Proceedings of CONF-SPML 2026 Symposium: The 2nd Neural Computing and Applications Workshop 2025
DOI: 10.54254/2755-2721/2026.TJ30413

Hybrid CNN-LSTM Framework for Lithium-Ion Battery Degradation and Remaining Useful Life Prediction

Canyang Yao

Jiangsu Tianyi High School, Wuxi, China
yao-harry@163.com

Abstract. Lithium-ion battery reliability is dependent upon an accurate prediction of battery remaining useful life (RUL). There are several challenges to exact prediction such as the nonlinear degradation, constant measurement noise, etc. That standalone models are unable to provide a full answer. For example, Convolutional Neural Network (CNN) can capture local temporal features but often fail to take care of the long-range dependencies, and Long Short-Term Memory (LSTM) can model the development of sequences while they usually miss the significant, detailed structures. This paper introduces a deep hybrid architecture composed of CNNs, LSTMs and a multi-head attention mechanism that overcomes these complementary limitations. The resulting system learns to extract local features through the CNN, to notice long-term trends through the LSTM and, most importantly, to adaptively recalibrate the feature's relevance through attention over time steps. On the normalized sliding-window sequences of NASA battery dataset, this paper proposes a hybrid model evaluated using both Leave-One-Battery-Out (LOBO) and Temporal Split experiments that has been shown to exhibit a significant improvement in its prediction accuracy and robustness, relative to its individual components, allowing us to focus our attention towards only the important degradation phases with which to perform a better RUL estimation.

Keywords: Hybrid CNN-LSTM, RUL prediction, Multi-head self-attention, Battery

1. Introduction

The electric vehicle's fast growing rate, aerospace and portable electric devices has drive a new need for reliable lithium-ion batteries, which will be the key to overcome the advanced challenges: assurance of safety, control of cost and assurance of system dependability [1-4]. This necessitates an accurate modeling and prediction of a battery's health status and the battery's remaining useful life (RUL). Researchers still remain with a large task to accomplish this [5,6]. Data-driven methods have been put forward as alternatives to physics-based models [7-9].

Regarding CNN-based approaches, which excel at local-feature and spatial pattern recognition, numerous models have been developed [10-16]. Zheng et al. used CNN-GRU hybrid model to compute the battery SOH under random charge condition [10]. Li et al. proposed Multi-channel CNN-Times Net framework, the predication errors are smaller than 1.5% in several datasets [11]. Xia et al. combined CNN with BiGRU to handle degradation nonlinearity and noise, while Ma et al. adopted a two-stage CNN-GPR ensemble to improve early prediction [12,13].

DOI: 10.54254/2755-2721/2026.TJ30413

While CNNs are capable of extracting local features, they struggle to capture long-term dependencies in the temporal dimension. In contrast, LSTM-based approaches are specifically designed to model these long-term temporal dependencies. Ma et al. applied an optimized LSTM along with health indicator selection to voltage, current and temperature, which attain high SOH estimation performance on NASA and MIT datasets [17]. Sun et al. obtained similar results. The fusion of feature engineering and deep learning was showed to be beneficial by integrating Incremental Capacity Analysis (ICA) with bidirectional LSTM [18].

While LSTM captures long-term dependencies effectively, it exhibits limited sensitivity to local features. To address this limitation, this paper proposes a hybrid CNN-LSTM model with multi-head attention for battery degradation prediction. This study has not only technical value but also real world implications for making the battery management systems more reliable and safe.

2. Problem definition and research scope

As shown in Fig. 1, the health management system of the lithium-ion battery works in 2 stages. First, an offline computational phase using a CNN-LSTM hybrid neural network to train a high precision RUL prediction model on previous cycling history of the battery and then deploy it. The second stage is online computing phase in which the present runtime information is being processed by the previously described pipeline and analyzed by the deployed pipeline. Then the real-time operational data is processed by the deployed model.

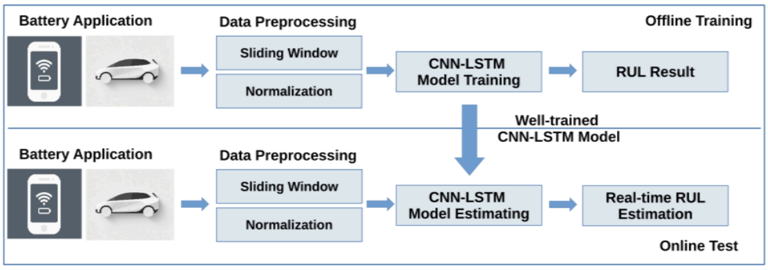

Figure 1. The framework of lithium-ion battery health management system

The non-linear degradation trend of lithium-ion batteries operating data renders the prediction of lithiumion battery RUL difficult. We approach this problem through a data-driven regression framework. Our model takes multivariate time-series data as input, which captures voltage, current, and temperature at a 1Hz sampling rate from standard battery cycles. Its aims to map these profiles directly to the remaining capacity per cycle and to predict the number of cycles left until capacity falls to 70% of the 2Ah nominal value. We use data from the NASA Prognostics Center of Excellence dataset, specifically cells B0005, B0006, B0007, and B0018. This selection ensures a consistent experimental setup, with all cells featuring a 18650 $LiCoO_2$ chemistry and being cycled at a 24°C ambient temperature.

3. CNN-LSTM model design

This paper proposes a hybrid CNN-LSTM model augmented with a self-attention mechanism. The CNN captures local characteristics in the capacity degradation data, while the LSTM models the long-term trends. The self-attention mechanism then identifies and highlights the importance of key time points for final RUL prediction. The network framework is shown in Fig. 2. The framework consists of input data, preprocessing data, the CNN model and the LSTM model.

DOI: 10.54254/2755-2721/2026.TJ30413

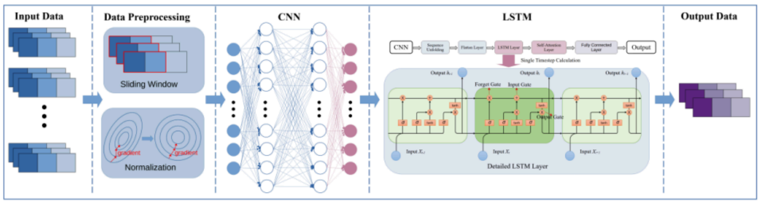

Figure 2. The framework of CNN-LSTM model

3.1. Input layer and data preprocessing

The model input is preprocessed as battery capacity time series. Specifically, a sliding window approach is employed to construct supervised learning samples. Denote the time step as kim (set to 2), meaning two consecutive historical capacity values (C_t, C_{t+1}) are used as input features to predict the capacity value C_{t+2} at a prediction horizon. Both training and testing set data are normalized to the interval [0, 1] to accelerate model convergence and improve generalization performance. Finally, the input data is reshaped into a 4-dimensional tensor format to suit subsequent convolutional operations.

3.2. Feature extraction layer (CNN)

One of the core innovations of the hybrid prediction model constructed in this paper lies in the introduction of a one-dimensional CNN as a front-end feature extractor. The CNN is specifically designed to automatically learn and extract discriminative local temporal features from the degradation sequence data of lithium-ion batteries. The attenuation of battery capacity is not a purely random process. Its variation patterns contain locally correlated patterns driven by the degradation of internal electrochemical characteristics. Traditional feature engineering methods heavily rely on domain prior knowledge and struggle to capture these complex, non-linear, deep-seated patterns. The introduction of the CNN structure effectively addresses this issue.

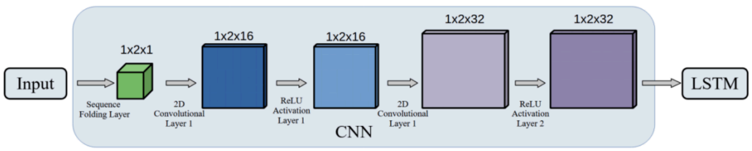

Figure 3. The framework of CNN

As shown in Fig. 3, the CNN section of the model consists of the following layers in sequence:

Sequence Folding Layer folds the input time series samples along the time dimension, converting them into a two-dimensional structure, thereby simulating the height and width of an image. This prepares the data for the subsequent application of standard two-dimensional convolutional layers.

The model stacks two 2D convolutional layers. The kernel size is strategically set to [3, 1]. The first convolutional layer uses 16 filters, aiming to extract 16 basic primary local feature patterns from the raw sequence (such as specific segments of capacity decline rate). The second convolutional layer uses 32 filters, which performs deeper combination and abstraction based on the primary features to capture more complex and advanced temporal features.

A ReLU activation function, i.e. f(x)= max(0, x), is applied after every convolutional layer, which generates non-linear transformation. This enables the CNN to learn mappings between the input data and the target. The ReLU activation function allows the model to fit the non-linear dynamics of

DOI: 10.54254/2755-2721/2026.TJ30413

battery degradation. This layer transform raw sequences $X=[x_1, x_2, \ldots, x_T]$ into feature maps through learnable kernels. The specific formula is as follows:

$$h_j^{(l)}=\mathrm{ReLU}\left(\sum_{k=1}^{K_{l-1}} W_{jk}^{(l)} * h_k^{(l-1)} + b_j^{(l)}\right) \tag{1}$$

Where $W^{(l)} \in \mathbb{R}^{3\times 1}$ implements temporal convolution matching the [3,1] kernel configuration.

3.3. Sequence modeling and attention layer (LSTM + attention)

In this section, this study proposes a hybrid model incorporating an LSTM network and a self-attention mechanism. The hybrid model extracts features from the CNN module. Its purpose is to capture the long-term temporal dependencies within the battery degradation process. The design addresses two critical challenges: learning the non-linear trends along the cycles causing the capacity fading, and adjusting to the decisive time steps affecting the RUL estimates. As shown in Fig. 4, the LSTM's gating mechanism (including the forget, input, and output gates) selectively retains critical degradation signatures. This design also mitigates the gradient vanishing/exploding problems inherent in standard RNNs. Afterwards, this study applies a multi-head self-attention layer to further fine-tune feature importance, making it more resistant to noise-induced perturbation.

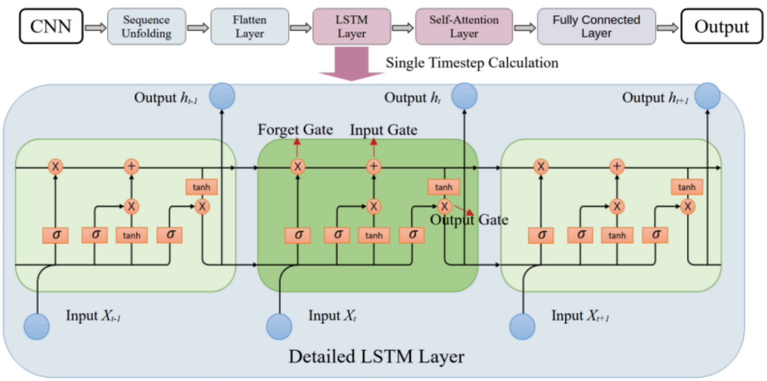

Figure 4. The framework of LSTM

3.3.1. Theory and analysis of the internal gating mechanism in the LSTM layer

The key component in this model is the LSTM network that learns the long-term dependency of battery capacity degradation. It possesses an advanced gating mechanism capable of addressing the vanishing/exploding gradient problems for conventional RNNs and, therefore, permits the learning and memorizing the major patterns for a long time.

As shown in Fig. 4, the LSTM unit's core is Cell State, and there are 3 gate structures which include the Forget Gate, the Input Gate, and the Output Gate. All gate structures use the Sigmoid activation function (output range 0 to 1) to implement the “gate” function, so as to control the flow of information.The forward propagation process and formulas are as follows:

The Forget Gate can learn to forget early historical information irrelevant to the current capacity degradation (such as stable capacity during initial cycles) and focus on recent decay trends. The forget gate decides what information to discard from the previous cell state C_{t-1} . It looks at the current input x_t and the previous hidden state h_{t-1} , and outputs a value between 0 and 1 (via the

DOI: 10.54254/2755-2721/2026.TJ30413

Sigmoid function) to the cell state C_{t-1} . A value closer to 0 means "completely forget," while a value closer to 1 means "completely retain".

The Input Gate allows the model to selectively memorize new patterns present in the current input x_t (i.e., the latest features extracted by the CNN) that are valuable for predicting future capacity. The input gate decides which new information will be stored in the cell state. It consists of two parts: first, a Sigmoid layer (the input gate) decides which values to update. Second, a Tanh layer creates a new candidate value vector C_{ds} which might be added to the cell state.

The Output Gate determines the final output based on the current cell state C_t . First, a Sigmoid layer decides which parts of the cell state will be output. Then, the cell state is processed through the Tanh function (to obtain a value between -1 and 1), and this result is multiplied by the output of the Sigmoid gate, yielding the final hidden state output h_t . The h_t serves as the "short-term memory" or the hidden output at the current time step, containing both the essence of the long-term memory C_t and reflecting the influence of the current input. In this model's configuration (OutputMode = 'last'), only the hidden state h_t from the last time step is output. It is designed to encode the condensed contextual information of the entire input sequence.

The LSTM layer processes CNN-extracted features to model long-term dependencies across charge-discharge cycles. The gating architecture are described as follows:

$$C_t = f_t \odot C_{t-1} + i_t \odot \widetilde{C_t} \tag{2}$$

$$h_t = o_t \odot \tanh(C_t) \tag{3}$$

In case $|\Delta C_t| < \epsilon$, the forget gate ($f_t = \sigma(\bullet)$) excluded unnecessary historical information actively. The input gate i_t emphasizes learning meaningful degradation feature. C_t as a long time electrochemical memory in the cell state could handle the gradient vanishing problem frequently existing in traditional RNNs.

3.3.2. Self-attention layer

The feature vector from the LSTM context h_t is passed through several self-attention networks in order to extract dependencies between the feature vector channels, dynamically re-weighting the relative importance of each feature channel. This process makes the model focus on salient failure modes, and hence achieves higher accurate RUL prediction. This functionality is implemented as follows:

$$\text{Attention} = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \tag{4}$$

$$z = W^O \bullet [\text{head}_1 \oplus \text{head}_1] \tag{5}$$

where queries Q , keys K , and values V are linearly projected from the LSTM output. The softmax weights identify pivotal degradation phases, with dual attention heads capturing complementary temporal contexts.

DOI: 10.54254/2755-2721/2026.TJ30413

4. Experiments

This study evaluates the proposed CNN-LSTM model on the NASA dataset with two experimental designs to assess its generalization capability. First, the LOBO cross-validation tests generalization to unseen battery units by training on multiple cells and testing on a held-out one. Second, a Temporal Split experiment, using the first 70% of cycles for training and the last 30% for testing, evaluates the model's ability to predict long-term degradation of a known cell. Performance is measured using MAE, MSE, RMSE, RPD, and MBE, with results benchmarked against CNN-only and LSTM-only baselines under identical conditions.

4.1. Leave-One-Battery-Out (LOBO) experiments

To rigorously assess the generalisation capability of the proposed hybrid CNN+LSTM architecture, a LOBO cross-validation protocol was implemented. In each fold, data from three cells (e.g., B0006, B0007 and B0018) were used for training, while the remaining cell (e.g., B0005) was retained as a completely unseen test set. As shown in Fig. 5, subfigures (a), (b), and (c) present the training-set results for cells B0006, B0007, and B0018, respectively, while subfigure (d) displays the test-set data for cell B0005.

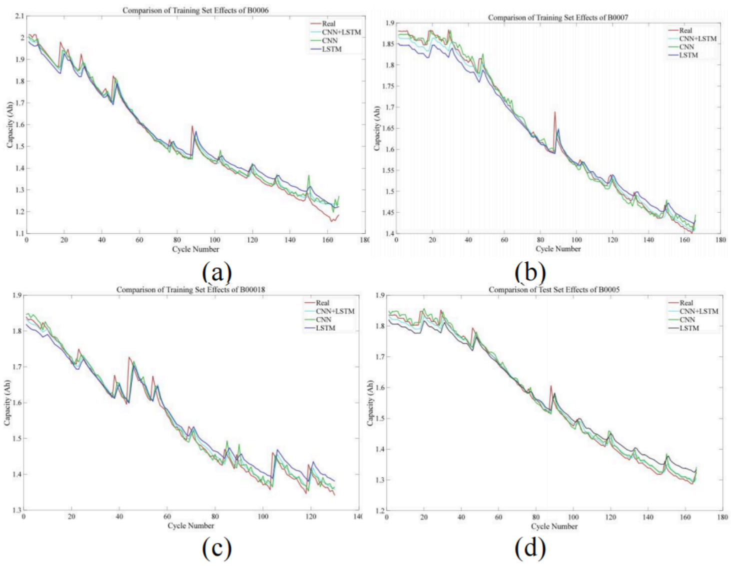

Figure 5. The experimental results of the LOBO experiment on the training and testing sets

A quantitative analysis of these outcomes is provided in Table 1. The hybrid CNN-LSTM model demonstrates superior performance across all metrics, drastically reducing RMSE, MAE, and MSE by approximately 96% compared to the single-branch networks (e.g., RMSE of 0.046 Ah vs. ≈1.09 Ah). Its near-zero MBE and higher RPD (4.31) confirm an unbiased and more predictive fit, contrasting with the standalone models' tendency to regress to a naive historical average. These results collectively affirm the hybrid architecture's capability to accurately capture the battery's temporal degradation.

DOI: 10.54254/2755-2721/2026.TJ30413

Table 1. Data analysis of LOBO experimental results

	MAE	MSE	RMSE	RPD	MBE
CNN+LSTM train	0.010431	0.00023071	0.015189	12.44	0.0010599
CNN+LSTM test	0.018015	0.00215	0.046368	4.3115	-0.0010707
CNN train	1.0912	1.192	1.0918	6.1774	1.0912
CNN test	1.0856	1.182	1.0872	3.9751	1.0856
LSTM train	1.0914	1.1942	1.0928	4.0066	1.0914
LSTM test	1.0854	1.1836	1.0879	3.1449	1.0854

4.2. Temporal Split experiment

To evaluate the effectiveness of the proposed LSTM+CNN hybrid model for lithium-ion battery life prediction, we adopt a Temporal Split strategy. Specifically, the first 70% of cycles from each cell are used for training, while the remaining 30% are reserved for testing. This mimics real-world deployment, where only historical data are available and future measurements are unknown. The results for both the training and test sets are displayed in every sub-figure of Fig. 6. It can be seen that—apart from using LSTM alone—the other two approaches achieve strong performance.

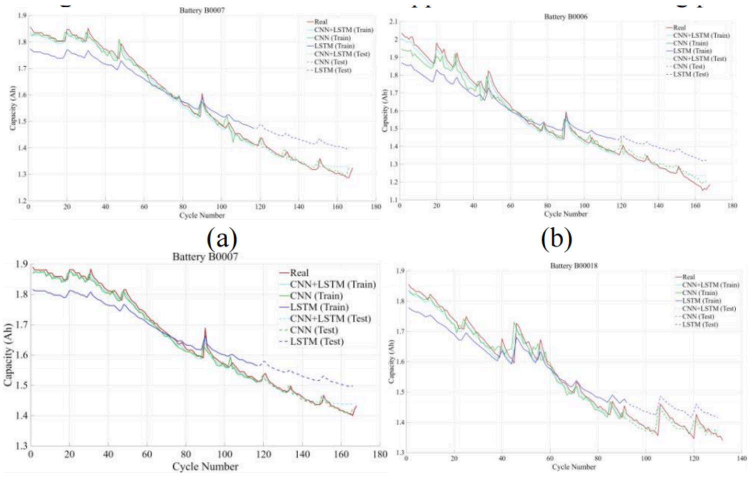

(a) (b)

(c) (d)

Figure 6. The experimental results of the Temporal Split experiment on the training and testing sets

As shown in Table 2, the proposed CNN-LSTM hybrid demonstrates superior generalization with the lowest test RMSE (≈0.045), highest RPD (≈4.42), and negligible bias (MBE≈+0.0038). In contrast, the pure CNN shows high training accuracy but significant overfitting, while the standalone LSTM performs poorest. The hybrid architecture thus achieves the best performance by balancing accuracy and robustness.

Table 2. Data analysis of Temporal Split experimental results

	MAE	MSE	RMSE	RPD	MBE
CNN+LSTM train	0.020162	0.0030088	0.054852	3.5884	-0.00078551
CNN+LSTM test	0.020471	0.0020656	0.045448	4.419	0.0037696
CNN train	0.018915	0.0019086	0.043688	4.5053	0.00053011
CNN test	0.022825	0.003382	0.058155	3.4683	0.0072082
LSTM train	0.053794	0.0036856	0.061764	2.498	0.012386
LSTM test	0.075675	0.0028631	0.078391	1.212	0.074953

5. Conclusion

This research tackles a persistent problem in predicting battery lifespan: the complex, non-linear nature of battery degradation. This study proposes a hybrid model that combines CNN, LSTM and a multi-head self-attention mechanism. To prove the efficacy of proposed model, the model is evaluated to check its performance. This method is more accurate than the CNN or LSTM models on the both experiments Leave-One-Batch-Out (LOBO) and Temporal Split.

References

[1] Zhao, B., Zhang, W., Zhang, Y., Zhang, C., Zhang, C., & Zhang, J. (2024). Lithium-ion battery remaining useful life prediction based on interpretable deep learning and network parameter optimization. Applied Energy, 379, 124713. https: //doi.org/10.1016/j.apenergy.2024.124713

[2] Wang, S., Wang, P., Wang, L., Li, K., Xie, H., & Jiang, F. (2024). An enhanced deep learning framework for state of health and remaining useful life prediction of lithium-ion battery based on discharge fragments. Journal of Energy Storage, 107, 114952. https: //doi.org/10.1016/j.est.2024.114952

[3] Wang, C., Huang, Z., He, C., Lin, X., Li, C., & Huang, J. (2024). Research on remaining useful life prediction method for lithium-ion battery based on improved GA-ACO-BPNN optimization algorithm. Sustainable Energy Technologies and Assessments, 73, 104142. https: //doi.org/10.1016/j.seta.2024.104142

[4] Liu, L., Huang, J., Zhao, H., Li, T., & Li, B. (2025). PatchFormer: A novel patch-based transformer for accurate remaining useful life prediction of lithium-ion batteries. Journal of Power Sources, 631, 236187. https: //doi.org/10.1016/j.jpowsour.2025.236187

[5] Wang, X.-T., Zhang, S.-B., Wang, J.-S., Liu, X., Sun, Y.-C., Shang-Guan, Y.-P., & Zhang, Z.-Z. (2024). Ensemble learning prediction model for lithium-ion battery remaining useful life based on embedded feature selection. Applied Soft Computing, 169, 112638. https: //doi.org/10.1016/j.asoc.2024.112638

[6] Yang, M., Hu, X., Liu, Y., Zhang, H., & Li, S. E. (2024). An on-device personalized charging strategy with an aging model for lithium-ion batteries using deep reinforcement learning. IEEE Transactions on Industrial Informatics, 20(2), 1435–1446. https: //doi.org/10.1109/TII.2023.3312305

[7] Zhang, W., He, H., Zhang, X., Sun, L., & Li, J. (2024). Rapid assessing cycle life and degradation trajectory based on transfer learning for lithium-ion battery. IEEE Transactions on Industrial Informatics, 20(7), 9422–9432. https: //doi.org/10.1109/TII.2023.3340359

[8] Sundaresan, S., Pillai, P., Pattipati, K. R., & Balasingam, B. (2024). Theoretical performance bound and its experimental validation of battery capacity estimates in rechargeable batteries. IEEE Transactions on Industrial Electronics, 71(1), 669–679. https: //doi.org/10.1109/TIE.2023.3255953

[9] Deng, Z., Li, S., Zhang, L., & Xu, L. (2024). Attention-based deep learning method for lithium-ion battery remaining useful life prediction under variable loading profiles. Applied Energy, 354, 121896. https: //doi.org/10.1016/j.apenergy.2023.121896

[10] Zheng, Y., Hu, J., Chen, J., Deng, H., & Hu, W. (2023). State of health estimation for lithium battery random charging process based on CNN-GRU method. Energy Reports, 9, 1–10. https: //doi.org/10.1016/j.egyr.2022.12.093

[11] Li, Y., Qin, X., Chai, M., Wu, H., Zhang, F., Jiang, F., & Wen, C. (2025). SOH evaluation and RUL estimation of lithium-ion batteries based on MC-CNN-TimesNet model. Reliability Engineering & System Safety, 261, 111125.

DOI: 10.54254/2755-2721/2026.TJ30413

https: //doi.org/10.1016/j.ress.2025.111125
[12] Xia, W., Xu, J., Liu, B., et al. (2024). A novel denoising autoencoder hybrid network for remaining useful life estimation of lithium-ion batteries. Energy Science & Engineering, 12(8). https: //doi.org/10.1002/ese3.1823
[13] Ma, G., Wang, Z., Liu, W., et al. (2022). A two-stage integrated method for early prediction of remaining useful life of lithium-ion batteries. Knowledge-Based Systems, 259, 110012. https: //doi.org/10.1016/j.knosys.2022.110012
[14] Safavi, V., Vaniar, A. M., Bazmohammadi, N., et al. (2024). Early prediction of battery remaining useful life using CNN-XGBoost model and Coati optimization algorithm. Journal of Energy Storage, 98. https: //doi.org/10.1016/j.est.2024.113176
[15] Zhou, K., & Zhang, Z. (2024). Remaining useful life prediction of lithium-ion batteries based on data denoising and improved transformer. Journal of Energy Storage, 100. https: //doi.org/10.1016/j.est.2024.113749
[16] Ding, P., Liu, X., Li, H., et al. (2021). Useful life prediction based on wavelet packet decomposition and two-dimensional convolutional neural network for lithium-ion batteries. Renewable and Sustainable Energy Reviews, 148. https: //doi.org/10.1016/j.rser.2021.111287
[17] Ma, Y., Shan, C., Gao, J., & Chen, H. (2022). A novel method for state of health estimation of lithium-ion batteries based on improved LSTM and health indicators extraction. Energy, 251, 123973. https: //doi.org/10.1016/j.energy.2022.123973
[18] Sun, H., Sun, J., Zhao, K., Wang, L., & Wang, K. (2022). Data-driven ICA-Bi-LSTM-combined lithium battery SOH estimation. Mathematical Problems in Engineering, 2022, Article ID 9645892, 1–8. https: //doi.org/10.1155/2022/9645892.