic time-series and permuted noisy data, respectively. As  $C$  increases, both scores improve before reaching a plateau. This stabilization indicates a suitable trade-off between model bias and variance, beyond which additional model complexity does not provide further gains and may lead to overfitting. Based on this trend,  $C$  values of  $10^6$  and  $10^7$  are selected for the linear kernel.

Beyond hyperparameter selection, a clear structural difference between the two datasets is evident in the latent space. For ideal periodic inputs (row 1), the LS points are distributed in a more diffuse, cloud-like pattern, resulting in lower and more variable cross-validation scores. In contrast, the permuted noisy inputs (row 2) yield LS points forming a clearer, trajectory-like structure, leading to higher and more stable cross-validation scores. One possible explanation is that permuting the temporal sequence reduces the influence of local temporal correlations that are not directly related to battery aging, such as periodic excitation patterns or measurement artifacts. By removing the chronological order, the autoencoder is forced to rely on intrinsic, time-invariant features of the individual responses, which are more closely linked to the battery's health state. As a result, the latent space becomes more structured and aligns more clearly with the underlying aging trajectory. These results indicate that permuted data allows for the training of a higher-quality model, as the latent representations provide a more structured mapping of battery aging progression. The permuted dataset considered here corresponds to the case in which all time-series samples are fully scrambled, rather than truncated into shorter contiguous segments. This choice represents the most conservative scenario in terms of temporal information loss and demonstrates that meaningful latent-space structure can still be learned without relying on complete periodic waveforms. The improved trajectory-like organization observed in the latent space indicates that permutation acts as a regularization mechanism, encouraging the autoencoder to focus on aging-relevant features rather than precise temporal phase information.

A similar trend is observed for the RBF kernel. The validation curves for the RBF kernel, shown in Figure 4 column a, rows 3 and 4, examine the influence of the coefficient  $\gamma$ . For very small  $\gamma$  values, both training and validation scores remain low, indicating underfitting. As  $\gamma$  increases to moderate levels, both scores improve, reflecting effective classifier performance. However, excessively large  $\gamma$  values lead to overfitting, where training performance remains high while validation performance deteriorates. Based on these observations,  $\gamma$  values of  $10^2$  and  $10^8$  are selected for the ideal and permuted datasets, respectively. This parameter selection ensures a well-balanced model capable of accurately distinguishing battery aging states in the latent space.

A clear difference is observed in the selected  $\gamma$  values for the two datasets, whereas the optimal  $C$  values remain comparable. The discrepancy is caused by scaling differences in the latent representations derived from the respective inputs. While linear decision boundaries are invariant to such scaling, nonlinear boundaries are sensitive to these variations.

The classifier results shown in Figure 4 column c, demonstrate the decision boundaries in the 2D latent space. These boundaries demonstrate that the extracted features are sufficiently discriminative to separate *fresh*, *aged*, and *damaged* battery states. Notably, similar classification performance is obtained for latent spaces derived from both ideal time-series data and permuted noisy inputs. In both cases, the aging progression follows a consistent trajectory, transitioning from the bottom-left blue region (*fresh cell*), through the center-right orange region (*aged cells*), and eventually to the upper-center green region (*damaged cells*). This agreement indicates that reducing the frequency content, as applied in stochastic pulse signal analysis, may not degrade classification performance.



**Figure 4.** Comparison of SVM-based battery health classification in the two-dimensional latent space obtained from different autoencoder training strategies. Rows correspond to: (1) linear-kernel SVM with ideal periodic data; (2) linear-kernel SVM with permuted noisy data; (3) RBF-kernel SVM with ideal periodic data; (4) RBF-kernel SVM with permuted noisy data. Columns correspond to: (a) validation curves showing training and cross-validation accuracy as a function of the SVM hyperparameters ( $C$  for linear kernel,  $\gamma$  for RBF kernel); (b) confusion matrices evaluated on an independent test dataset; (c) decision boundaries and class distributions in latent space. Compared to ideal periodic data, permuted noisy training produces a more structured, trajectory-like latent-space geometry that improves cross-validation stability and classifier robustness, particularly for linear decision boundaries. (reproduced from own work under permissive licensing conditions [22]).

To assess the accuracy of the trained classifiers, a set of independent simulation-based battery datasets, corresponding to *fresh*, *aged*, and *damaged* cells and not used during training, is employed for testing. These test datasets are mapped into the latent space using the pretrained autoencoder, without retraining, and subsequently fed into the already trained SVM classifiers. While the true health-state labels of the test data are known to the user, they are not provided to the classifier during inference. The classifier assigns each latent-space point to one of the predefined categories, and the predicted labels are then

compared against the known ground truth to evaluate classification performance. The resulting confusion matrix is displayed in Figure 4 column b. The diagonal elements of the matrix represent the number of correctly classified data points, showing the highest values. Misclassified samples are reflected in the off-diagonal entries, which reach a maximum of 4, suggesting a low error level. Notably, these misclassified points are located close to the decision boundaries between neighboring SVM domains, suggesting that classification errors occur primarily for latent representations near transitions between battery health states. A higher value of diagonal elements reflects accurate predictions by the classifier. Accuracy, as a measure of overall classification performance, indicates that the RBF kernel achieves slightly better results than the linear kernel. However, the difference between the two approaches is minor, demonstrating that both classifiers perform effectively in distinguishing battery aging states.

To provide a more comprehensive view of the classifier's performance, Table 1 summarizes the key classification metrics.

**Table 1.** Classification metrics for SVM classifiers.

Data Type	Kernel	Accuracy	Precision	Recall	F1-Score
Ideal Periodic	Linear	0.95	0.94	0.95	0.95
	RBF	0.96	0.96	0.96	0.96
Permuted Noisy	Linear	0.95	0.94	0.95	0.94
	RBF	0.96	0.96	0.96	0.96

The autoencoder architecture employed in this work is intentionally kept simple. While more complex networks could be used, the emphasis here is on a workflow that enables low-cost deployment during battery operation. Once trained offline for a given class of cells, the autoencoder is used only for inference, mapping new measurements into the latent space with minimal computational overhead. This inference step can be executed efficiently on commodity hardware, making the approach suitable for embedded or in-field applications. The evaluation on independent test datasets serves to validate this deployment strategy. These datasets are not used during training and are processed using the pretrained autoencoder without retraining. The resulting latent representations are classified by the SVM to assess generalization performance. The strong agreement between predicted and true health states confirms that the learned latent-space structure is robust and transferable, supporting the feasibility of the proposed low-cost, inference-focused battery health monitoring framework.

## 5. Conclusions

This paper presents a framework for classifying battery health states based on a limited set of experimental data. The approach utilizes a two-dimensional latent space representation obtained from an autoencoder to both distinguish between these states and analyze the progression of degradation. By incorporating ideal periodic inputs alongside permuted, noise-affected time-domain signals, the method successfully transitioned from linear to non-linear regimes, reflecting better practical operating conditions. Traditional impedance analysis served as a benchmark to validate the robustness and efficiency of the proposed methodology. The results demonstrate a stable aging trajectory in the latent space, with the distribution of battery states remaining consistent across both time-domain and frequency-domain inputs, as well as under noisy conditions. This consistency underscores the potential of the classifier for diverse data inputs, reinforcing its reliability. Importantly, this study highlights the opportunity to develop stochastic pulse sequences that optimize

charging strategies, potentially extending battery lifetimes by leveraging these insights into battery aging behavior.

Despite the promising outcomes, the current methods are not without limitations, which could influence their efficacy and broader applicability. These limitations also provide clear directions for future research to further validate and extend our approach. One notable limitation is the reliance on a relatively small dataset of only three button cells (fresh, aged, and damaged). While this dataset was sufficient for the initial development and proof-of-concept, it restricts the generalizability of the results. To address this, our future work will include an extensive comparison between our numerical models and a more extensive experimental dataset on battery aging, including different battery types and various operating conditions. This will not only validate our existing models but also refine them, enhancing their robustness and adaptability.

The use of small-capacity coin cells in this study also raises questions about the model's applicability to high-capacity cells, such as those used in electric vehicles (EVs). Our methodology, which focuses on extracting latent space features from time-series data, is a direct response to the challenges of performing impedance measurements on large-format batteries. Since our method is independent of cell geometry and capacity, we believe it can be scaled and validated for EV battery diagnostics, provided the necessary time-series data is available. Future work will focus on validating the methodology using lithium nickel manganese cobalt oxide (NMC) cylindrical and LFP prismatic cells, thereby assessing its robustness across different chemistries and formats and supporting its applicability for real-world EV battery diagnostics and large-scale deployment. The inclusion of LFP systems is particularly important due to their flat voltage profiles, which pose additional challenges for conventional diagnostics.

Another limitation pertains to the exclusion of temperature variations, which play a pivotal role in battery performance and aging. In this study, only SOC and SOH were considered, under the assumption of a controlled room temperature environment. Our future work is planned to include temperature as an additional parameter in both equivalent models and neural networks. This integration would provide a more complete description of battery dynamics and enhance accuracy under realistic operating conditions.

**Author Contributions:** L.J.: Writing—original draft, Methodology, Software, Visualization, Validation, Resources, Formal analysis, Data curation. F.P.B.: Investigation, Resources. R.-A.E.: Project administration, Funding acquisition. J.G.: Supervision, Writing—review and editing, Funding acquisition. C.S.: Supervision, Writing—review and editing, Conceptualization, Funding acquisition. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research has been financially supported by the Helmholtz AI Cooperation Unit (HAICU), project “Intelligent, individual battery management using spectroscopy and machine learning” (i2Batman). Open Access funding provided by the Max Planck Society.

**Data Availability Statement:** The data presented in this study are available on request from the corresponding author due to IP protection restrictions.

**Conflicts of Interest:** Authors Limei Jin, Franz Bereck, Rüdiger-A. Eichel, Josef Granwehr and Christoph Scheurer were employed by the Institute of Energy Technologies, Fundamental Electrochemistry (IET-1), Forschungszentrum Jülich. The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as potential conflicts of interest.

## References

1. Tian, Y.; Zeng, G.; Rutt, A.; Shi, T.; Kim, H.; Wang, J.; Koettgen, J.; Sun, Y.; Ouyang, B.; Chen, T.; et al. Promises and Challenges of Next-Generation “Beyond Li-ion” Batteries for Electric Vehicles and Grid Decarbonization. *Chem. Rev.* **2021**, *121*, 1623–1669. [[CrossRef](#)] [[PubMed](#)]
2. Fuller, T.F.; Doyle, M.; Newman, J. Simulation and Optimization of the Dual Lithium Ion Insertion Cell. *J. Electrochem. Soc.* **1994**, *141*, 1. [[CrossRef](#)]
3. Doyle, M.; Fuller, T.F.; Newman, J. Modeling of Galvanostatic Charge and Discharge of the Lithium/Polymer/Insertion Cell. *J. Electrochem. Soc.* **1993**, *140*, 1526. [[CrossRef](#)]
4. Carelli, S.; Bessler, W.G. Coupling Lithium Plating with SEI Formation in a Pseudo-3D Model: A Comprehensive Approach to Describe Aging in Lithium-Ion Cells. *J. Electrochem. Soc.* **2022**, *169*, 050539. [[CrossRef](#)]
5. Qu, W.; Deng, H.; Pang, Y.; Li, Z. An Improved Gaussian Process Regression Based Aging Prediction Method for Lithium-Ion Battery. *World Electr. Veh. J.* **2023**, *14*, 153. [[CrossRef](#)]
6. Tran, M.K.; Mathew, M.; Janhunen, S.; Panchal, S.; Raahemifar, K.; Fraser, R.; Fowler, M. A comprehensive equivalent circuit model for lithium-ion batteries, incorporating the effects of state of health, state of charge, and temperature on model parameters. *J. Energy Storage* **2021**, *43*, 103252. [[CrossRef](#)]
7. Halleman, N.; Howey, D.; Battistel, A.; Saniee, N.F.; Scarpioni, F.; Wouters, B.; La Mantia, F.; Hubin, A.; Widanage, W.D.; Lataire, J. Electrochemical impedance spectroscopy beyond linearity and stationarity—A critical review. *Electrochim. Acta* **2023**, *466*, 142939. [[CrossRef](#)]
8. Itagki, M.; Honda, K.; Hoshi, Y.; Shitanda, I. In-situ EIS to determine impedance spectra of lithium-ion rechargeable batteries during charge and discharge cycle. *J. Electroanal. Chem.* **2015**, *737*, 78–84. [[CrossRef](#)]
9. Krewer, U.; Röder, F.; Harinath, E.; Braatz, R.D.; Bedürftig, B.; Findeisen, R. Review—Dynamic Models of Li-Ion Batteries for Diagnosis and Operation: A Review and Perspective. *J. Electrochem. Soc.* **2018**, *165*, A3656. [[CrossRef](#)]
10. Kramer, M.A. Nonlinear principal component analysis using autoassociative neural networks. *AIChE J.* **1991**, *37*, 233–243. [[CrossRef](#)]
11. Che, Y.; Zheng, Y.; Sui, X.; Teodorescu, R. Boosting battery state of health estimation based on self-supervised learning. *J. Energy Chem.* **2023**, *84*, 335–346. [[CrossRef](#)]
12. Ide, H.; Kurita, T. Improvement of learning for CNN with ReLU activation by sparse regularization. In Proceedings of the 2017 International Joint Conference on Neural Networks (IJCNN), Anchorage, AK, USA, 14–19 May 2017; pp. 2684–2691. [[CrossRef](#)]
13. Tang, P.; Qiu, Z.; Yao, Z.; Pan, J.; Cheng, D.; Gu, X.; Sun, C. Lithium-ion battery RUL prediction based on optimized VMD-SSA-PatchTST algorithm. *Sci. Rep.* **2025**, *15*, 26824. [[CrossRef](#)] [[PubMed](#)]
14. Li, W.; Liu, Z. A method of SVM with Normalization in Intrusion Detection. *Procedia Environ. Sci.* **2011**, *11*, 256–262. [[CrossRef](#)]
15. Corinna, C.; Vladimir, V. Support-vector networks. *Mach. Learn.* **1995**, *20*, 273–297. [[CrossRef](#)] [[PubMed](#)]
16. Asrol, M.; Papilo, P.; Gunawan, F.E. Support Vector Machine with K-fold Validation to Improve the Industry’s Sustainability Performance Classification. *Procedia Comput. Sci.* **2021**, *179*, 854–862. [[CrossRef](#)]
17. Bartsch, C. Scalable Data Aggregation and Interpretation of Electrochemical Impedance Spectroscopy for Battery Characterization. Master’s Thesis, RWTH Aachen University, Aachen, Germany, 2022.
18. Jin, L.; Bereck, F.P.; Granwehr, J.; Scheurer, C. Extending Equivalent Circuit Models for State of Charge and Lifetime Estimation. *Electrochem. Sci. Adv.* **2025**, *5*, e202400024. [[CrossRef](#)]
19. Vashisht, S.; Rakshit, D.; Panchal, S.; Fowler, M.; Fraser, R. Thermal behaviour of Li-ion battery: An improved electrothermal model considering the effects of depth of discharge and temperature. *J. Energy Storage* **2023**, *70*, 107797. [[CrossRef](#)]
20. Edge, J.S.; O’Kane, S.; Prosser, R.; Kirkaldy, N.D.; Patel, A.N.; Hales, A.; Ghosh, A.; Ai, W.; Chen, J.; Yang, J.; et al. Lithium ion battery degradation: What you need to know. *Phys. Chem. Chem. Phys.* **2021**, *23*, 8200–8221. [[CrossRef](#)] [[PubMed](#)]
21. Moradpour, A.; Kasper, M.; Hoffmann, J.; Kienberger, F. Measurement Uncertainty in Battery Electrochemical Impedance Spectroscopy. *IEEE Trans. Instrum. Meas.* **2022**, *71*, 1006209. [[CrossRef](#)]
22. Jin, L. AI-Based Simulation of Battery System Combined with Advanced Spectroscopy. Ph.D. Thesis, RWTH Aachen University, Aachen, Germany, 2024.

**Disclaimer/Publisher’s Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.