This article has been accepted for publication in IEEE Transactions on Transportation Electrification. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTE.2026.3683190

A GPT-Powered Automated Feature Extraction Framework for State of Health Estimation of Fast-Charging Batteries

Laijin Luo, Yu Wang, *Senior Member*, *IEEE*, Yang Li, *Senior Member*, *IEEE,* Changfu Zou, *Senior Member*, *IEEE*

***Abstract*—Efficient and reliable feature extraction engineering plays a crucial role in improving the accuracy and real-time operational capability for data-driven state-of-health (SOH) estimation of lithium-ion batteries. However, conventional feature extraction methods are time-intensive and prone to manual bias, particularly under the multi-step constant current fast-charging protocols prevalent in practical scenarios. To address this problem, this study introduces a generative pre-trained transformer (GPT)-powered automated feature extraction (AFE) method that identifies a series of health features from voltage and capacity aging data in both the time domain and frequency domain. Observation and analysis reveal that with increasing temperature, capacity-related features exhibit a more stable correlation with the SOH compared to voltage-related features. Based on this insight, we propose a two-stage kernel extreme learning machine-autoencoder (TS-KELM-AE) model, which integrates deep feature extraction and nonlinear mapping capabilities to estimate SOH under varying temperatures. Compared to manual feature extraction (MFE) and mainstream deep learning as a feature extractor, the proposed AFE method is feasible and efficient, and the TS-KELM-AE model achieves higher accuracy. Furthermore, analyses of feature importance, model multicollinearity, and cross-battery generalization under diverse operating conditions demonstrate that integrating AFE with the TS-KELM-AE framework establishes a robust and scalable solution for practical applications.**

***Index Terms*—Lithium-ion batteries, state of health (SOH), automated feature extraction (AFE), pre-trained transformer (GPT), data-driven, two-stage kernel extreme learning machine-autoencoder (TS-KELM-AE).**

I. Introduction

THE green energy transition increasingly relies on rechargeable lithium-ion batteries (LIBs), valued for their high energy density, durability, recyclability, and reduced environmental impact [1]. With the rapid proliferation of electric vehicles (EVs), energy storage systems, and consumer electronic devices, traditional slow-charging technologies have increasingly struggled to keep pace with users' expectations for shorter charging times and improved overall efficiency [2]. Consequently, the consumer demand for faster charging capabilities is also growing. Fast charging technology has become a critical technology to resolve this challenge, with its core advantages of higher charging rates and enhancing user experience [3]. For EVs, fast charging technology typically enables the boost of the LIB state of charge (SOC) up to 80% in a short period of time, significantly reducing user wait times [4].

Unfortunately, the widespread application of fast charging technology has brought about a new set of technical challenges [5]. Under high-rate charging conditions, the chemical and physical reactions within LIBs are significantly accelerated [6]. This fact leads to an increased pace of battery degradation, which directly impacts its service life and safety [7]. As a result, battery management systems (BMSs) face greater demands to monitor and predict the battery state of health (SOH) in real-time to prevent overcharging and prolong battery lifetime. Exploring and enhancing the real-time operational capability and precision of SOH estimation is thus a critical task to ensure battery performance and safety.

Generally, SOH estimation methods can be categorized into model-based and data-driven methods [8]. The model-based methods typically rely on constructing an empirical or physics-based model and identifying its key parameters, such as those based on electrochemical models (EMs) [9] and equivalent circuit models (ECMs) [10]. These models are then coupled with filtering algorithms, including the Kalman filter (KF) [11], the particle filter [12], and their variants [13], to evaluate SOH. Increasing these models' complexity can improve estimation precision, but it also poses more challenges to the parameter identification process. Furthermore, under real-world operating conditions involving extreme charging/discharging rates and varying temperature environments, the accuracy of these model-based methods may be limited, which significantly reduces their suitability for practical applications.

In comparison, data-driven approaches eliminate the reliance on empirical and physical modeling by employing machine learning techniques to directly establish the relationship between SOH and crucial health features from historical data. Under variable operating conditions, data-driven methods tend to yield greater SOH estimation precision. Such an approach is primarily composed of two core components: health feature extraction and estimation model construction [14].

As a foundation of the data-driven approaches, the quality

This work was supported by the National Key R&D Program of China (Sci-Tech Innovation 2030 Agenda) (Grant No. 2025ZD0805405). (Corresponding author: Yu Wang)

Laijin Luo and Yu Wang are with the School of Electrical Engineering, Chongqing University, Chongqing, China (e-mail: 20241101048@stu.cqu.edu.cn; yu_wang@cqu.edu.cn).

Yang Li is with the School of Electrical Engineering and Automation, Wuhan University, Wuhan 430072, China (e-mail: yangli@ieee.org)

Changfu Zou is with the Department of Electrical Engineering, Chalmers University of Technology, 41296 Gothenburg, Sweden (e-mail: changfu.zou@chalmers.se).

of health feature extraction directly impacts the performance of the subsequent estimation models and the reliability of the results. Existing methods for health feature extraction can broadly be classified into manual feature extraction (MFE) [15], [16], [17] and automated feature extraction (AFE) [18], [19]. Traditional MFE methods for analyzing and extracting battery degradation information commonly depend on expert knowledge and experience-driven interpretation of massive volumes of time-series data, typically involving analyzing the inflection point, slope, area, ratio, and internal resistance trends in battery charge/discharge aging curves [15], [16], [17]. Although these methods are interpretable and physically meaningful, MFE approaches inherently require extensive manual intervention and expert judgement. Such dependence leads to limited scalability, high computational and labor costs, and inconsistent feature quality across datasets.

To address these limitations, AFE methods, such as autoencoder (AE) [18] and convolutional neural network (CNN) [19], have seen increasing adoption for their advantages in flexibility and efficiency. More recently, the emergence of large language models (LLMs), in particular the generative pre-trained transformer (GPT), has opened new opportunities for AFE [20]. As an advanced natural language processing and data analysis model, GPT is known for its powerful data comprehension capabilities, which can be leveraged for pattern recognition and feature extraction from time-series data [21]. This makes it suitable for complex tasks such as battery degradation analysis and prediction.

Although benefiting from GPT can easily extract massive features for SOH estimation, the input of these features is usually associated with high computational complexity, presenting new challenges to the real-time performance, efficiency, and accuracy of BMS in fast charging scenarios. Furthermore, prevailing data-driven algorithms such as extreme learning machine (ELM) [22], support vector machine (SVM) [23], Gaussian process regression (GPR) [24], long short-term memory (LSTM) [25], and Transformer [26] each have inherent limitations. Shallow models like ELM lack the capacity to model complex nonlinear relationships in high-dimensional data. Classical machine learning methods, such as SVM and GPR, require MFE and struggle to uncover deep patterns. Deep learning models like LSTM and Transformer have powerful representational ability, but their high dependence on large-scale data and computational resources poses challenges in many practical battery applications.

To summarize, there remains a critical need for SOH estimation frameworks that 1) can efficiently extract meaningful features from large-scale battery data using modern LLMs like GPT, and 2) enable fast, accurate, and computationally efficient prediction mechanisms suitable for real-time deployment under fast charging conditions. To address these challenges, we propose a novel SOH estimation framework that integrates GPT-powered AFE with a two-stage kernel extreme learning machine-autoencoder (TS-KELM-AE) model. This framework facilitates highly efficient battery health feature extraction and adapts to the demand of real-time SOH estimation under large-scale data input. The main contributions are reflected in the following:

1) A novel SOH estimation framework is proposed that integrates LLMs to enable intelligent, prompt-driven, and user-friendly feature extraction. Task-specific prompts effectively combine battery aging data with external knowledge to automatically generate diverse and meaningful aging features, significantly improving the efficiency and flexibility of feature extraction engineering.

2) We proposed a two-stage kernel ELM-autoencoder (TS-KELM-AE) model to accommodate SOH estimation under massive GPT-generated feature input. The first stage is dedicated to extracting deep features, while the second stage focuses on making decisions for accurate SOH estimation.

3) We systematically reveal the importance of GTP-powered aging features on SOH estimation performance by integrating the SHAP model. The results demonstrate that frequency-domain feature families consistently dominate the attribution rankings across different temperatures.

The remainder of this paper is organized as follows: Section II describes two battery datasets used in this study, along with the fast-charging protocols. Section III presents the extraction of GPT-powered capacity and voltage-related features and introduces the TS-KELM-AE model for SOH estimation. Section IV comprehensively evaluates feature extraction and input selection, further validating feature importance, model multicollinearity, and cross-battery generalization under different conditions. Finally, Section V concludes the paper and outlines future research directions.

II. Battery Dataset Description

In battery aging research, SOH serves as a key metric for evaluating battery degradation. Among various definitions, the maximum available capacity-based SOH has gained widespread adoption due to its interpretability and suitability across different battery types and aging scenarios [27]. The SOH is typically expressed as a percentage:

$$\mathrm{SOH}_i = \frac{Q_i}{Q_{\mathrm{rated}}} \times 100\% \tag{1}$$

where Q_i and Q_{rated} represent the maximum available capacity at the ith cycle and the nominal capacity, respectively.

To comprehensively characterize battery aging behavior under different conditions, in this study, we employed two complementary datasets provided by TBSI-Sunwoda and MIT [28], [29]. These two datasets differ in battery chemistry, temperature, and fast charging protocol, enabling a reliable assessment of model accuracy and cross-battery generalization.

The first dataset comprises 32 lithium-nickel-manganese-cobalt oxide (NCM) batteries with a nominal capacity of 1.1 Ah, tested under four controlled temperatures. Among them, nine batteries were tested at 25 °C and 35 °C, respectively, while seven batteries were tested at 45 °C and 55 °C. The capacity degradation trajectories for each battery are presented in Fig. 1(a), which primarily simulates the impact of fast charging strategies on battery degradation across a wide temperature range.

The second dataset was used to further validate the generalization of the proposed method under actual high-rate charging conditions. It contains 124 commercial $LiFePO_4$ (LFP) batteries, each cycled to the end of life under controlled and different fast charging protocols. All batteries were charged using one-step or two-step constant current (CC) fast charging protocols of the form C1(Q1)–C2, where the current transitions at a prescribed state-of-charge Q1. The second stage terminated at 80% SOC, followed by a standard 1C constant current-constant voltage (CC-CV) procedure with voltage limited to 2.0–3.6 V and finally discharged at 4C. The capacity degradation curves of batteries under various charging strategies are illustrated in Fig. 1(b). Furthermore, Table I summarizes the detailed parameter information for both datasets. Compared to the TBSI-Sunwoda dataset, the MIT dataset is primarily distinguished by its extensive testing of fast charging protocols under constant temperature.

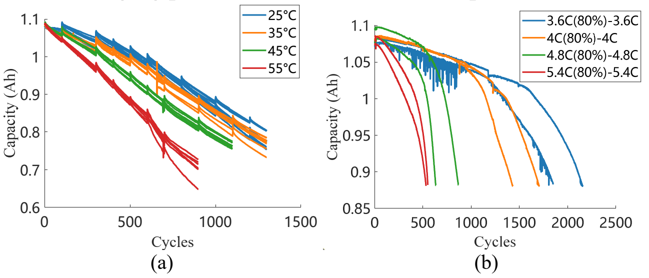

Fig. 1. Battery capacity degradation curves for (a) the TBSI-Sunwoda dataset with NCM batteries and (b) the MIT dataset with LFP batteries.

TABLE I
SPECIFICATIONS OF THE TWO BATTERY DATASETS

Dataset	TBSI-Sunwoda	MIT
Cathode	NCM	LFP
Anode	Graphite	Graphite
Nominal capacity	1.1 Ah	1.1 Ah
Nominal voltage	4.2 V	3.6 V
Temperature	25, 35, 45, 55 °C	30 °C
Charging scheme	9-step fast charging	2-step charging
Discharging scheme	Constant current (1C)	Constant current (4C)

As seen in Fig. 1(a), under the same temperature conditions, the degradation trajectories exhibit high similarity during the early aging period, demonstrating high consistency, while obvious deviations gradually appear during the middle and late aging periods. Additionally, under different temperatures, batteries exhibit distinct degradation trajectories, particularly at the initial aging stage. As the temperature increases, the rate of battery capacity degradation to the failure threshold also accelerates significantly, which indicates that high temperatures exacerbate multiple aging mechanisms, thereby further shortening the battery's service life. While Fig. 1(b) reveals the variety of degradation patterns from another perspective. For batteries employing the same fast-charging protocol, initial cycles typically exhibit good trajectory overlap. However, as cycles towards accelerated degradation, individual differences rapidly enlarge. In cross-protocol comparisons, aging trajectories diverge almost from the cycle start, which indicates that fast charging strategies are also key factors influencing SOH degradation patterns.

In this study, a cell is selected from each temperature in the TBSI-Sunwoda datasets for subsequent feature extraction and SOH estimation modeling. Eight representative cells were selected from the MIT datasets to validate robustness and generalization under cross-cell SOH prediction. These cells correspond to four distinct fast charging protocols, each with two parallel test cells: Cells 1 and 2 (3.6C(80%)–3.6C); Cells 4 and 5 (4C(80%)–4C); Cells 7 and 8 (4.8C(80%)–4.8C); Cells 17 and 18 (5.4C(80%)–5.4C).

To intuitively illustrate the fast charging process adopted in this study, an example of the current and voltage profiles with 9-step fast charging during one aging cycle is shown in Fig. 2. Notably, the batteries are charged by applying CC with varying C-rates after setting a fixed incremental SOC value for each charging step, e.g., C_1 is in the charging of 0−8% SOC range, and C_2 is in the charging of 8−20% SOC range, C_3 is in the charging of 20−30% SOC range, etc. The more detailed protocol for SOC allocation is summarized in Table II.

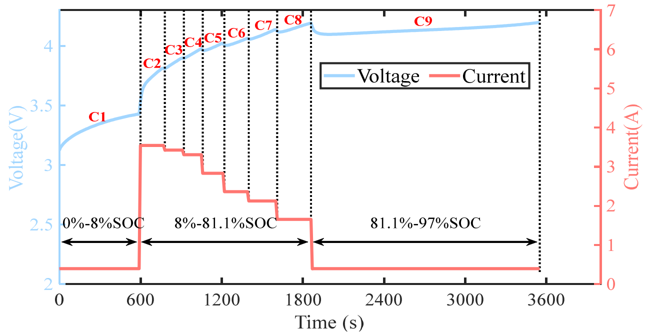

Fig. 2. Voltage and current profiles of the 9-step fast charging protocol.

TABLE II
SOC ALLOCATION FOR THE 9-STEP FAST CHARGING PROTOCOL

Step	Charging rate (C-rate)	Time (min)	SOC (%)
1	+0.33	14.54	+8.0
2	+3.00	2.40	+12.0
3	+2.90	2.07	+10.0
4	+2.80	2.14	+10.0
5	+2.40	2.50	+10.0
6	+2.00	3.00	+11.1
7	+1.80	3.33	+10.0
8	+1.40	4.29	+10.0
9	+0.33	28.93	+15.9
Total:	——	61.20	+97.0

III. PROPOSED METHODOLOGY

A. Proposed SOH Estimation Framework

The overview of the battery SOH estimation framework and the corresponding pseudocode are presented in Fig. 3 and Supplementary Table S1, integrating GPT-powered feature extraction with the proposed TS-KELM-AE model. During the feature extraction process, we directly upload the Q–V aging sequence data and send task-specific textual prompts to GPT, which automatically extracts 46 aging-related features. This task-prompt driven approach is highly effective, more intelligent, and user-friendly through text-based interaction. It overcomes the need for complex aging data preprocessing and the associated hyperparameter tuning involved in the design of traditional feature engineering algorithms. Using the extracted features, the proposed TS-KELM-AE architecture is introduced to form a coherent and efficient SOH estimation pipeline, offering both high estimation accuracy and computational efficiency. Specifically, the model operates in two stages: the first stage is dedicated to extracting deep features, while the second stage focuses on making decisions for accurate SOH estimation. Finally, the feasibility, effectiveness, and generalization capability of GPT-powered feature extraction and the TS-KELM-AE model are evaluated under various scenarios.

This article has been accepted for publication in IEEE Transactions on Transportation Electrification. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTE.2026.3683190

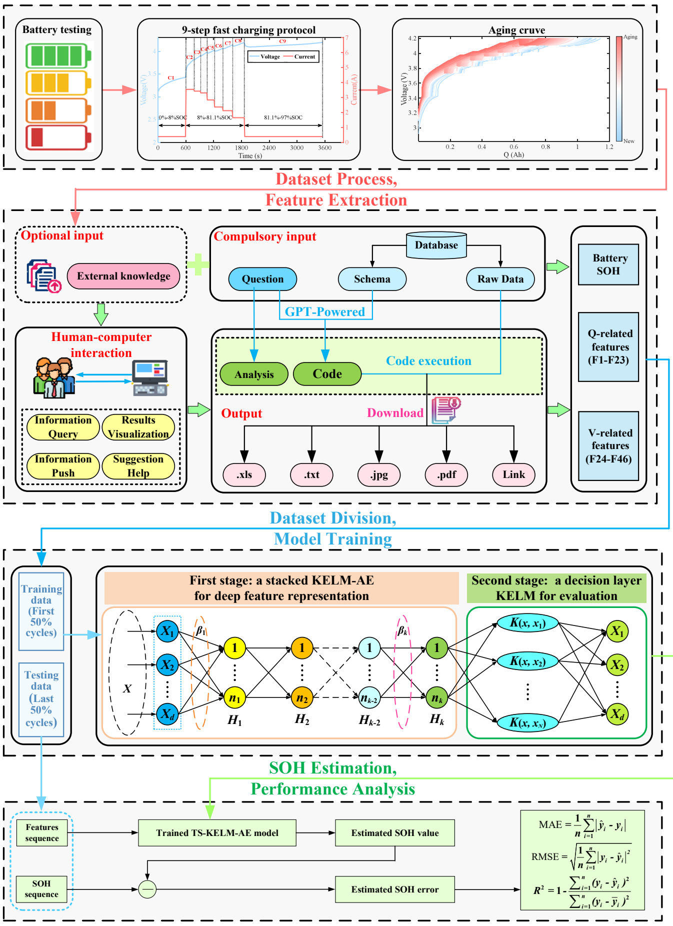

Fig. 3. Overview of the proposed battery SOH estimation framework.

B. GPT-Powered Automatic Feature Extraction

For data-driven SOH estimation approaches, charging data such as current (I), voltage (V), etc., which contain extensive battery degradation information, can serve as inputs [24]. Since the I data remains constant during different charging steps in Fig. 2, it is not an ideal candidate for direct input. In contrast, charging capacity (Q), obtained by integrating current over time, is more suitable for reflecting the process of cumulative battery degradation. Therefore, we input the Q and V sequences into the GPT model to automatically extract battery aging features from both time and frequency domain perspectives. Figs. 4 and 5 illustrate a general prompt structure and a specific dialogue example flow for battery aging feature extraction in this study, respectively.

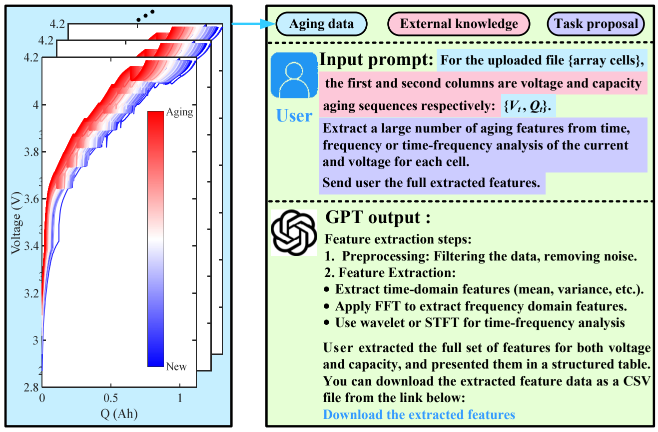

Fig. 4. Feature extraction example using a user-input prompt in GPT.

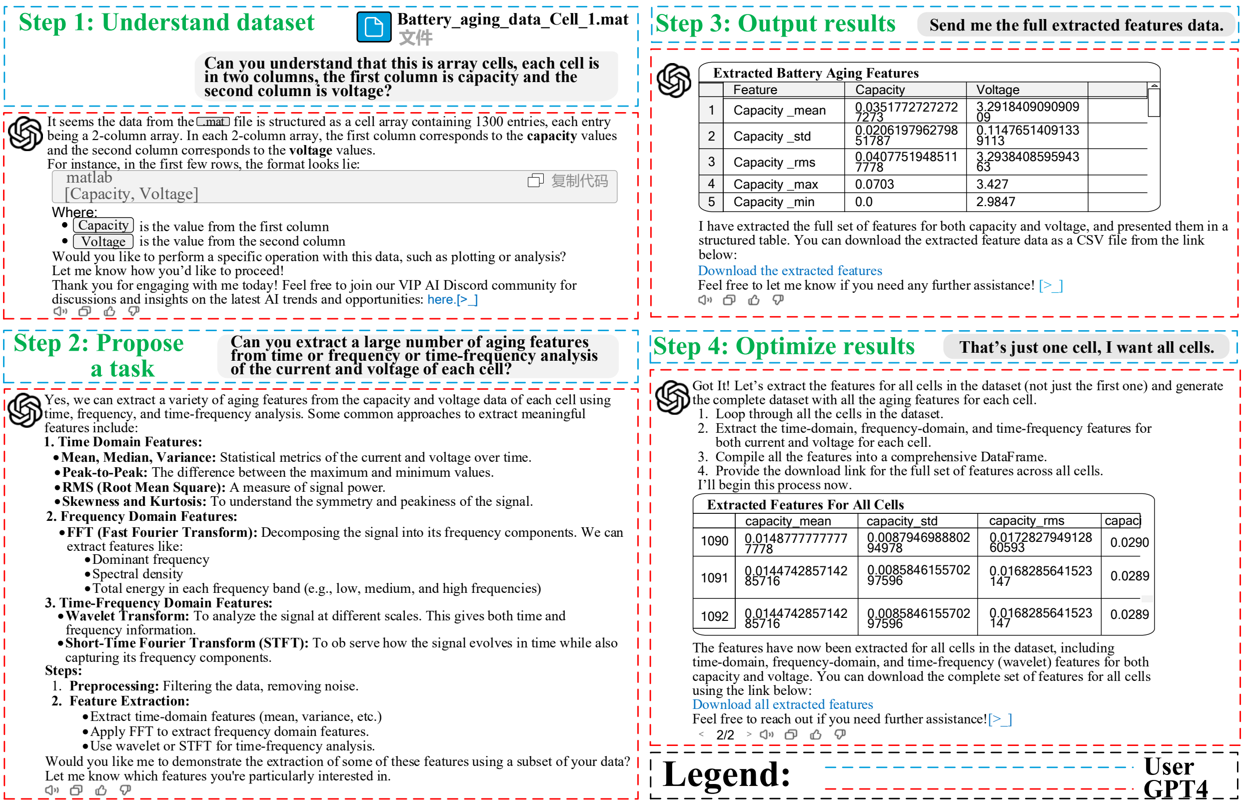

Fig. 5. Example dialog flow for battery feature extraction via GPT-4 prompt engineering.

This article has been accepted for publication in IEEE Transactions on Transportation Electrification. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTE.2026.3683190

Specifically, the user input prompt consists of three components: **1) aging data**, delivering charging information throughout the entire battery lifecycle; **2) external knowledge**, providing domain background and contextual information that assist GPT in interpreting data and enhancing the feature extraction process; and **3) task proposal**, defining the objectives that guide GPT in generating the required features and returning them to the user.

Conventional feature engineering relies on manually predefined features based on fixed formulas. In contrast, GPT acts as a semantic reasoning agent that automatically filters and reorganizes feature candidates using contextual dependencies in both the data and the task description. Conditioned on this structured input, GPT's output can synthesize task-specific feature representations by integrating domain interpretability with data-driven statistical relevance.

The GPT's outputs include the steps for performing the feature extraction tasks, as well as the download links (CSV files) of the required features. The returned tables are subsequently parsed in MATLAB. We discarded any rows containing non-numeric entries or inconsistent lengths, and performed validation checks on finite values and basic physical consistency (e.g., non-negativity requirements for square features, and bounded ranges for normalized quantities). Based on this, we can directly utilize the validated numerical features for subsequent feature correlation analysis and input them into data-driven models for battery SOH estimation. The extracted time-domain features include mean, standard deviation, maximum, minimum, root mean square (RMS), kurtosis, skewness, median, and energy. Frequency-domain features, computed using fast Fourier transform (FFT), include peak frequency, spectrum entropy, RMS spectrum, mean spectrum, bandwidth, and average frequency. Energy in the frequency-domain features consists of mean energy, maximum energy, and energy standard deviation. Frequency band energy is classified into low-frequency band energy (≤ 50 Hz), mid-frequency band energy (50-200 Hz), and high-frequency band energy (> 200 Hz). In addition, based on wavelet analysis, it also includes wavelet energy and correlation. In this study, 23 types of features are available from GPT-powered feature extraction for Q and V aging sequences, and are summarized in Table S2 in the Supplementary Materials. The features related to capacity Q are labeled F1-F23, and corresponding features related to voltage V are labeled F24-F46. To validate the reliability of the features extracted by GPT, we performed the Pearson correlation analysis between these features and battery SOH according to

$$P_{XY} = \frac{\sum_{i=1}^{n}\left(X_i - \bar{X}\right)\left(Y_i - \bar{Y}\right)}{\sqrt{\sum_{i=1}^{n}\left(X_i - \bar{X}\right)^2}\sqrt{\sum_{i=1}^{n}\left(Y_i - \bar{Y}\right)^2}} \tag{2}$$

where X_i and Y_i indicate GPT extracted features (F1-F46) and battery SOH at ith cycle, $\bar{X}$ and $\bar{Y}$ denote their average values. Figs. 6(a) and (b) illustrate the correlation heatmaps of the aging features extracted from the first CC charging step at 25 °C for Q and V features, respectively. This charging step was selected due to its highly stable degradation behavior and repeatable thermodynamic region to ensure cross-cycle correlation analysis. For each cycle, the number of data rows corresponds to the full sampling window of this CC stage, and details are provided in Fig. R5 in the Supplementary Materials.

In Fig. 6(a), the correlation values in the first column and the area proportion of pie charts in the first row indicate that among all the GPT-powered feature types, approximately 60% of the capacity-related features exhibit a correlation above 0.9, while Fig. 6(b) indicates that around 70% of the voltage-related features exhibit a correlation greater than 0.9.

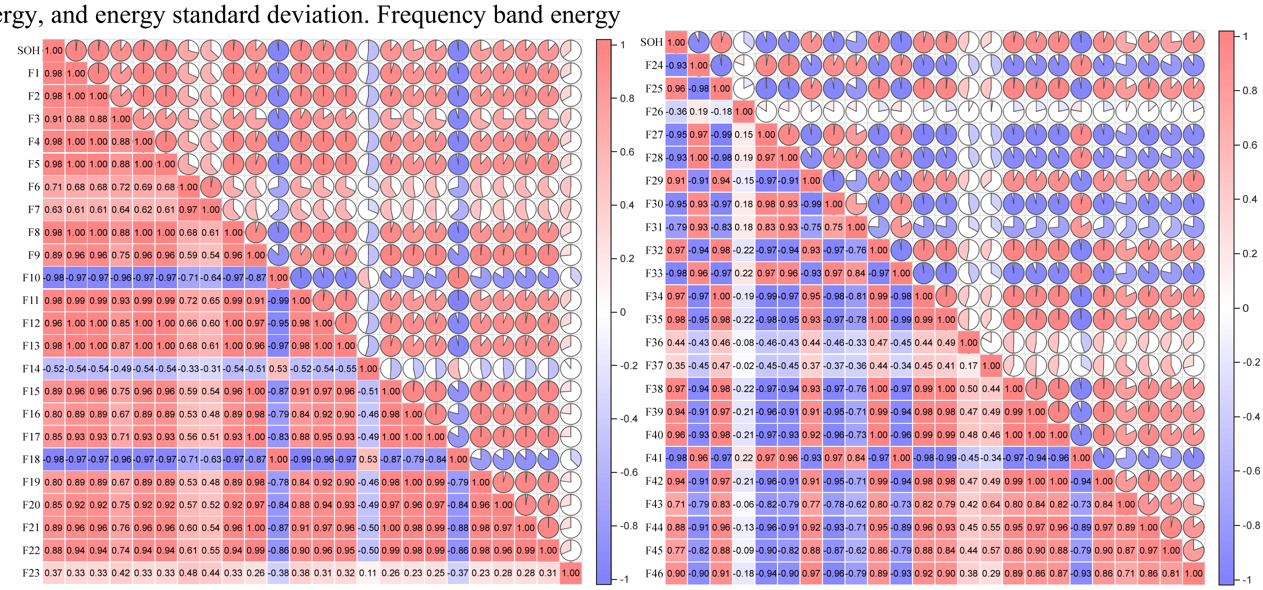

(a) (b)

Fig. 6. Correlation heatmaps of 23 types of GPT-extracted aging features and battery SOH. (a) Capacity-related features. (b) Voltage-related features. The lower triangular matrix shows the Pearson correlation coefficients, while the upper triangular matrix represents the correlation strength by the area of the corresponding pie charts.

We further extended the correlation analysis from 25 °C to three different temperatures (35 °C, 45 °C, and 55 °C) and compared the correlation evolution trends in Fig. 7. From Fig. 7(a), we discover that the color difference of correlation between capacity-related features and SOH is relatively small with increasing temperature, while in Fig. 7(b), the correlation

This article has been accepted for publication in IEEE Transactions on Transportation Electrification. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTE.2026.3683190

color difference between voltage-related features and SOH fluctuates more significantly. This observation suggests that voltage-related features are more sensitive to temperature changes, which may disadvantage the accurate SOH estimation. For instance, voltage features F24, F26, F27, F28, F31, F33, and F41 show a clear shift from negative to positive correlation with increasing temperature, while most capacity-related features maintain a relatively stable positive correlation. Consequently, we naturally hypothesize that in SOH estimation across varying temperatures, relatively stable capacity-related features may be more effective than voltage-related features.

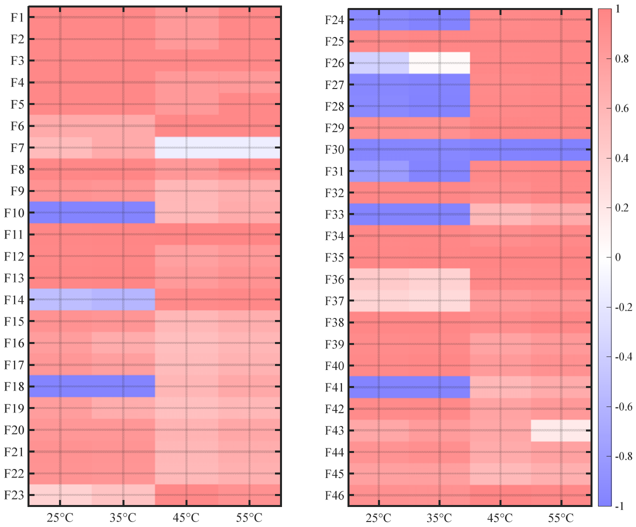

(a) (b)

Fig. 7. Evolution of correlation heatmap of 23 aging features at different temperatures (a) Capacity-related features. (b) Voltage-related features.

C. *Two-Stage KELM-AE Model*

1) *Kernel Extreme Learning Machine (KELM)*

As an efficient learning algorithm, ELM is distinguished from gradient descent-based algorithms. In ELM, the weights and biases of its single hidden layer are randomly initialized with an arbitrary continuous probability distribution. This training approach eliminates the necessity for time-consuming and computationally expensive parameter fine-tuning [22]. The standard ELM is described as follows. Given N training samples $\{(x_i, y_i)\}_{i=1}^{N}$, and denoting $X=[x_1,x_2,\cdots,x_N]^{\mathrm{T}}$ and $Y=[y_1,y_2,\cdots,y_N]^{\mathrm{T}}$ the input and corresponding output target data, respectively, the output of ELM with L hidden nodes can be expressed as:

$$y_i=\sum_{j=1}^{L}\beta_j g(\omega_j^{\mathrm{T}}x_i+b_j),\quad i=1,2,\ldots,N, \tag{3}$$

where ω_j and b_j are the connection weight and bias, respectively, β_j denotes the jth output weight, and $g(\cdot)$ represents an activation function. Equation (3) can be further expressed in matrix form as:

$$H\beta=Y \tag{4}$$

where $H\in\mathbb{R}^{N\times L}$ represents the matrix format of the hidden state output, and $\beta\in\mathbb{R}^{L\times 1}$ denotes the weight vector of the output layer, and $Y\in\mathbb{R}^{L\times 1}$ corresponds to the final output target matrix, i.e.,

$$H=\begin{bmatrix}h_1\\ \vdots\\ h_N\end{bmatrix}=\begin{bmatrix}g(\omega_1x_1+b_1) & \cdots & g(\omega_Lx_1+b_L)\\ \vdots & \ddots & \vdots\\ g(\omega_1x_N+b_1) & \cdots & g(\omega_Lx_N+b_L)\end{bmatrix}$$

$$\beta=[\beta_1,\beta_2,\cdots,\beta_L]^{\mathrm{T}}$$

By introducing a regularization coefficient λ into (3), ELM can be seen as a joint optimization least squares problem:

$$\min_{\beta}=\|H\beta-Y\|_2^2+\lambda\|\beta\|_2^2 \tag{5}$$

where $\lambda > 0$ is the regularization parameter that controls the balance between data fitting and weights.

Based on Karush–Kuhn–Tucker conditions, the closed-form primal solution $\hat{\beta}$ for output weights β can be expressed as:

$$\hat{\beta}=H^{\mathrm{T}}(HH^{\mathrm{T}}+\frac{I}{\lambda})^{-1}Y, \tag{6}$$

where I is the identity matrix. Since the parameters ω_j and b_j are randomly generated after pre-setting the number of hidden layer nodes in ELM, it can lead to defects such as unstable output and large predictive errors, especially when dealing with nonlinear data. The problem of the nonlinear learning capacity and stability of ELM can be solved using kernel techniques in KELM. The kernel matrix in KELM can be defined as:

$$\Omega_{\mathrm{ELM}(i,j)}=h_i\bullet h_j=K(x_i,x_j)=\exp(-\frac{\|x_i-x_j\|^2}{2\sigma^2}) \tag{7}$$

where $\bullet$ represents the dot product operation, K represents the radial basis function kernel, and σ denotes the corresponding width of the kernel function. According to (3), (6), and (7), the output function of the KELM is as follows:

$$f(x)=\begin{bmatrix}K(x,x_1)\\ \vdots\\ K(x,x_N)\end{bmatrix}^{T}\left(\frac{I}{\lambda}+K_{ELM}\right)^{-1}Y \tag{8}$$

The corresponding optimization problem for (8) and its solution are:

$$\min_{\beta}=\|K\beta-Y\|_2^2+\lambda\beta^{\top}K\beta \tag{9}$$

$$\hat{\beta}_{\mathrm{KELM}}=(\frac{I}{\lambda}+K)^{-1}Y \tag{10}$$

2) *Autoencoder (AE)*

AE is commonly applied in unsupervised learning tasks, which can automatically identify an effective data compression approach to discover potential features and structures within the data [15]. Specifically, the fundamental principle of AE can be divided into two main components: 1) the encoder, which maps the input data to a potential low-dimensional space, effectively compressing the input data into a compact representation; and 2) the decoder, which reconstructs the original input data from the encoder as accurately as possible, which captures the essential characteristics of the data. The mathematical formulation of the AE process is as follows:

$$Z=g^{*}(E\bullet X+B) \tag{11}$$

$$\tilde{X} = g^*(D \bullet Z + \tilde{B}) \tag{12}$$

where Z represents the compressed form of input data X produced by the encoder, while $\tilde{X}$ denotes the reconstructed form of input Z generated by the decoder; E and D represent weight matrices for encoding and decoding, respectively, while B and $\tilde{B}$ denote the corresponding bias matrices, $g^*(\cdot)$ refers to the activation function. Similarly, by extending the concept of AE to KELM, it can be further derived that the output weights of the KELM-AE from (8) are as follows:

$$\hat{\beta}_{\text{KELM-AE}} = (\frac{I}{\lambda} + K)^{-1} X \tag{13}$$

Despite the KELM-AE as a specific variant of ELM, it remains a shallow architecture, which constrains its expressive ability to capture deep correlations and underlying patterns in high-dimensional data, particularly with multi-type aging characteristics as input. Consequently, there is an urgent need for more advanced network architectures and deeper models that can address these challenges.

3) Two Stage KELM-AE (TS-KELM-AE)

Inspired by the hierarchical greedy training idea of deep learning, this study adopts the KELM-AE constructed by KELM and AE as the fundamental building block, and based on it, we propose a two-stage KELM-AE model for estimating battery SOH, which is demonstrated in Fig. 8.

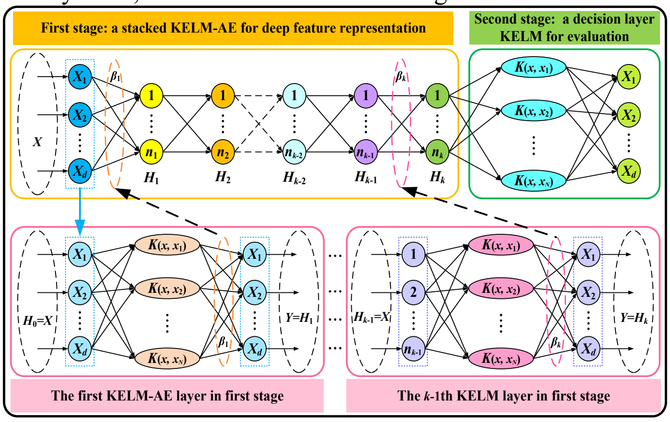

Fig. 8. Architecture and computational process of the TS-KELM-AE model.

The proposed model can be divided into two stages: in the first stage, a stacked KELM-AE network with k layers is constructed for deep feature extraction, which enables the model to learn hierarchical and abstract representations of the input data; in the second stage, the output H_k of the final hidden layer is transferred to the decision-making KELM to estimate SOH.

Assuming the TS-KELM-AE contains a total of k+1 hidden layers, according to (3) and (12), its kth layer output weights and hidden layer outputs can be described as follows:

$$\hat{\beta}_k = (\frac{I}{\lambda_{k-1}} + K)^{-1} H_{k-1} \tag{14}$$

$$H_k = g(H_{k-1} \bullet \hat{\beta}_k) \tag{15}$$

Specifically, the data sample X is simultaneously considered as the input and target output of the 1st KELM-AE unit, i.e., $H_0 = X$, the output weights β_1 and the hidden nodes output H_1 are derived according to (13) and (14), respectively. Subsequently, H_1 is used as the input and target output for the 2nd KELM-AE unit, i.e., $H_1 = X$, and the output weights β_2 and hidden nodes output H_2 are calculated respectively for the 2nd KELM-AE unit, and so on, iterative training until H_k for the kth KELM-AE unit is obtained. Eventually, the hidden layer output H_k is fed into (7) to solve the TS-KELM-AE model. More importantly, within each KELM-AE module, the regularization coefficient λ_{k-1} plays a crucial role in balancing model approximation and generalization performance. In the explored range of 0.001 to 1000, smaller coefficients impose stronger regularization, which may lead to underfitting and increased SOH error, while excessively large coefficients can cause overfitting and unstable aging trajectory reconstruction. Therefore, we increment its value layer by layer to ensure hierarchical control over the degradation representation.

D. SHAP Model

To quantify the contribution of each input feature to the SOH estimation, the Shapley additive explanations (SHAP) method is employed in this study. The SHAP provides a unified model-agnostic approach for interpreting the proposed TS-KELM-AE model by assigning consistent contribution values to each feature based on cooperative game theory [30].

The Shapley value of the jth feature Φ_j is defined as:

$$\Phi_j = \sum_{S \subseteq F \setminus \{j\}} \frac{|S|!(|F|-|S|-1)!}{|F|!} \left[f(S \cup \{j\}) - f(S) \right] \tag{16}$$

where F denotes the full set of input features, $S \subseteq F \setminus \{j\}$ represents any subset that excludes feature j; $f(S)$ and $f(S \cup \{j\})$ represent the outputs of the TS-KLEM-AE model when only the feature subset S is present and when feature i is added to the subset S, respectively.

Based on these Shapley values, the surrogate output $m(z)$ is expressed as a linear combination of binary feature indicators and their corresponding Shapley contribution values:

$$m(z) = \Phi_0 + \sum_{j=1}^{M} \Phi_j z_j \tag{17}$$

where $z_j \in \{0,1\}$ is a binary indicator reflecting whether feature j is included in the coalition, M is the number of features.

E. Performance Metrics

To comprehensively evaluate the performance of the TS-KELM-AE model in estimating battery SOH, it is essential to use a set of well-established metrics. In this study, three evaluation metrics are employed, including the mean absolute error (MAE), the root mean square error (RMSE), and the coefficient of determination (R^2). These metrics provide insights into the model's estimation accuracy, magnitude of errors, and ability to explain the variance in the true SOH values.

$$\text{MAE} = \frac{1}{M} \sum_{i=1}^{M} |y_i - \hat{y}_i| \tag{18}$$

$$\text{RMSE} = \sqrt{\frac{1}{M} \sum_{i=1}^{M} |y_i - \hat{y}_i|^2} \tag{19}$$

$$R^2 = 1 - \frac{\sum_{i=1}^{M} (y_i - \hat{y}_i)^2}{\sum_{i=1}^{M} (y_i - \bar{y}_i)^2} \tag{20}$$

where M indicates the overall number of aging cycles; y_i, $\hat{y}_i$, and $\bar{y}_i$ denote the actual SOH, the estimated SOH, and the average of the actual SOH values, respectively.

IV. Results and Discussion

A. Selection of Model Parameters

In constructing the TS-KELM-AE model, the number of layers selected for the stacked KELM-AE network during the first stage directly impacts the depth of feature extraction and the model's complexity. Therefore, we tested different layer configurations to develop a high-accuracy SOH estimation model with low computational overhead, and evaluated the effect of network layers on the SOH estimation performance by four metrics, namely MAE, RMSE, R^2, and training time. All experiments were performed on a workstation equipped with a 13th Gen Intel® Core™ i7-13700H @ 2.40 GHz CPU, 16 GB RAM, and an NVIDIA GeForce RTX 4060 GPU, using MATLAB 2022b for SOH estimation model construction and training. The GPT-powered feature generation module was executed independently and accessed through the cloud API; it does not require GPU acceleration or local model deployment.

Using aging-battery data at 25 °C as a case in point, we employ all 46 features (F1-F46) as inputs and the battery SOH as the output, aiming to validate the capability of the proposed model to extract deep features. To ensure a fair assessment, a forward validation scheme is adopted based on chronological ordering. The first 50% of cycles are used for model training, while the remaining 50% are strictly reserved for testing and represent a more degraded and challenging scenario. During model training, the validation set defaults to the final 20% of the training interval, and the evaluation results are summarized in Table III. In addition, we present the normalized evaluation metrics using a radar chart in Fig. 9 for further intuitive comparison.

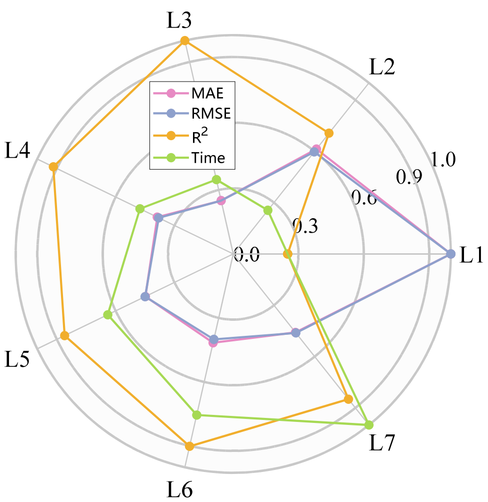

Fig. 9. Radar chart of four normalized evaluation metrics for different network layer configurations.

TABLE III
COMPARISON OF DIFFERENT LAYER CONFIGURATIONS

Number of layers	MAE (%)	RMSE (%)	R^2 (%)	Time (s)
1 (L1)	3.8551	4.8676	55.4773	**0.3431**
2 (L2)	2.3626	2.8958	80.4917	0.3524
3 (L3)	**0.9589**	**1.2036**	**96.6297**	0.5139
4 (L4)	1.4898	1.8397	92.1264	0.7336
5 (L5)	1.7243	2.1789	88.9558	1.0178
6 (L6)	1.6017	1.9391	91.2546	1.2152
7 (L7)	1.7649	2.2383	88.3459	1.6385

TABLE IV
KEY HYPERPARAMETERS FOR 3-LAYER CONFIGURATIONS

Layer	Regularization coefficient	Kernel function width
Layer 1	0.0251	6.7376
Layer 2	1.2664	8.2342
Layer 3	156.9548	33.0308

Table III and Fig. 9 show that the model accuracies are improved as the number of layers increases from L1 to L3. Particularly, L3 exhibits the best performance in MAE, RMSE, and R^2 metrics with 0.9589%, 1.2036%, and 96.63%, respectively, and the training time is relatively short, which indicates L3 has smaller estimation errors and can effectively explain most of the data regression. However, further increasing the number of layers from L4 to L7 results in a slight negative effect on model performance, along with a significant increase in computational burden, especially after L5, where the training time exceeds 1 second. Therefore, the number of KELM-AE network layers in the first stage is selected as 3 to achieve the best balance between model performance and computational efficiency. Correspondingly, Table IV summarizes the key hyperparameter settings under the L3 layer optimal configuration, including the kernel function width and regularization coefficient, which provides a clear basis for the subsequent SOH estimation model.

B. Comparative Evaluation of Different Feature Extractors

In this section, three representative deep learning baselines operating directly on V–Q sequences, namely 1D CNN, LSTM, and Transformer networks, were implemented for SOH estimation under four temperature conditions. For a fair architectural comparison, all three models were designed as compact two-layer encoders with ReLU activation in the hidden layers, ensuring that performance differences are attributable to the underlying representation mechanisms. The configuration of these baselines follows design choices that have been reported as effective in recent SOH and knee-prediction studies [19], [25], [26], including batch size of 32, a maximum of 300 training epochs, convolution kernel size of 3 with 32 filters for the CNN, 180 hidden units for the LSTM, and 4 attention heads for the Transformer. All deep baselines were trained under the same training/validation split scheme and optimized using the Adam optimizer. For the critical hyperparameter of learning rate, we performed multiple trials over the range {1e−4, 3e−4, 1e−3} and selected 0.0015 based on validation performance. Across all four temperature datasets, the CNN, LSTM, and Transformer models exhibited smooth convergence of both training loss without significant optimization instability issues, and the corresponding results are provided in Fig. R7 of the Supplementary Material. The average runtime for each algorithm execution was approximately 550 seconds, primarily involving the offline training and hyperparameter optimization processes required for deep learning benchmark models. It should be noted that the computational overhead is acceptable in an offline deployment model aligned with cloud-based digital twin architectures, as model training and optimization occur on high-performance servers rather than resource-constrained edge devices. In this setting, the digital twin continuously synchronizes aging-related features from the physical battery and updates them in the cloud environment. Under this framework, extracted degradation features can be provided as input to the proposed TS-KELM-AE model for systematically comparing GPT-powered SOH estimation under varying temperature conditions. The SOH estimation results under different feature extractors are shown in Fig. 10, and the four metrics are summarized in Table V.

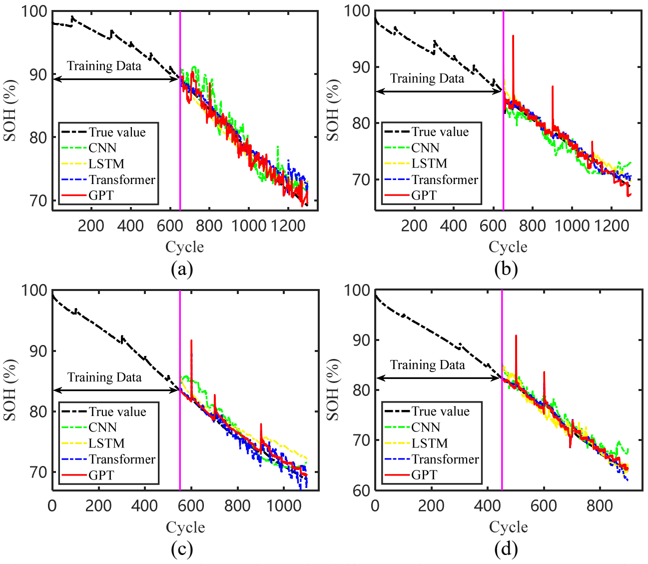

Fig. 10. Experimental results of different feature extractors at four temperatures. (a) 25 °C. (b) 35 °C. (c) 45 °C. (d) 55 °C.

From Fig. 10, it can be observed that the CNN-based feature extractor is capable of producing an acceptable fit to the SOH degradation trend under high-temperature conditions (55 °C). However, it fails to accurately track the degradation behavior at lower temperatures (25 °C). This discrepancy highlights the inherent limitation of CNN; the local receptive fields restrict their capacity to capture long-term degradation features in the V-Q sequence. The performance of LSTM and Transformer in tracking SOH degradation is found to be comparable, with both exhibiting a certain capability to capture long-term aging features. However, as indicated by the evaluation metrics presented in Table V, Transformer consistently outperforms LSTM in terms of SOH estimation accuracy at multiple temperature conditions, which indicates that its self-attention mechanism offers more advantages for extracting global degradation features in V-Q sequences. Inspiringly, the GPT-powered feature extractor achieves the best SOH estimation accuracy and degradation trend tracking at all four temperatures. Compared with the mainstream deep learning mentioned above, the GPT-powered feature extraction method demonstrates a superior ability to integrate domain-specific prior knowledge with sequential aging patterns, which enables highly efficient extraction of global degradation features across different cycles and temperatures.

Notably, the model is implemented without utilizing any labels or performing online updates during testing. The SOH spike phenomenon of "self-correction" in the model originates from a transient feature scale shift induced by cycle duration variations within the 9-step fast charging protocol. Once the operational state (feature distribution) returns to the dominant regime, the robust global mapping learned by TS-KELM-AE naturally guides the output values back to the main degradation trajectory.

TABLE V
EVALUATION METRICS OF DIFFERENT FEATURE EXTRACTORS

Operating temperature	Feature extractor	MAE (%)	RMSE (%)	R^2 (%)
25 °C	CNN	1.8593	2.2361	88.3676
	LSTM	1.2789	1.5315	94.5431
	Transformer	1.1639	1.6259	93.8496
	GPT	**0.9589**	**1.2036**	**96.6297**
35 °C	CNN	1.5506	1.8821	86.7759
	LSTM	0.8795	1.0809	95.3102
	Transformer	0.7287	**0.8831**	**96.8692**
	GPT	**0.6407**	0.9741	96.1909
45 °C	CNN	1.6381	1.9964	81.5718
	LSTM	0.9426	1.0568	94.8367
	Transformer	0.8253	1.1007	94.3986
	GPT	**0.7898**	**1.0122**	**95.2626**
55 °C	CNN	1.4894	1.8849	90.7044
	LSTM	1.0182	1.2587	95.8545
	Transformer	0.6712	0.9329	97.1846
	GPT	**0.5758**	**0.8993**	**97.8837**

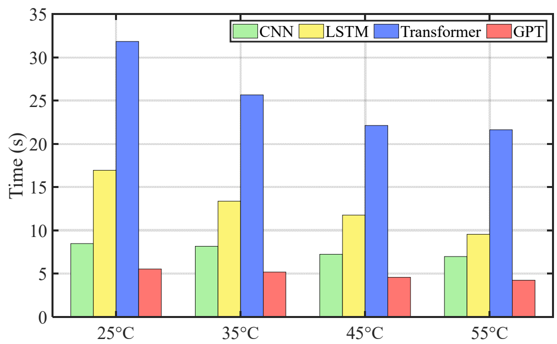

Fig. 11. Comparison of running times of different feature extractors at four temperatures. (a) 25 °C. (b) 35 °C. (c) 45 °C. (d) 55 °C.

Based on the above comparisons, we further evaluate the efficiency of different feature extractors from the perspective of runtime, as illustrated in Fig. 11. Compared with mainstream deep learning-based feature extractors that directly process high-dimensional V-Q sequences, the GPT-based approach rapidly compresses each charging cycle into a 46-dimensional feature vector. This efficiency benefits from large-scale pre-training that can achieve effective pattern abstraction and signal representation. As a result, it is capable of rapidly compressing and extracting battery aging data information without the need for complex parameter tuning or a backpropagation network structure.

To further validate that the GPT-extracted features are correct and effective, we conducted an additional benchmark experiment with one battery SOH researcher who had not accessed our manuscript. The researcher was provided only the names and physical interpretation of 46 features, without formulas, and was asked to extract them from the same battery datasets. The results are shown in Figs. R1-R4, with the corresponding metrics summarized in Table S3, all in the Supplementary Materials. Analysis indicates that GPT not only reproduces expert-level features, but also autonomously discovers physically meaningful frequency-domain descriptors (e.g., spectral mean frequency for capacity and voltage features). The development time reduced from one week to 1.5 hours, reflecting a substantial reduction in time cost and labor dependency. Moreover, the GPT feature set yields comparable or superior SOH estimation accuracy under different operating conditions, and further detailed comparisons of advantages are summarized in Table S4 in the Supplementary Materials.

C. *Selection of Input Feature Type*

In this section, to further validate the effectiveness of the GPT-powered feature extraction for battery SOH estimation, as well as the conjecture proposed in Section III-B that the capacity-related features provide more accurate SOH estimation compared to voltage-related features under varying temperature conditions. We performed four types of feature ablation experiments: the combined voltage and capacity-related features F1-F46 (Type 1); the capacity-related features F1-F23 (Type 2); the voltage features F24-F46 (Type 3); and the manual feature extraction method for voltage (Type 4), as discussed in [19].

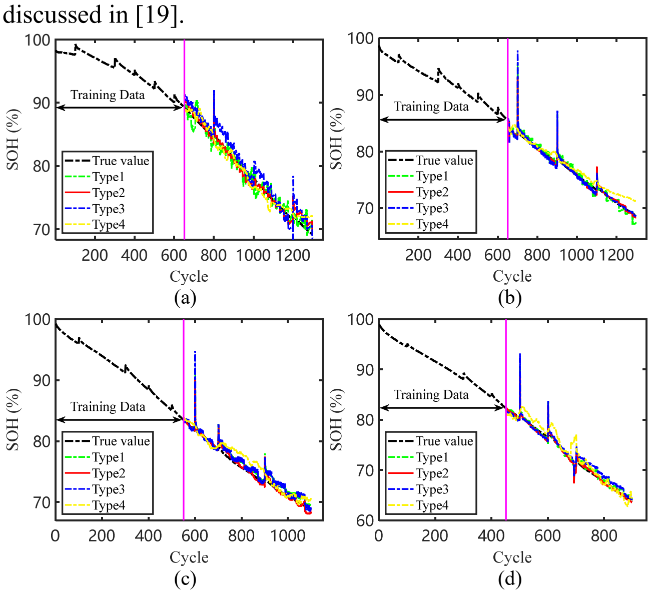

Fig. 12. Experimental results of feature ablation by GPT and manual extraction at four different temperatures. (a) 25 °C. (b) 35 °C. (c) 45 °C. (d) 55 °C.

Fig. 12 presents the SOH estimation results at four different temperatures. From comparing the estimated results of SOH regeneration points with the real values, the aging features of Types 1, 2, and 3 can effectively match the battery SOH regeneration points, which indicates these features are able to accurately capture the SOH regeneration cycle across different temperatures. Particularly at higher temperatures, these features still reliably track the SOH degradation trajectory, showing strong adaptability. Although the Type 4 estimation results at the SOH regeneration point show a smaller amplitude deviation, from the perspective of the overall degradation trend, the estimation results of Types 1, 2, and 3 are closer to the true value. This phenomenon may be attributed to the large amplitude difference between charging cycle data in regeneration points and earlier charging cycle data, which further leads to excessive variation for aging features extracted by GPT in SOH regeneration points.

TABLE VI
EVALUATION METRICS OF FEATURE ABLATION EXPERIMENTS

Operating temperature	Input feature	MAE (%)	RMSE (%)	R^2 (%)
25 °C	Type 1	0.9589	1.2036	96.6297
	Type 2	**0.7207**	**0.9405**	**97.9424**
	Type 3	1.4408	1.6783	93.4468
	Type 4	0.9931	1.2176	96.5508
35 °C	Type 1	0.6407	0.9741	96.1909
	Type 2	**0.3864**	**0.6716**	**98.1891**
	Type 3	0.4129	0.8616	97.0196
	Type 4	1.0184	1.2654	93.5718
45 °C	Type 1	0.7898	1.0122	95.2626
	Type 2	**0.4153**	**0.6553**	**98.0151**
	Type 3	0.8837	1.2047	93.2898
	Type 4	1.1118	1.4209	90.6653
55 °C	Type 1	0.5758	0.8993	97.8837
	Type 2	**0.3245**	**0.6473**	**98.9036**
	Type 3	0.9085	1.2701	95.7789
	Type 4	1.2334	1.6419	92.9462

D. *Comparison With Different SOH Estimation Methods*

To further validate the superiority of the TS-KELM-AE model, we conducted comparative experiments with other methods, including ELM, SVM, LSTM, and KELM. Notably, we use capacity-related features as model inputs, namely Type 2, as mentioned in the previous section. Fig. 13 displays comparative experiments for SOH estimation results at four different temperatures. For brevity, only the RMSE results of all estimation methods are illustrated in Fig. 14. As we can see from Fig. 13, during the early SOH estimation, other methods attempt to track the SOH degradation trend observed in the training data. However, the deviations arise more significantly in the later stage of estimation, especially after the SOH regeneration point. In contrast, the TS-KELM-AE model remains closer to the true SOH variation throughout most of the cycle and demonstrates superior SOH tracking ability.

It is noteworthy that in Fig. 13(a), all methods exhibit significant fluctuations in their estimated results. A possible reason for this is the frequent occurrence of SOH regeneration points in the training data, which increases the variability of the extracted features. Additionally, compared with Figs. 13(b), (c), and (d), we observe that the SOH degradation rates of the training and testing data in Fig. 13(a) differ significantly. Specifically, the training part exhibits a relatively gentle degradation trend, while the testing set shows a steeper one, which poses a challenge for SOH estimation. Satisfyingly, our TS-KELM-AE method is always fluctuating around the true value, while the other methods are fluctuating away from the true value. This further demonstrates the significant

This article has been accepted for publication in IEEE Transactions on Transportation Electrification. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTE.2026.3683190

advantages of our proposed TS-KELM-AE in deep feature extraction and nonlinear mapping capabilities.

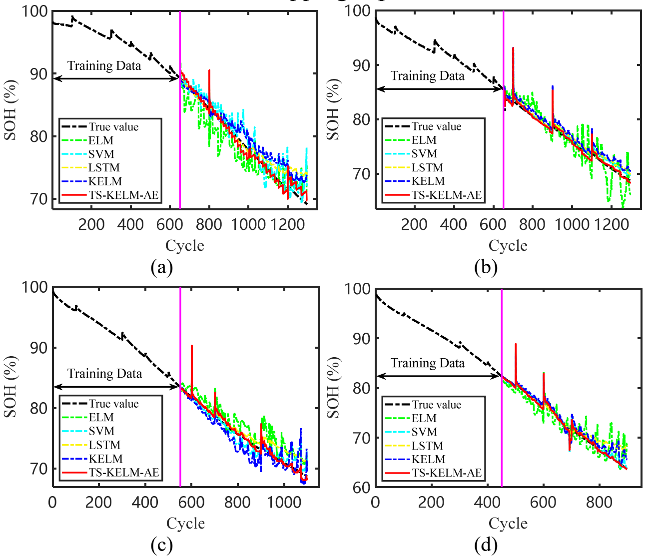

Fig. 13. Estimated comparison results at four different temperatures. (a) 25 °C. (b) 35 °C. (c) 45 °C. (d) 55 °C.

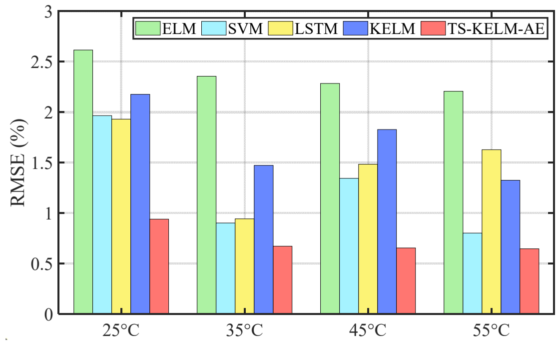

Fig. 14. Performance evaluation metrics of the comparative experiment at four different temperatures.

Specifically, the stacked KELM-AE structure of the first stage is a deep learning model that enables multi-level feature extraction to capture higher-order nonlinear relationships for describing the dynamic process of battery degradation. Moreover, the introduced kernel trick enhances its regression ability for nonlinear data. Therefore, the proposed method can more accurately characterize the degradation patterns of battery SOH and has a broad application potential.

Furthermore, to further validate the robustness and data efficiency of the proposed method, SOH estimation experiments were conducted under four operating temperatures using four training samples (50%, 40%, 30%, and 20%), and the corresponding estimation results are summarized in Table VII. As shown in the table, the proposed model maintains high accuracy across all settings, with MAE and RMSE confined to 0.32–2.30% and 0.65–3.01%, respectively, and R^2 consistently above 78% even under the lowest data size. As the training dataset size increases, model performance exhibits a monotonically increasing trend. For instance, at 25 °C, increasing the training size from 20% to 50% reduces MAE and RMSE by 68.6% and 68.8%, respectively, while R^2 improves by 24.0%. Comparable gains are seen across other temperatures: at 35 °C, MAE decreases from 0.85% to 0.39%, and R^2 rises from 95.98% to 98.19%; at 45 °C, MAE drops from 1.40% to 0.42%; and at 55 °C, the error reduction is even more pronounced, decreasing from 2.14% to 0.32%.

TABLE VII
PERFORMANCE OF THE TS-KELM-AE MODEL UNDER DIFFERENT TRAINING SCENARIOS

Operating temperature	Training size	MAE (%)	RMSE (%)	R^2 (%)
25 °C	50%	**0.7207**	**0.9405**	**97.9424**
	40%	0.8044	1.0191	97.2578
	30%	2.0665	2.5703	84.6297
	20%	2.2972	3.0111	78.9065
35 °C	50%	**0.3864**	**0.6716**	**98.1891**
	40%	0.4557	0.7221	97.7188
	30%	0.5399	0.9001	96.5501
	20%	0.8479	1.1945	95.9807
45 °C	50%	**0.4153**	**0.6553**	**98.0151**
	40%	0.9332	1.1636	94.0631
	30%	1.1335	1.4021	93.1009
	20%	1.4041	1.6256	90.7821
55 °C	50%	**0.3245**	**0.6473**	**98.9036**
	40%	1.1813	1.5921	94.2851
	30%	1.1587	1.6351	95.2478
	20%	2.1387	2.7012	91.7354

These results demonstrate that the proposed framework achieves accurate, stable, and highly generalizable SOH estimation across temperatures and data sizes. Even with only 20–30% of the available samples, the model maintains high estimation performance.

E. Feature Importance Analysis and Model Multicollinearity Assessment

To enhance the interpretability of the extracted 46 features, a SHAP-based feature importance analysis was conducted under four temperature conditions, as shown in Figs. 15–18. The results consistently show that frequency-domain features contribute substantially more than time-domain features. In particular, the peak frequency (F10) and spectral mean frequency (F18), which represent the dominant spectral energy location, consistently rank among the most important features in both capacity-based and voltage-based feature sets. This result demonstrates that GPT-powered spectral features can effectively capture nonlinear SOH degradation behaviors. In contrast, most of the remaining features show strong temperature-dependent variability, indicating they are associated with specific thermal pathways of degradation rather than serving as universal features.

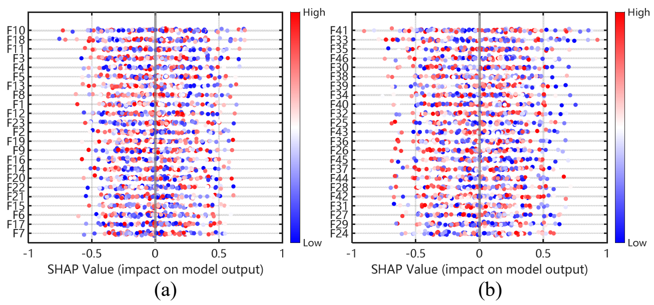

This article has been accepted for publication in IEEE Transactions on Transportation Electrification. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTE.2026.3683190

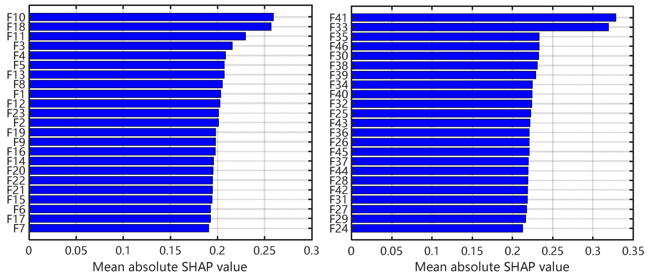

(c) (d)

Fig. 15. Contributions of SHAP-based 46 features at 25°C. (a) Beeswarm plot for F1-F23. (b) Beeswarm plot for F24-F46. (c) ranking of F1-F23 importance. (d) ranking of F24-F46 importance.

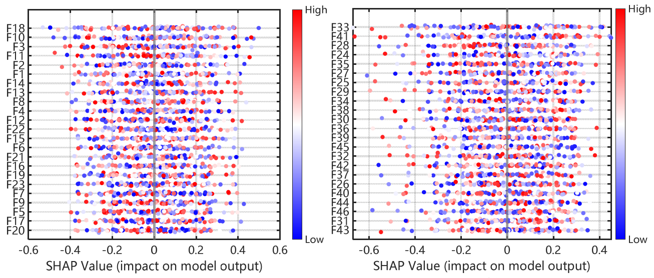

(a) (b)

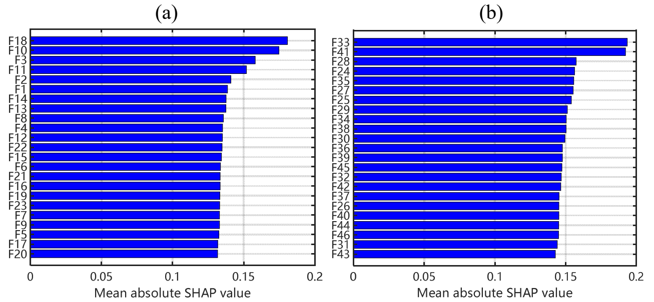

(c) (d)

Fig. 16. Contributions of SHAP-based 46 features at 35°C. (a) Beeswarm plot for F1-F23. (b) Beeswarm plot for F24-F46. (c) ranking of F1-F23 importance. (d) ranking of F24-F46 importance.

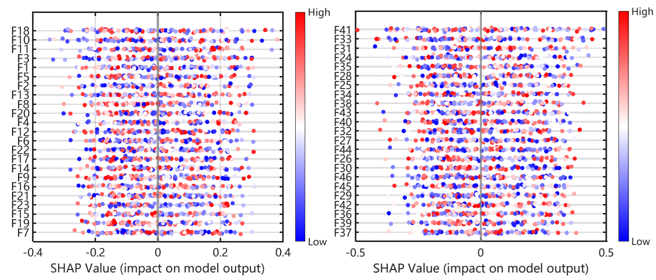

(a) (b)

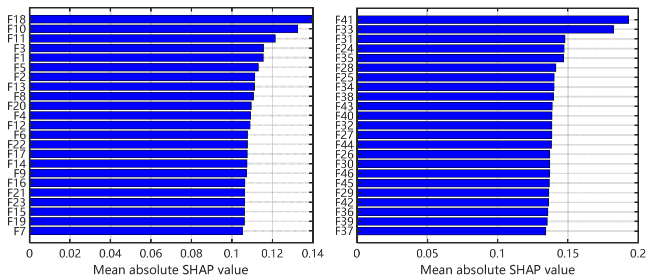

(c) (d)

Fig. 17. Contributions of SHAP-based 46 features at 45°C. (a) Beeswarm plot for F1-F23. (b) Beeswarm plot for F24-F46. (c) ranking of F1-F23 importance. (d) ranking of F24-F46 importance.

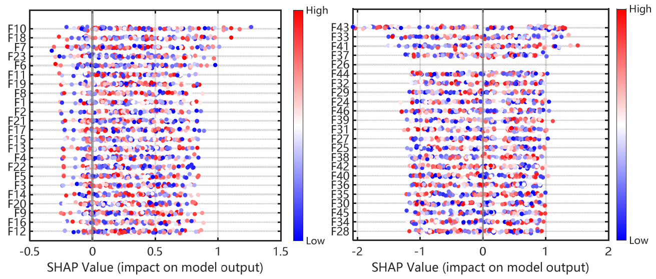

(a) (b)

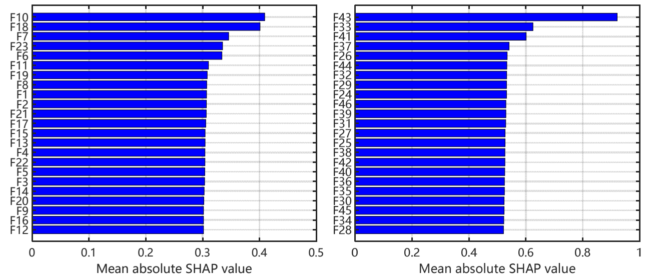

(c) (d)

Fig. 18. Contributions of SHAP-based 46 features at 55°C. (a) Beeswarm plot for F1-F23. (b) Beeswarm plot for F24-F46. (c) ranking of F1-F23 importance. (d) ranking of F24-F46 importance.

Furthermore, spectral entropy (F11) shows high importance only for capacity-based features, but contributes minimally to voltage-based features, revealing that entropy-based spectral dispersion plays a role only within certain signal modalities. These SHAP analysis results significantly improve the interpretability of the 46 features extracted by GPT and strengthen the physical consistency of the results under different thermal conditions. More importantly, they further indicate that frequency-domain features are generally more effective than time-domain features in capturing the nonlinear degradation associated with SOH evolution. Although the SHAP impact score ranges of different features are relatively close, the ranking still reflects their relative average marginal contributions to the model output. Therefore, SHAP-based interpretation is not used as an independent standard, but is further validated through subsequent feature elimination analysis and statistical validation of individual features. To further remove the possibility that the advantage originates from a few good features, we performed a feature elimination experiment guided by SHAP rankings. The top-K most contributive GPT-powered features were progressively removed, and the SOH estimation performance variation was evaluated accordingly. Specifically, when K=2, the two highest-ranking spectral features were removed from the capacity-derived feature set, namely peak frequency (F10) and spectral mean frequency (F18), both of which encode dominant SOH-related frequency shifts. When K=4, an equal number of the most contributive features were simultaneously removed from both capacity and voltage sequences (F10, F18, F33, and F41), resulting in four eliminated features. Finally, when K=6, an additional degenerate sensitive feature was removed from each domain, namely spectral entropy in the capacity-derived set (F11) and mid-band spectral energy (F43), forming a comprehensive progressive elimination scheme. The performance of SOH estimation results with progressively removed features is summarized in Table VIII.

This article has been accepted for publication in IEEE Transactions on Transportation Electrification. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TTE.2026.3683190

Table VIII more clearly and consistently demonstrates that performance declines sharply when only the first two key features are removed. For instance, at 25°C, the RMSE increases from 0.94% (K=0) to 1.51% (K=2), while the R^2 decreases from 97.94% to 94.73%; and similar decreases are observed at 35 °C, 45 °C, and 55 °C. As more high-ranking features are discarded, the R^2 value continued to decline by 2.2% to 10.2%, particularly under high-temperature conditions (e.g., 55 °C). Removing six features leads to an abrupt drop in the R^2 value from 98.90% to 88.79%, which indicates that high-temperature degradation modeling is highly dependent on the frequency-domain features captured by GPT. These results confirm that GPT can effectively construct nonlinear and higher-order derivative feature representations, thereby linearizing the nonlinear SOH degradation process of batteries and enabling TS-KELM-AE models to capture degradation dynamics with greater accuracy. Furthermore, we conducted single-feature regression models for the F18, F41, and representative manual features, and statistically compared the SOH estimation results. The corresponding results are presented in Fig. R6, Table S5, and Table S6 of the Supplementary Materials. The F18 and F41 yield significantly lower RMSE and higher R^2 values, with Wilcoxon signed-rank tests confirming statistical significance (p < 0.01 in most cases). These results demonstrate that the advantage of F18 and F41 extends beyond high correlation to translate into linearization and error reduction benefits.

TABLE VIII
MODEL PERFORMANCE AFTER REMOVING THE TOP-K FEATURES BY SHAP

Operating temperature	Removing features	MAE (%)	RMSE (%)	R^2 (%)
25 °C	K=0	**0.7207**	**0.9405**	**97.9424**
	K=2	1.1914	1.5052	94.7293
	K=4	1.2783	1.6419	93.7281
	K=6	1.4255	1.9206	91.4182
35 °C	K=0	**0.3864**	**0.6716**	**98.1891**
	K=2	0.6644	0.9195	96.6058
	K=4	0.8408	1.2331	93.8958
	K=6	0.9294	1.3424	92.7651
45 °C	K=0	**0.4153**	**0.6553**	**98.0151**
	K=2	0.7533	0.9892	95.4759
	K=4	0.8805	1.2034	93.3039
	K=6	1.0641	1.2868	92.3442
55 °C	K=0	**0.3245**	**0.6473**	**98.9036**
	K=2	0.7291	1.4075	95.7405
	K=4	0.8622	1.6050	93.2600
	K=6	1.1417	2.0698	88.7902

In addition, to demonstrate the capability of the proposed TS-KELM-AE model in mitigating multicollinearity among features, we compared the correlation structures of the raw GPT-extracted features (Fig. 6) with the latent representations learned by our proposed TS-KELM-AE model. As illustrated in Fig. 6, the raw GPT-extracted feature heatmaps exhibit strong pairwise correlations, with a mean absolute Pearson correlation of 0.8547. After passing through three layers of KELM-AE, the correlation heatmap of the TS-KELM-AE hidden layer feature progressively decreases, as shown in Figs. 19(a), (b), and (c); with the corresponding mean absolute Pearson correlation decreases from 0.7544 to 0.5695 and 0.4541. These results demonstrate that TS-KELM-AE can effectively perform nonlinear decorrelation and mitigate multicollinearity before the SOH regression stage.

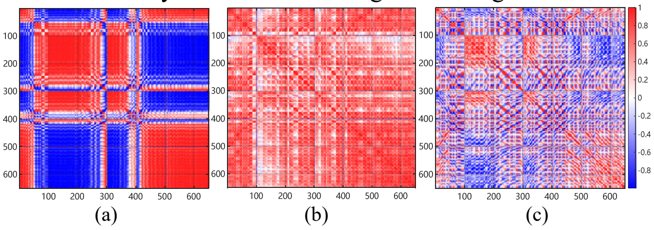

Fig. 19. Correlation heatmap of the TS-KELM-AE hidden layer feature. Weights for (a) the first hidden layer, (b) the second hidden layer, and (c) the final hidden layer.

F. Cross-Cell Verification Under Different Fast Charging Protocols and Temperatures

To further enhance the model's generalization analysis, cross-cell validation experiments were conducted under multiple fast-charging protocols. Building upon the original protocol, four supplementary fast charging conditions were considered: 3.6C(80%)–3.6C (Cells 1 and 2), 4C(80%)–4C (Cells 4 and 5), 4.8C(80%)–4.8C (Cells 7 and 8), and 5.4C(80%)–5.4C (Cells 17 and 18).

For each charging protocol, one cell was used for model training and validation, while another cell was fully reserved as an independent test set to assess cross-cell generalization. Throughout the training process, 70% of the cycles were used for training and the remaining 30% for validation. The test cell was never exposed to the model during training, ensuring that the evaluation strictly reflects the ability of the model to transfer to a different battery under the same fast charging protocol. In addition, we also adopted a bidirectional cross-cell protocol and its corresponding averaged results to enhance the rigor of SOH assessment. For instance, Cell 1 was trained while Cell 2 was tested, and Cell 2 was trained while Cell 1 was tested, and so on. This design helps verify whether the learned representations remain stable when the roles of training cells and test cells are reversed.

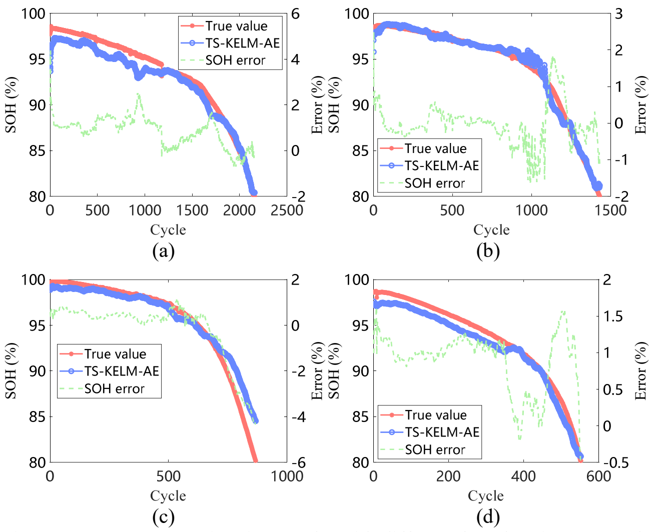

Fig. 20. Battery SOH estimation results with different fast charging protocols. (a) Cell 1 (b) Cell 4. (c) Cell 7. (d) Cell 17.

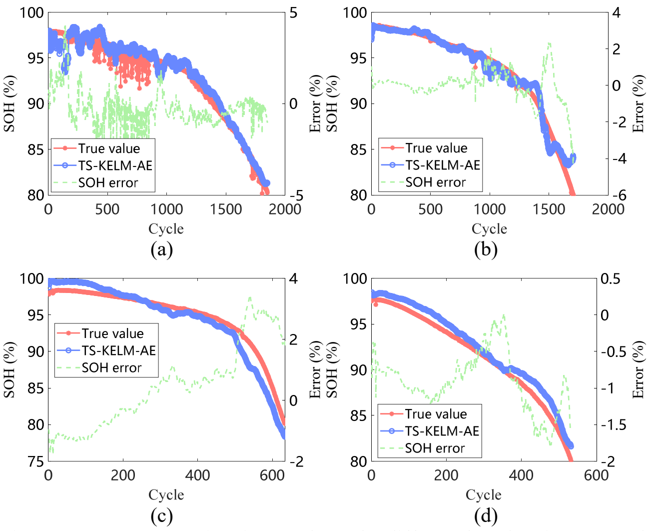

Fig. 21. Battery SOH estimation results under different fast charging protocols. (a) Cell 2 (b) Cell 5. (c) Cell 8. (d) Cell 18.

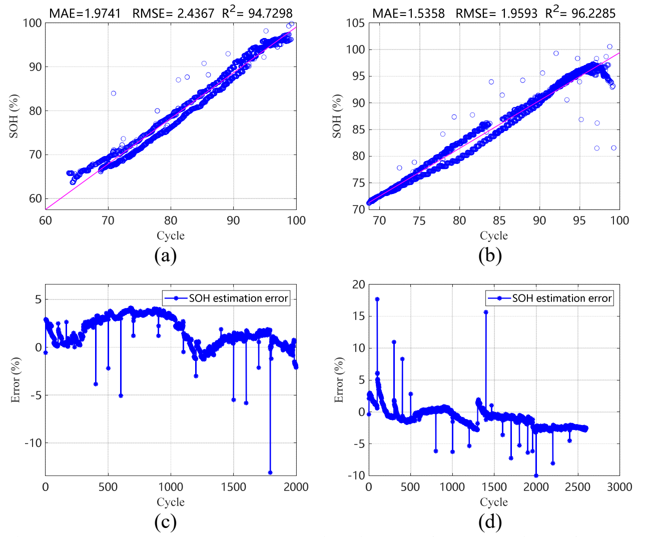

Fig. 22. Cross-temperature SOH estimation results. (a) estimated SOH of 45°C and 55°C (b) estimated SOH of 25°C and 35°C. (c) errors of 45°C and 55°C. (d) errors of 25°C and 35°C.

The corresponding SOH estimation trajectories are illustrated in Figs. 20 and 21, respectively, and the error metrics are summarized in Table IX. For the 3.6C(80%)–3.6C condition, both test directions yield low MAE and RMSE values, with R^2 exceeding 0.94, demonstrating stable transferability between the two cells. Similarly, under the 4C(80%)–4C and 5.4C(80%)–5.4C protocols, the cross-cell MAE remains low, and the R^2 values stay above 0.95 on average, indicating reliable performance even when evaluated on an unseen battery. Although the 4.8C(80%)–4.8C protocol exhibits more pronounced differences between the two cells, the proposed method still maintains competitive performance with an average R^2 above 0.90.

TABLE IX
CROSS-CELL GENERALIZATION PERFORMANCE UNDER DIFFERENT FAST CHARGING PROTOCOLS

Charging Protocol	Cell Index	MAE (%)	RMSE (%)	R^2 (%)
3.6C(80%)–3.6C	Cell 1	0.9136	1.0662	94.3902
	Cell 2	0.7339	0.9966	95.3734
	Average	**0.8238**	**1.0314**	**94.8819**
4C(80%)–4C	Cell 4	0.3903	0.5710	98.6391
	Cell 5	0.5681	0.8493	96.6848
	Average	**0.4792**	**0.7102**	**97.6619**
4.8C(80%)–4.8C	Cell 7	1.1401	1.4243	88.1130
	Cell 8	0.8555	1.3002	92.9208
	Average	**0.9978**	**1.3623**	**90.5169**
5.4C(80%)–5.4C	Cell 17	0.9155	0.9969	95.3938
	Cell 18	0.9087	0.9813	95.0661
	Average	**0.9121**	**0.9891**	**95.2298**

Furthermore, to evaluate the model's robustness under domain shift conditions, we also performed cross-temperature validation, where the model was trained on low-temperature data (25 °C and 35 °C) and tested on high-temperature data (45 °C and 55 °C), and vice versa, the SOH estimation results are illustrated in Fig. 22. The estimated and true SOH values are distributed near the diagonal, indicating that the model exhibits good generalization ability across unknown thermal conditions.

Overall, these experiments demonstrate that the proposed SOH estimation framework with fused LLM is not restricted to a single dataset or specific operating conditions. In contrast, it can learn degradation-related underlying features that effectively transfer across different batteries, charging rates, and temperature conditions. These results consistently confirm the strong cross-cell and cross-temperature generalizability of the proposed method.

G. Comparison with LLM-Based Feature Engineering and Prompt Sensitivity Analysis

Although LLM applications in battery SOH estimation are still in early stage, prior studies have already explored two preliminary LLM-driven strategies: (1) pre-trained large language models (PLMs) as generic knowledge sources for text-based reasoning or auxiliary feature interpretation; (2) LLM-based feature prompt engineering, where researchers design prompts (no-prompt, casual, and template forms in this work) to guide the model in generating feature descriptors related to battery degradation. As summarized in Table X, PLM-based models such as BERT-base, DeepSeek-R1, and TimeGPT provide a useful baseline, achieving RMSE values in the range of 0.21–1.06% without any manually defined features. However, their lack of explicit battery-specific knowledge limits their ability to consistently capture subtle aging signatures embedded in Q–V curves. This motivates the LLM-based feature prompt engineering paradigm, where the LLM is explicitly instructed to enumerate degradation-sensitive features. For ChatGPT-4, the No-prompt configuration yields an RMSE of 1.34%, and the extracted features are dominated by simple time-domain statistics. Introducing an informal ("casual") prompt reduces the RMSE to 0.96% and produces feature sets that begin to reflect degradation patterns. Satisfyingly, the structured template prompt achieves the best SOH estimation performance, with an RMSE of 0.72%, by providing explicit, domain-aware instructions that reduce semantic ambiguity and stabilize the feature generation behavior. Under the same template design, DeepSeek-V3 attains an RMSE of 0.88%, indicating that the overall prompt structure is transferable across strong LLMs, albeit with some performance gap relative to GPT-4.

Conceptually, the prompt plays the role of a task-specific interface that aligns the LLM's latent knowledge with our structured feature extraction objectives. When an alternative LLM is employed (e.g., DeepSeek-V3 instead of ChatGPT-4), the interaction between the prompt and the model may differ along two main dimensions: (i) semantic alignment indicates how well the instruction phrasing and examples of instructions align with the model's pre-training distribution and internal reasoning patterns; and (ii) output format adherence measures how strictly the model adheres to numerical and tabular format constraints when generating feature values from Q-V data [37], [38]. Both aspects directly influence the consistency, physical plausibility, and usability of the extracted features.

Overall, this analysis indicates that the prompt is a critical design element, determining how LLMs transform raw V-Q trajectories into high-level aging descriptors. By combining template-based prompt engineering with TS-KELM-AE, the proposed framework remains both effective and reasonably robust to the choice of underlying LLM, which is essential for future deployment scenarios where different organizations may favor different commercial or open-source LLMs.

TABLE X
COMPARISON OF LLM-DRIVEN STRATEGIES IN SOH ESTIMATION TASKS

Strategy	Algorithm	Prompt	RMSE (%)
Pre-trained large language model（PLM）	BERT-base[31]	--	0.54
	Deepseek-R1[32]	--	0.21
	TimeGPT[33]	--	1.06
LLM-based feature prompt engineering	FourierGNN[34]	--	0.86
	ChatGLM-6B[35]	--	1.04
	ChatGPT-4 [36]	No	1.34
	ChatGPT-4	Casual	0.96
	ChatGPT-4	Template	0.72
	DeepSeek-V3	Template	0.88

Note: Template prompt is shown in Fig. 5; and casual prompt is "This is battery data containing capacity and voltage values. Please create some features that may help analyze battery aging or health".

V. CONCLUSION

To improve the processing efficiency and utilization of battery aging data, this paper combines GPT-powered time-domain and frequency-domain feature extraction for the first time. To validate GPT's capability to capture high-order nonlinear degradation patterns and transform them into a more linearized and learnable feature space, we conducted SHAP-based feature importance analysis and progressive feature elimination experiments under different operating conditions. Results demonstrate that GPT-powered aging features consistently highlight capacity-related factors while mitigating the detrimental sign reversal behavior of voltage-related features at elevated temperatures, particularly spectral mean frequency and spectral entropy, confirming GPT's advantage in extracting robust degradation descriptors. Based on these features, we developed a two-stage kernel extreme learning machine autoencoder (TS-KELM-AE) model, which possesses robust nonlinear mapping and deep representation abilities. Future research will leverage knowledge distillation methods to transfer reasoning capabilities from open-source or locally hosted large language models (e.g., Llama and the Gemini series) to lightweight on-device models, and combine them with parameter-efficient fine-tuning techniques to reduce computational costs and enhance data privacy protection.

REFERENCES

[1] Z. Wei, K. Liu, X. Liu, et al., "Multi-level data-driven battery management: From internal sensing to big data utilization," *IEEE Trans. Transport. Electrific.*, vol. 9, no. 4, pp. 4805-4823, Dec. 2023.

[2] A. Tomaszewska et al., "Lithium-ion battery fast charging: A review," *eTransp.*, vol. 1, Art. no. 100011, Aug. 2019.

[3] Y. Liu, Y. Zhu, Y. Cui, "Challenges and opportunities towards fast-charging battery materials," *Nat. Energy*, vol. 4, no. 7, pp. 540-550, Jul.2019.

[4] C. She, Y. Shen, G. Bin, J. Hong, Y. Peng, "Accurate state of health estimation of battery system based on multistage constant current charging and behavior analysis in real-world electric vehicles," *IEEE Trans. Transport. Electrific.*, vol. 11, no. 1, pp. 97-108, Feb. 2025.

[5] A. G. Olabi, Q. Abbas, P. A. Shinde, et al., "Rechargeable batteries: Technological advancement, challenges, current and emerging applications," *Energy*, vol. 266, Art. no. 126408, Mar.2023.

[6] S. Ma, M. Jiang, P. Tao, C. Song, J. Wu, J. Wang, et al., "Temperature effect and thermal impact in lithium-ion batteries: A review," *Prog. Nat. Sci. Mat. Int*., vol. 28, no. 6, pp. 653-666, Dec. 2018.

[7] Y. Wang, X. Zhang, K. Li, et al., "Perspectives and challenges for future lithium-ion battery control and management," *eTransp.*, vol. 18, Art. no. 100260, Oct. 2023.

[8] Y. Che, X. Hu, X. Lin, J. Guo, and R. Teodorescu, "Health prognostics for lithium-ion batteries: Mechanisms, methods, and prospects," *Energy Environ. Sci*., vol. 16, pp. 338-371, 2023.

[9] W. Hu and Q. Qian, "Lithium-ion battery state of health and failure analysis with mixture Weibull and equivalent circuit model," *iScience*, vol. 27, no. 6, Jun.2024.

[10] Y. Gao, K. Liu, C. Zhu, X. Zhang, and D. Zhang, "Co-estimation of state-of-charge and state-of-health for lithium-ion batteries using an enhanced electrochemical model," *IEEE Trans. Ind. Electron*., vol. 69, no. 3, pp. 2684-2696, Mar. 2022.

[11] J. Qiao, S. Wang, C. Yu, et al., "A chaotic firefly-particle filtering method of dynamic migration modeling for the state-of-charge and state-of-health co-estimation of a lithium-ion battery performance," *Energy*, vol. 263, Art. no. 126164, Jan.2023.

[12] Z. Xue, Y. Zhang, C. Cheng, et al., "Remaining useful life prediction of lithium-ion batteries with adaptive unscented Kalman filter and optimized support vector regression," *Neurocomput.*, vol. 376, pp. 95-102, Sep. 2020.

[13] X. Li, C. Yuan, Z. Wang, et al., "Lithium battery state-of-health estimation and remaining useful lifetime prediction based on non-parametric aging model and particle filter algorithm," *eTransp.*, vol. 11, Art. no. 100156, Feb. 2022.

[14] C. She, Y. Li, C. Zou, et al., "Offline and online blended machine learning for lithium-ion battery health state estimation, " *IEEE Trans. Transport. Electrific*., vol. 8, no. 2, pp. 1604-1618, Aug. 2022.

[15] C. Lin, J. Xu, J. Hou, et al., "Ensemble method with heterogeneous models for battery state-of-health estimation," *IEEE Trans. Ind. Informat.,* vol. 19, no. 10, pp. 10160-10169, Oct. 2023.

[16] S. Ji, J. Zhu, Z. Zhang, et al., "Degradation prognosis for fast-charging batteries via improved domain adaptation," *IEEE Trans. Ind. Informat.*, vol. 21, no. 2, pp. 1734-1743, Feb. 2025.

[17] D. Zhang et al., "Multistep fast charging-based state of health estimation of lithium-ion batteries," *IEEE Trans. Transport. Electrific.*, vol. 10, no. 3, pp. 4640-4652, Sep. 2024.

[18] J. Obregon, Y. R. Han, C. W. Ho, et al., "Convolutional autoencoder-based SOH estimation of lithium-ion batteries using electrochemical impedance spectroscopy," *J. Energy Storage*, vol. 60, Art. no. 106680, Apr.2023.

[19] J. Lee and J. H. Lee, "Simultaneous extraction of intra- and inter-cycle features for predicting lithium-ion battery's knees using convolutional and recurrent neural networks," *Appl. Energy*, vol. 356, p. 122399, Feb. 2024.

[20] C. Bian, X. Han, Z. Duan, C. Deng, S. Yang, and J. Feng, "Hybrid prompt-driven large language model for robust state-of-charge estimation of multitype Li-ion batteries," *IEEE Trans. Transport. Electrific.*, vol. 11, no. 1, pp. 426-437, Feb. 2025.

[21] Z. Zheng, Z. He, O. Khattab, et al., "Image and data mining in reticular chemistry powered by GPT-4V," *Digital Discov.*, vol. 3, no. 3, pp. 491-501, 2024.
[22] D. Ge, G. Jin, J. Wang, et al., "A novel data-driven IBA-ELM model for SOH/SOC estimation of lithium-ion batteries," *Energy*, vol. 305, Art. no. 132395, Oct. 2024.
[23] B. Gou, Y. Xu, and X. Feng, "An ensemble learning-based data-driven method for online state-of-health estimation of lithium-ion batteries," *IEEE Trans. Transport. Electrific.*, vol. 7, no. 2, pp. 422-436, Jun. 2021.
[24] Y. Zhang, T. Wik, J. Bergström, and C. Zou, "State of Health Estimation for Lithium-Ion Batteries Under Arbitrary Usage Using Data-Driven Multimodel Fusion," *IEEE Trans. Transport. Electrific.*, vol. 10, no. 1, pp. 1494-1507, Mar. 2024
[25] B. Zraibi, C. Okar, H. Chaoui, M. Mansouri, "Remaining useful life assessment for lithium-ion batteries using CNN-LSTM-DNN hybrid method," *IEEE Trans. Veh. Technol.*, vol. 70, no. 5, pp. 4252-4261, May 2021.
[26] K. Luo, H. Zheng, and Z. Shi, "A simple feature extraction method for estimating the whole life cycle state of health of lithium-ion batteries using transformer-based neural network," *J. Power Sources*, vol. 576, Art. no. 233139, Aug.2023.
[27] K. Liu, Y. Liu, S. Zhao, X. Li, and Q. Peng, "An ultrasonic wave-based method for efficient state-of-health estimation of Li-ion batteries," *IEEE Trans. Ind. Electron.*, vol. 72, no. 4, pp. 3792-3802, Apr. 2025.
[28] S. Tao, M. Zhang, Z. Zhao, et al., "Non-destructive degradation pattern decoupling for early battery trajectory prediction via physics-informed learning," *Energy Environ. Sci.*, vol. 18, pp. 1544-1559, 2025.
[29] K. A. Severson, "Data-driven prediction of battery cycle life before capacity degradation," *Nat. Energy*, vol. 4, no. 5, pp. 383–391, 2019.
[30] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 1–10.
[31] Z. Zhang, Y. Zhu, Q. Zhang, et al., "Multi-cycle charging information guided state of health estimation for lithium-ion batteries based on pre-trained large language model," *Energy*, vol. 313, Art. no. 133993, Jan. 2024.
[32] Q. Zhao, Y. Li, L. Kong, L. Chang, G. Kou, and M. Mu, "Multi-feature battery state of health linear estimation based on low-cost pretrained inference LLM," *IEEE Trans. Transport. Electrific*., 2025.
[33] W. Sun, C. Wu, C. Xie, et al., "Fine-tuning enables state of health estimation for lithium-ion batteries via a time series foundation model," *Energy*, vol. 318, Art. no. 134177, Jan. 2025.
[34] R. Lai, Y. Tian, J. Wang, et al., "Battery charging curve prediction via Fourier graph neural network fused prompt features," *Energy*, Art. no. 137565.
[35] Z. Du, Y. Qian, X. Liu, et al., "GLM: General language model pretraining with autoregressive blank infilling," in *Proc. 60th Annu. Meeting Assoc. Comput. Linguistics*, 2022, pp. 320–335.
[36] S. Tuncel, H. Cinar, M. Gucyetmez, and N. Erdogan, "Large language model-based methodology for data-driven health prediction of lithium-ion batteries," *Energy Rep.*, vol. 13, pp. 6425–6434, 2025.
[37] T. Brown, B. Mann, N. Ryder, et al., "Language models are few-shot learners," in *Adv. Neural Inf. Process. Syst.*, vol. 33, pp. 1877–1901, 2020.
[38] L. Ouyang, J. Wu, X. Jiang, et al., "Training language models to follow instructions with human feedback," in *Adv. Neural Inf. Process. Syst.*, vol. 35, pp. 27730–27744, 2022.

Laijin Luo received the M.S. degree in information and communication Engineering from Anqing Normal University, Anhui, China, in 2024.

He is currently pursuing the Ph.D. degree with the School of Electrical Engineering, Chongqing University, Chongqing, China. His research interests include battery modeling, state estimation and deep learning.

Yu Wang (Senior Member, IEEE) received B.Eng. degree in Electrical Engineering and Automation from Wuhan University, China in 2011, M.Sc. and Ph.D. degree in Power Engineering from Nanyang Technological University (NTU), Singapore in 2012 and 2017, respectively.

Currently, he is a professor at the School of Electrical Engineering, Chongqing University, China. He was a Marie Skłodowska-Curie Individual Fellow at Control and Power Group, Imperial College London.

Dr. Wang serves as the Associated Editor of *Journal of Modern Power Systems and Clean Energy*, *IET Generation, Transmission and Distribution*, and *IET Renewable Power Generation*. His research interests include microgrid operation and control, power system dynamics and control, and cyber-physical power systems.

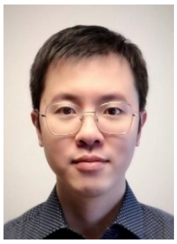

Yang Li (Senior Member, IEEE) received the B.E. degree in electrical engineering from Wuhan University, Wuhan, China, in 2007, and the M.Sc. and Ph.D. degrees in power engineering from Nanyang Technological University (NTU), Singapore, in 2008 and 2015, respectively.

He was with the Energy Research Institute, NTU, the School of Electrical Engineering and Computer Science, Queensland University of Technology, Brisbane, QLD, Australia, and the School of Automation, Wuhan University of Technology, Wuhan, the Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden, from 2016 to 2025. Since 2025, he has been a Professor with the School of Electrical Engineering and Automation, Wuhan University. His research interests include modeling, control, and management of energy storage systems in the power and transport sectors.

Dr. Li serves as an Associate Editor for several journals, such as IEEE Transactions on Industrial Electronics, IEEE Transactions on Transportation Electrification, and IEEE Transactions on Energy Conversion.

Changfu Zou (Senior Member, IEEE) received the Ph.D. degree in automation and control engineering from the University of Melbourne, Melbourne, VIC, Australia, in 2017.

He was a Visiting Student Researcher at the University of California, Berkeley, Berkeley, CA, USA, from 2015 to 2016. He has been with the Automatic Control Group, Chalmers University of Technology, Gothenburg, Sweden, since 2017, progressing from a Researcher to an Assistant Professor, then an Associate Professor, and is currently a Professor. His research interests include advanced modeling and automatic control of energy storage systems, particularly batteries.

Dr. Zou has been awarded several prestigious grants from the European Commission and Swedish national agencies, including the Swedish Research Council Starting Grant and Project Grant. He serves as an Associate Editor for several journals, such as IEEE Transactions on Vehicular Technology and IEEE Transactions on Transportation Electrification.