LiB-TRAD: A Lithium Battery Thermal Runaway Acoustic Dataset for Anomaly Detection

Wang Xiaoliang[1], *Ming Ao*[1], *Chen Meixin*[1], *Jin Hao*[1]

[1]Zhejiang University, Hangzhou, China

***Abstract*— Efficient detection of lithium battery thermal runaway is a critical factor in promoting the large-scale application of lithium batteries in energy storage and electric transportation. Traditional methods rely heavily on contact-based techniques such as temperature, current, voltage, impedance, or structural deformation monitoring, which have limitations in terms of cost, real-time performance, and scalability. In contrast, acoustic detection, with its non-contact nature, low cost, and suitability for large-scale monitoring, is emerging as a promising alternative. While previous studies have demonstrated the effectiveness of machine learning-based acoustic methods for thermal runaway detection, there is still a lack of an open acoustic dataset covering the entire process of lithium battery thermal runaway. To address this, this paper introduces the first lithium battery acoustic dataset containing both normal and thermal runaway events, annotated with abnormal events. We further evaluate several baseline models and state-of-the-art acoustic event detection models using this dataset. Experimental results show that this dataset holds strong potential for thermal runaway anomaly detection and provides a valuable data foundation and benchmark for future research.**

***Index Terms*—Thermal Runaway, Acoustic Anomaly Detection, Lithium Battery Safety, Audio Dataset**

1. INTRODUCTION

With the development of renewable energy sources such as solar and wind power [1–2], large-scale energy storage stations and grid facilities have advanced significantly in recent years, leading to the construction of large stations equipped with a vast number of energy storage batteries [3]. Due to their long cycle life, high output voltage, high energy density, and low self-discharge rate [4–8], lithium-ion batteries have gradually become the most promising option among all energy storage battery technologies. Unfortunately, in large-scale energy storage stations, tens of thousands of closely packed battery cells are often deployed, and each individual cell has the potential to undergo thermal runaway due to heat accumulation and material characteristics, which may lead to cascading thermal events and even fire hazards [9]. In recent years, numerous cases of thermal runaway have occurred in both the energy storage and electric vehicle sectors, posing significant challenges to the further development of lithium-ion batteries. Thermal runaway has become the most critical safety issue in the energy storage domain [10–11], and how to detect and intervene in the early stages has become a shared focus of both academia and industry.

Overcharging, overdischarging, and mechanical damage can all trigger the risk of thermal runaway in lithium-ion batteries [12–13]. Fig. 1(a)(b)(c) illustrates the heat generated by mild overcharge and overdischarge conditions [14], while Fig. 1(d)(e)(f) depicts several stages of thermal runaway induced by excessive overcharge. In stage (d), the initial phase, a large amount of heat is generated inside the cell, and the electrolyte begins to produce bubbles, leading to internal pressure buildup. In stage (e), the onset phase, the pressure reaches the threshold of the safety valve, causing it to open and release a significant amount of gas, accompanied by a distinct venting sound. Finally, in stage (f), thermal runaway escalates and propagates violently, during which the electrolyte is expelled and flames may occur and spread [12].[1]

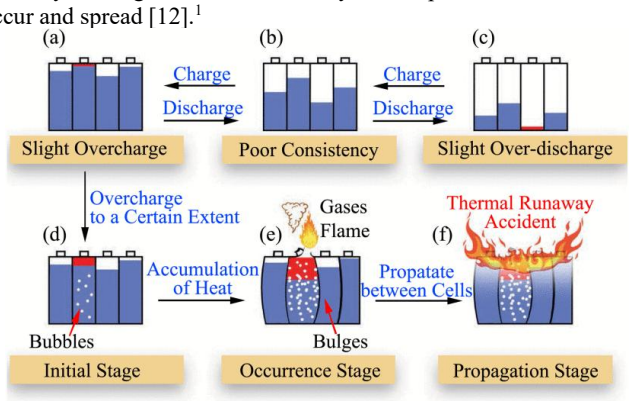

Fig. 1: Mechanism of thermal runaway initiation and propagation [14].(a) and (c): Mild overcharge and overdischarge conditions in individual cells;(b): Module inconsistency;(d)–(f): Representative stages of thermal runaway progression.

Various feasible methods have been proposed for monitoring thermal runaway in lithium batteries, which can be broadly categorized into approaches based on changes in temperature, pressure, and electrical characteristics [15–18]. Temperature-based monitoring typically involves placing thermocouple sensors on the battery surface to track temperature variations, thereby enabling prediction and detection of thermal runaway events [19]. While effective, this approach requires extensive wiring and suffers from significant latency. Pressure-based methods rely on installing pressure sensors between battery cells to detect swelling, thus providing early warnings of abnormal behavior [20]; however, this method also demands complex wiring and suffers from limited detection accuracy. Monitoring based on electrical characteristics involves using battery management systems (BMS) to measure parameters such as impedance, current, and voltage to assess the battery's health status [15]. Nonetheless, such methods may lack direct correlation with thermal runaway events and are often costly.

As previously mentioned, the opening of the safety valve during a thermal runaway event generates a distinct venting sound. Consequently, acoustic-based monitoring methods have attracted growing attention. In the study by Su [21], a venting sound detection method for individual lithium cells was proposed, combining XGBoost and wavelet transform to validate the existence of characteristic acoustic signals during thermal runaway and the feasibility of monitoring such events via sound. Similarly, in the work by Lyu [12], four microphones were deployed inside a battery storage

This work is supported by Science and Technology Project of the State Grid Corporation of China (Evolution mechanism of performance degradation and status sensing methods for lithium-ion battery energy storage system based on advanced acoustic sensing technology, Grant No. 520627230016).

cabin, and a combination of wavelet transform and cross-correlation algorithms was used to detect and locate the safety valve venting sound. These studies demonstrate that acoustic sensing offers an effective approach for detecting thermal runaway in lithium batteries [12].

Although acoustic methods have proven effective for identifying thermal runaway, growing attention has been paid to predicting battery states using machine learning based on acoustic signals. While several datasets exist for machine learning tasks related to battery impedance and capacity [22], there is still a lack of publicly available acoustic datasets specifically focused on abnormal sounds during thermal runaway. Such a dataset is essential for better understanding the acoustic evolution during thermal runaway and for enabling prediction and diagnosis using sound-based methods.

In this work, we address this gap by collecting thermal runaway sound data from lithium-ion batteries using a microphone array and annotating abnormal events throughout the process. We construct a lithium battery thermal runaway sound dataset and validate its effectiveness through a binary classification task using a CNN-based model. Furthermore, to better reflect real-world scenarios where abnormal data is scarce, we also explore an unsupervised anomaly detection approach based on acoustic signals.

Specifically, our contributions are as follows:

1. We present the first acoustic dataset for lithium-ion batteries that includes both normal and thermal runaway states.
2. We provide detailed annotations of abnormal events within the dataset.
3. We evaluate various baseline models and mainstream acoustic event detection methods, offering a valuable benchmark and data foundation for future research.

2. DATASET CONSTRUCTION

2.1. Data Acquisition System Design

We employed a four-microphone array to collect acoustic data. Specifically, our data acquisition system consists of the microphone array and a central processing unit. The microphones used are ICS-43434 digital MEMS microphones, and a DSP chip serves as the processing core for signal acquisition and transmission. The configuration of the microphone array and the data acquisition system is shown in Fig. 2.

Due to varying scales of battery modules in our thermal runaway experiments, the microphone array needed to maintain a certain distance from the center of the runaway cell. Therefore, the shape of the array was not fixed but adjusted according to the size of the battery module. For single-cell thermal runaway experiments, the microphones were placed at the four corners of a rectangular battery module. For experiments involving a full row of cells, where reactions are more intense, the array was mounted on a rectangular metal frame approximately one meter away from the battery pack.

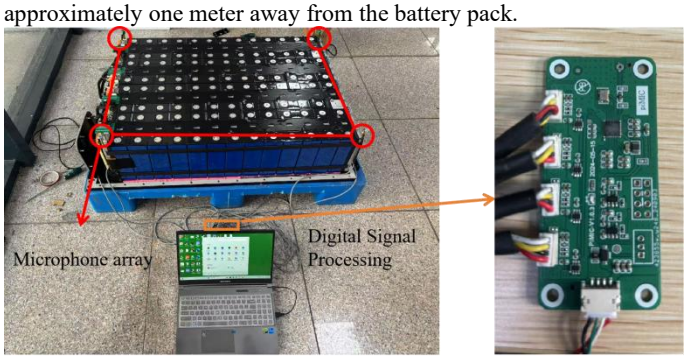

Fig. 2: Microphone array configuration and data acquisition system. (a) microphone array and data acquisition setup; (b) close-up view of the data processing circuit board.

2.2. Thermal Runaway Data Collection

Here, we used overcharging to trigger thermal runaway in the battery cells. The cell model used was 314-0.5C (Narada, China), and the charging method applied was 157 A, 0.5 C direct current. Data acquisition—including sound, temperature, and other physical parameters—began at the start of charging and ended when water cooling was activated. In total, we collected three sets of thermal runaway acoustic signals, with corresponding cell configurations and microphone array layouts summarized in Table 1. The actual experimental site setups are shown in Fig. 3.

Table 1: Battery configurations and microphone array layouts in thermal runaway experiments.

Thermal Runaway ID	Battery Configuration at Test Site	Triggered Cell Configuration	Microphone Array Layout
1	Full Row	Single Cell	Placed at four corners of the full cell row
2	Full Row	Full Cell Row	Placed on rectangular iron frame surrounding the cells
3	Entire Battery Pack	Single Cell	Placed at four corners of the entire battery pack

Fig. 3: On-site thermal runaway cell and microphone array configuration. (a) single-cell thermal runaway triggering experiment on a single-row cell array; (b) whole-row cell triggering experiment on a single-row cell array; (c) single-cell triggering experiment on the entire battery module.

We collected three complete sets of thermal runaway sound data using a four-channel microphone array, with a sampling rate of 16,000 Hz. Each recording segment is 10 seconds long. Specifically, we obtained 100×10 s, 185×10 s, and 186×10 s of thermal runaway audio samples, and recorded corresponding timestamps for synchronization with video footage and subsequent annotation. The recordings from each channel of the microphone array were treated as separate audio files for data processing.

Additionally, we collected two sets of normal charging sound data, comprising 1000×10 s and 41 × 10 s segments, to verify that batteries produce almost no audible noise during normal charging. Due to variations in field conditions, normal charging sounds were not included in the training dataset for the subsequent experiments.

2.3. Dataset Construction

Since the acoustic signals during thermal runaway are primarily caused by the activation of the safety valve and the subsequent venting process, we define the safety valve activation and the following sound as anomalous events associated with thermal runaway. We name it LiB-TRAD, which is, to the best of our knowledge, the first sound dataset specifically dedicated to thermal runaway in lithium-ion

batteries. The annotated format of our thermal runaway dataset is shown in Fig. 4.

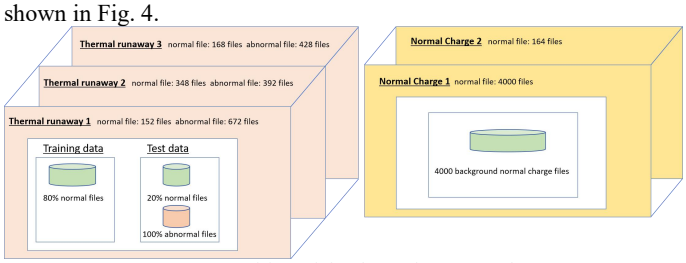

Fig. 4: Composition of the thermal runaway dataset.

We extracted mel-spectrograms from both normal and abnormal signals for comparison, as shown in Fig. 5. For illustrative purposes, the spectrograms shown here are derived from a single representative channel of the microphone array. It can be observed that the normal signals contain a large amount of low-frequency components, which are attributed to the noise generated by fans and operating machinery. In contrast, the abnormal signals exhibit higher-frequency and more energy-concentrated components. This is because the sound source of the abnormal venting signal from the safety valve is closer to the microphone array, and both the opening of the valve and the subsequent venting process produce significant high-frequency components.

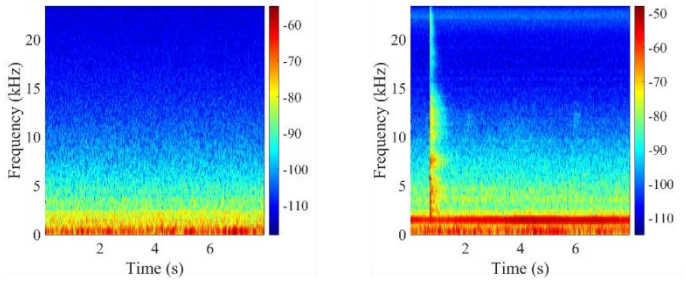

Fig. 5: Comparison of mel-spectrograms between normal and abnormal signals. (a) Mel-spectrogram of a normal signal. (b) Mel-spectrogram of an abnormal signal.

3. DATASET VALIDATION

To demonstrate the effectiveness of the collected data and our abnormal sound labeling method, we first formulate the abnormal sound detection task as a binary classification problem, i.e., a supervised classification task. By analyzing the results on the validation set, we can evaluate the validity of our labeling approach.

We designed a CNN-based network to perform this task, with the network architecture shown in Fig. 6. Specifically, each audio segment is read at a sampling rate of 16 kHz, and a 64-dimensional Mel-spectrogram is extracted (with a window size of 1024 and hop length of 512). The Mel-spectrogram is then converted into a log-mel power spectrogram, which serves as the input to the model.

The CNN model consists of three convolutional blocks, each containing a convolutional layer, Batch Normalization, and a ReLU activation function. Downsampling is performed using MaxPooling layers. Finally, an Adaptive Average Pooling layer followed by a fully connected layer outputs the prediction probability, representing the confidence that the current input is an abnormal sound.

During training, normal samples are labeled as 0 and abnormal samples as 1. The model is optimized using a binary cross-entropy loss function with the Adam optimizer and a learning rate of 1e-3. The dataset is randomly split into 80% training and 20% validation sets. At the end of each epoch, we compute the Area Under the Curve (AUC) metric on the validation set. A higher AUC indicates better performance in distinguishing between abnormal and normal sounds, thereby validating the effectiveness of our data collection and labeling methodology. The training loss curve is shown in Fig. 7.

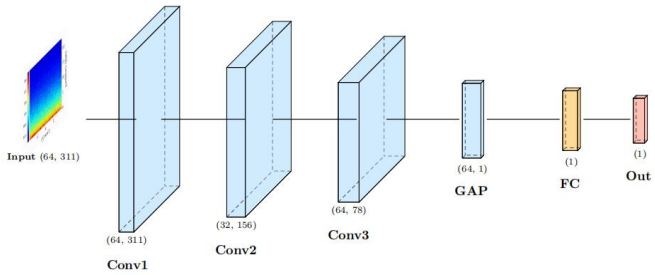

Fig. 6: Schematic diagram of the designed CNN-based binary classification network.

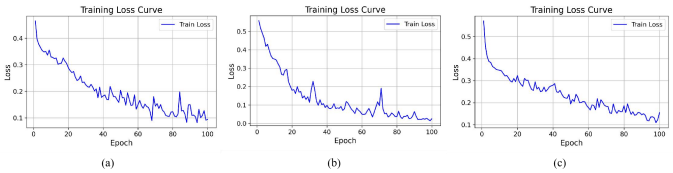

Fig. 7: Loss curves. (a) Training loss curve of Thermal Runaway Experiment 1. (b) Training loss curve of Thermal Runaway Experiment 2. (c) Training loss curve of Thermal Runaway Experiment 3.

The AUC values of our model on the test sets from different thermal runaway experiments are reported in Table 2.

Table 2: Binary classification test results of the CNN model on different thermal runaway experiments in the LiB-TRAD dataset.

Thermal Runaway ID	Test set AUC
1	0.9971
2	0.9876
3	0.9893

The results on the test set indicate that our abnormal sound labeling is highly accurate. However, in real-world production environments, it is often difficult to obtain abnormal sound data. In such cases, abnormal sound detection must be conducted in an unsupervised manner [23], which is similar to Task 2 of DCASE 2025 [24]. Therefore, in the following section, we adopt an unsupervised training approach, where the model is trained using only normal sound data, and abnormal sounds are introduced only during testing.

4. UNSUPERVISED ABNORMAL SOUND DETECTION

The task of unsupervised abnormal sound detection for machine condition monitoring has attracted considerable attention [25]. Abnormal sound detection refers to identifying whether the sound emitted by a target machine is normal or abnormal. In this work, we evaluate several baseline and pre-trained models that have demonstrated strong performance in previous DCASE tasks, including the autoencoder (AE) model [26], the BEATs model [27], the EATs model [28], and the Dasheng model [29]. Additionally, we propose an improved AE model as the baseline specifically tailored for the dataset presented in this paper.

In this study, we further designed and implemented an improved AE model, as illustrated in the flowchart in Fig. 8, for unsupervised detection of abnormal sound data. The model evaluates anomaly scores based on the Mahalanobis distance. The proposed processing pipeline consists of three key stages: data preprocessing, model training, and anomaly evaluation.

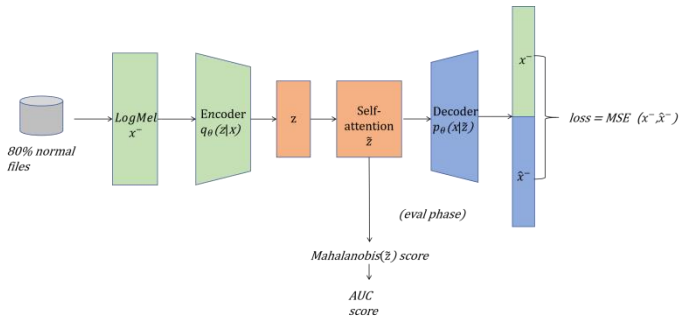

Fig. 8: Flowchart of the autoencoder model with integrated self-attention mechanism.

4.1. Data Preprocessing

Our training set consists of 80% of the normal audio data, while the test set includes the remaining 20% of the normal audio along with 100% of the abnormal audio, enabling an unsupervised detection task.

We use audio data in .wav format sampled at 16 kHz (SR=16000), organized into two folders: normal/ for recordings during regular operation and abnormal/ for recordings during thermal runaway events. To extract time-frequency features from the audio, we use Librosa to compute 64-dimensional Log-Mel spectrograms (with n_mels=64).

The extracted features form a 2D matrix of shape [T, 64], where T denotes the number of time frames. For unified modeling, each [1, 64] frame is treated as an individual input sample during training.

4.2. Model Training

We employ a symmetrical two-layer fully connected architecture and introduces an attention mechanism to enhance the model's ability to focus on key features. The overall structure is as follows:

Encoder:

Linear(64 → 128), Activation: ReLU.

Linear(128 → 64), Output: encoded latent vector z.

Self-Attention Mechanism:

We introduce a self-attention module on the latent vector z, formulated as follows:

$$Q = W_q \cdot z, K = W_k \cdot z, V = W_v \cdot z \quad (1)$$

$$Attention(Q, K, V) = Softmax\left(\frac{QK^T}{\sqrt{d}}\right) \cdot V \quad (2)$$

This module enhances the model's ability to focus on time-frequency components with significant abnormal characteristics, thereby improving the robustness and sensitivity of anomaly detection.

Decoder:

Linear(64 → 128), Activation: ReLU

Linear(128 → 64), Output: reconstructed features

The model uses MSELoss as the reconstruction loss function and is optimized with the Adam optimizer. The learning rate is set to 1e-3, and the model is trained for a total of 200 epochs, with the data shuffled randomly in each epoch. The training batch size is 128, meaning that 128 frames are used for each model update.

4.3. Anomaly Evaluation

After training, during the inference phase, the model takes a test frame as input, extracts the latent representation z through the encoder and attention module, and reconstructs the output $\hat{x}$ using the decoder. During training, the reconstruction error (mean squared error) is used as the loss function to optimize the model's reconstruction ability, expressed as:

$$\mathcal{L}_{MSE} = \frac{1}{n}\sum_{i=1}^{n}(x_i^- - \hat{x}_i^-)^2 \quad (3)$$

To better characterize the difference between abnormal and normal samples in the latent space during inference, the model constructs a multivariate Gaussian distribution based on the mean (μ) and covariance matrix (Σ) of all z values from the training set. The Mahalanobis distance is then used as the final anomaly scoring metric:

$$Score(z) = \sqrt{(z-\mu)^T \Sigma^{-1}(z-\mu)} \quad (4)$$

This distance measures how far the latent representation of the current sample deviates from the distribution of normal training data. A higher score indicates a higher likelihood of being an anomaly.

5. EXPERIMENTAL RESULTS

We evaluated the anomaly detection performance of each model on the test set using the AUC (Area Under ROC Curve) metric to measure recognition accuracy. The results are shown in the table below:

Table 3: Performance comparison of different models.

	Thermal Runaway ID AUC			hmean	Average Inference Time/ms
	1	2	3		
AE	0.4953	0.5236	0.4864	0.5012	0.12
Beats	0.5861	0.4253	0.5742	0.5173	3.52
Eats	0.6324	0.5236	0.5637	0.5697	1.84
Dasheng	0.7125	0.5324	0.6152	0.6112	0.95
Proposed	**0.6892**	**0.6125**	**0.5936**	**0.6291**	**0.17**

The experimental results show that various models trained on this dataset exhibit different levels of performance in the thermal runaway sound anomaly detection task, which validates the effectiveness of the constructed dataset in evaluating model recognition capabilities in complex battery safety scenarios. Among them, the proposed autoencoder achieved stable and excellent results in all three experiments, with a harmonic mean AUC (hmean) of 0.6291, the highest among all models. This reflects strong overall detection ability and good robustness. Additionally, with an average inference time of 0.17 ms, it strikes a balance between performance and efficiency, demonstrating high practical deployment value.

6. CONCLUSION AND FUTURE WORK

This study presents the first publicly available multi-channel acoustic dataset covering the entire thermal runaway process of lithium-ion batteries, comprehensively recording crucial acoustic changes from normal operation to thermal runaway onset. The LiB-TRAD dataset was collected in real experimental scenarios with precise anomaly labeling and phase segmentation, providing a standardized benchmark for training and evaluating anomaly detection algorithms.

We systematically evaluated various typical unsupervised anomaly detection models, including AE, BEATs, EATs, and Dasheng, and proposed an improved autoencoder model as the baseline for our dataset. Evaluation results demonstrate that our dataset effectively distinguishes the stability and generalization capabilities of different models across various experimental scenarios. The proposed improved autoencoder achieved the most balanced performance across three thermal runaway experiments (hmean reaching 0.6291), not only

validating the dataset's applicability for acoustic anomaly detection tasks but also establishing an effective benchmark for future research.

Future work will focus on finer-grained sound event recognition, particularly the acoustic signature modeling of safety valves during early release processes to identify potential warning signals. We also plan to continuously expand the dataset scale, enrich experimental conditions and sensor configurations to enhance model robustness and practical deployment capabilities, ultimately promoting the real-world application of acoustic sensing technology in lithium battery safety monitoring.

REFERENCES

[1] E. Kabalci, "Design and analysis of a hybrid renewable energy plant with solar and wind power," Energy Convers. Manage., vol. 72, pp. 51–59, 2013.

[2] J. B. V. Subrahmanyam, P. Alluvada, K. Bhanupriya, et al., "Renewable energy systems: Development and perspectives of a hybrid solar-wind system," Eng. Technol. Appl. Sci. Res., vol. 2, no. 1, pp. 177–181, 2012.

[3] Y. Yang, C. Lian, C. Ma, et al., "Research on energy storage optimization for large-scale PV power stations under given long-distance delivery mode," Energies, vol. 13, no. 1, p. 27, 2019.

[4] G. Assat, J.-M. Tarascon, "Fundamental understanding and practical challenges of anionic redox activity in Li-ion batteries," Nat. Energy, vol. 3, pp. 373–386, 2018.

[5] L. Chen, M. Fiore, J. E. Wang, et al., "Readiness level of sodium-ion battery technology: a materials review," Adv. Sustain. Syst., vol. 2, 2018, Art. no. 1700153.

[6] J. W. Choi, D. Aurbach, "Promise and reality of post-lithium-ion batteries with high energy densities," Nat. Rev. Mater., vol. 1, pp. 1–16, 2016.

[7] R. Hagiwara, K. Matsumoto, J. Hwang, et al., "Sodium Ion Batteries using Ionic Liquids as Electrolytes," Chem. Rec., vol. 19, pp. 758–770, 2019.

[8] X. Lin, M. Salari, L. M. Arava, et al., "High temperature electrical energy storage: advances, challenges, and frontiers," Chem. Soc. Rev., vol. 45, pp. 5848–5887, 2016.

[9] Q. Wang, B. Mao, S. I. Stoliarov, et al., "A review of lithium ion battery failure mechanisms and fire prevention strategies," Prog. Energy Combust. Sci., vol. 73, pp. 95–131, 2019.

[10] G. Wang, D. Kong, P. Ping, et al., "Revealing particle venting of lithium-ion batteries during thermal runaway: A multi-scale model toward multiphase process," ETransportation, 2023, 16: 100237.

[11] J. Q. Li, D. N. Sun, X. Jin, et al., "Lithium-ion battery overcharging thermal characteristics analysis and an impedance-based electro-thermal coupled model simulation," Applied Energy, 2019, 254: 113574.

[12] N. Lyu, Y. Jin, S. Miao, et al., "Fault warning and location in battery energy storage systems via venting acoustic signal," IEEE J. Emerg. Sel. Topics Power Electron., vol. 11, no. 1, pp. 100–108, 2021.

[13] D. Kong, G. Wang, P. Ping, et al., "A coupled conjugate heat transfer and CFD model for the thermal runaway evolution and jet fire of 18650 lithium-ion battery under thermal abuse," Etransportation, 2022, 12: 100157.

[14] N. Lyu, Y. Jin, R. Xiong, et al., "Real-time overcharge warning and early thermal runaway prediction of Li-ion battery by online impedance measurement," IEEE Trans. Ind. Electron., vol. 69, no. 2, pp. 1929–1936, 2021.

[15] H. Rahimi-Eichi, U. Ojha, F. Baronti, et al., "Battery management system: An overview of its application in the smart grid and electric vehicles," IEEE Ind. Electron. Mag., vol. 7, no. 2, pp. 4–16, 2013.

[16] M. Duff and J. Towey, "Two ways to measure temperature using thermocouples feature simplicity, accuracy, and flexibility," Analog Dialogue, vol. 44, no. 10, pp. 1–6, 2010.

[17] J. Christensen, D. Cook, and P. Albertus, "An efficient parallelizable 3D thermoelectrochemical model of a Li-ion cell," J. Electrochem. Soc., vol. 160, no. 11, pp. A2258–A2264, 2013.

[18] H. Chikh-Bled, K. Chah, Á. González-Vila, et al., "Behavior of femtosecond laser-induced eccentric fiber Bragg gratings at very high temperatures," Opt. Lett., vol. 41, no. 17, pp. 4048–4051, 2016.

[19] L. H. J. Raijmakers, D. L. Danilov, R. A. Eichel, et al., "A review on various temperature-indication methods for Li-ion batteries," Appl. Energy, vol. 240, pp. 918–945, 2019.

[20] R. Li, D. Ren, D. Guo, et al., "Volume deformation of large-format lithium ion batteries under different degradation paths," J. Electrochem. Soc., vol. 166, no. 16, p. A4106, 2019.

[21] T. Su, N. Lyu, Z. Zhao, et al., "Safety warning of lithium-ion battery energy storage station via venting acoustic signal detection for grid application," J. Energy Storage, vol. 38, p. 102498, 2021.

[22] M. F. Ng, J. Zhao, Q. Yan, et al., "Predicting the state of charge and health of batteries using data-driven machine learning," Nat. Mach. Intell., vol. 2, no. 3, pp. 161–170, 2020.

[23] K. Dohi, K. Imoto, N. Harada, et al., "Description and discussion on DCASE 2022 challenge task 2: Unsupervised anomalous sound detection for machine condition monitoring applying domain generalization techniques," arXiv preprint arXiv:2206.05876, 2022.

[24] T. Nishida, N. Harada, D. Niizumi, et al., "Description and discussion on DCASE 2025 challenge task 2: First-shot unsupervised anomalous sound detection for machine condition monitoring," arXiv preprint arXiv:2506.10097, 2025.

[25] Y. Wang, Y. Zheng, Y. Zhang, et al., "Unsupervised anomalous sound detection for machine condition monitoring using classification-based methods," Appl. Sci., vol. 11, no. 23, p. 11128, 2021.

[26] C. Zhang, Y. Liu, and H. Fu, "Ae2-nets: Autoencoder in autoencoder networks," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit., 2019, pp. 2577–2585.

[27] S. Chen, Y. Wu, C. Wang, et al., "BEATs: Audio Pre-Training with Acoustic Tokenizers," in Proc. 40th Int. Conf. Mach. Learn., 2023, pp. 5178-5193.

[28] W. Chen, Y. Liang, Z. Ma, et al., "EAT: Self-supervised pre-training with efficient audio transformer," arXiv preprint arXiv:2401.03497, 2024.

[29] H. Dinkel, Z. Yan, Y. Wang, et al., "Scaling up masked audio encoder learning for general audio classification," arXiv preprint arXiv:2406.06992, 2024.