Proc. of the European Safety and Reliability Conference (ESREL 2026)
Edited by José C. Matos, Paulo B. Lourenço, Michael Beer, Daniel Oliveira, Edoardo Patelli, Hélder S. Sousa, Rui A. Silva and Monica Santamaria
 Published by Research Publishing, Singapore.
doi: 10.3850/ESREL2026061419_esrel26-p26338-cd

Stochastic Physics-Guided Graph–Liquid Neural Networks for Explainable and Real-Time Prognostics of Lithium-Ion Batteries

Jinrui Zhang

School of Electromechanical and Vehicle Engineering, Beijing University of Civil Engineering and Architecture, China. E-mail: Zjrui1227@163.com

Kehui Zhu

School of Electromechanical and Vehicle Engineering, Beijing University of Civil Engineering and Architecture, China. E-mail: 17630974192@163.com

Yanxue Wang*

School of Electromechanical and Vehicle Engineering, Beijing University of Civil Engineering and Architecture, China.E-mail: yan.xue.wang@gmail.com

Reliable and explainable prediction of the Remaining Useful Life (RUL) of lithium-ion batteries is essential for electric vehicles and aerospace systems. Accurate RUL estimation is often hindered by nonlinear degradation behavior, variable operating conditions, and the limited transparency of deep learning models. To address these challenges, a physics-consistent Graph–Liquid Neural Network (GLNN) framework is proposed for real-time battery prognostics. In this framework, capacity degradation is represented as a dynamic graph, where inter-state dependencies are captured through a Graph Neural Network. The nonlinear temporal evolution of degradation is learned adaptively by a Liquid Neural Network with time-varying dynamics. A Wiener-process constraint is incorporated to ensure physical consistency and to improve uncertainty-aware prediction. The model parameters are optimized through joint learning of physical priors and data-driven representations. Two public lithium-ion battery datasets are used to validate the proposed method. Experimental results demonstrate that the GLNN achieves higher prediction accuracy, stronger robustness, and better interpretability compared with state-of-the-art approaches. The physical degradation trend can be effectively preserved, and the prediction uncertainty can be reduced. The results indicate that integrating physics knowledge with adaptive neural reasoning provides a reliable and explainable solution for lithium-ion battery health management under complex and time-varying operating conditions.

Keywords: Remaining Useful Life, Graph–Liquid Neural Network, Physics-Informed Prognostics, Wiener Degradation Process, Lithium batteries.

1. Introduction

In safety-critical aerospace and advanced energy systems, lithium-ion batteries serve as core power and energy storage units, where accurate prediction of performance degradation and remaining useful life (RUL) is essential for ensuring operational safety, mission reliability, and economic efficiency. Fundamentally, battery RUL prognostics is a dynamic modeling and state inference problem for a complex degradation system governed by multi-physics coupling and intrinsic stochasticity. Although data-driven approaches such as support vector machines Mehraj et al. (2025), Gaussian process regression Li et al. (2020), and recurrent neural networks represented by long short-term memory (LSTM) Wang et al. (2023) and its bidirectional variant (BiLSTM) Wang et al. (2024) have demonstrated the ability to learn degradation trends from historical data, they suffer from inherent limitations. These models are largely black-box in nature, heavily dependent on large volumes of high-quality labeled data, and incapable of explicitly incorporating known degradation physics. As a result, their generalization capability and robustness degrade significantly in aerospace-relevant scenarios characterized by limited samples, varying operating conditions, and strong measurement noise, while the resulting predictions often lack physical inter-

pretability Zhang et al. (2024).

To improve predictive performance under complex conditions, signal decomposition and intelligent optimization techniques are frequently employed as auxiliary modules. Methods such as complete ensemble empirical mode decomposition with adaptive noise (CEEMDAN) can separate multi-scale degradation components and noise Wang et al. (2025), while genetic algorithms are commonly used for hyperparameter optimization to enhance prediction accuracy Shahjalal et al. (2022). However, these decomposition–optimization–prediction pipelines primarily operate at the data processing level and do not fundamentally address the core challenge of modeling spatiotemporally coupled degradation dynamics with embedded physical stochasticity. In practice, battery degradation is not merely a one-dimensional temporal process, but a spatiotemporal phenomenon driven by interactions among internal components, such as electrode particles and electrolyte interfaces, and influenced by random microscopic events including lithium dendrite formation and stochastic rupture of solid–electrolyte interphases Lyu et al. (2024). This complexity calls for a unified modeling framework capable of representing spatial structural dependence, continuous-time dynamics, and stochastic physical excitation.

Recent advances in graph neural networks (GNNs) have provided a promising paradigm for modeling systems with explicit or implicit topological relationships Li et al. (2025). Through message-passing mechanisms, GNNs naturally capture dependencies among interacting components, offering an effective means to model degradation propagation across battery units or material structures from a spatial perspective. In parallel, liquid neural networks (LNNs), inspired by biological neural dynamics, define neuron evolution in continuous time using ordinary differential equations and exhibit strong capability in temporal modeling and inherent robustness to noisy inputs Hasani et al. (2021). Their continuous-time formulation is particularly well suited for representing smooth degradation trajectories of battery health indicators, mitigating information loss induced by discrete-time sampling and enabling more faithful modeling of degradation dynamics. Despite these advances, existing GNN- and LNN-based approaches largely focus on deterministic settings and fail to explicitly account for the intrinsic uncertainty of battery degradation. In reality, battery ageing is inherently stochastic, with microscopic defect accumulation and fluctuating operating conditions introducing diffusion-like uncertainty around the nominal degradation trend Li et al. (2024). The Wiener process, or Brownian motion, serves as a classical mathematical foundation for describing such stochastic degradation phenomena in continuous time Zhang et al. (2023). Embedding Wiener-driven stochastic dynamics into neural models represents a critical step toward physics-informed probabilistic prognostics, enabling the model to capture not only the expected degradation trajectory but also the uncertainty evolution, which is essential for reliable RUL estimation.

Motivated by these considerations, a stochastic physics-informed liquid graph neural network (SPILGNN) is proposed for lithium-ion battery RUL prediction in aerospace and energy systems. A Wiener-driven stochastic differential equation is embedded into a continuous-time liquid neural architecture, while graph representations are employed to model spatial degradation dependencies. This formulation jointly captures inter-cell structural correlations, nonstationary degradation dynamics, and intrinsic stochasticity, enabling physically interpretable and uncertainty-aware prognostics. The main contributions of this work are summarized as follows: Stochastic physics-informed liquid graph neural framework. The Wiener degradation prior is incorporated through a physics-informed stochastic regularization term that constrains the learned degradation increments. Battery degradation is formulated as a graph-structured spatiotemporal process, where stochastic liquid units drive message passing and state evolution to capture structural dependencies and nonstationary temporal dynamics;Physics-constrained and uncertainty-aware learning. A learning strategy combining degradation-consistent regularization and varia-

tional inference is established to ensure physical plausibility and provide probabilistic RUL prediction under complex operating conditions.

2. Methodology

2.1. *Problem formulation and stochastic degradation prior*

Let $y(t)$ denote the battery capacity in continuous time, and let $\{y_t\}_{t=1}^{T}$ be the corresponding discrete observations during cycling. The degradation increment between two adjacent observations is defined as

$$\Delta y_t = y_t - y_{t-1}, \qquad t = 2, 3, \ldots, T, \quad (1)$$

which is generally non-positive during battery ageing. Given a historical window $\mathbf{X}_t = [x_{t-L+1}, x_{t-L+2}, \ldots, x_t]$, the model predicts the next degradation increment $\hat{\Delta y}_{t+1}$ and reconstructs the capacity recursively as

$$\hat{y}_{t+1} = \hat{y}_t + \hat{\Delta y}_{t+1}. \quad (2)$$

In this way, the model learns degradation evolution at the increment level rather than directly fitting the absolute capacity trajectory.

To describe the stochastic nature of battery degradation, the latent ageing process is modeled by a Wiener-type stochastic differential equation

$$dy(t) = \mu(t)\, dt + \sigma(t)\, dW(t), \quad (3)$$

where $\mu(t)$ and $\sigma(t)$ denote the drift and diffusion terms, respectively, and $W(t)$ is a standard Wiener process. Over a small interval Δt, the degradation increment follows the local form $\Delta y_t \approx \mu_t \Delta t + \sigma_t \sqrt{\Delta t}\epsilon_t$, where $\epsilon_t \sim \mathcal{N}(0,1)$. This implies that the expected local trend is governed by $\mu_t \Delta t$, while the uncertainty level is controlled by $\sigma_t^2 \Delta t$. Since capacity fading is dominant during battery ageing, the drift is constrained to remain non-positive. In the proposed framework, this Wiener prior is used to regularize the predicted degradation increments rather than directly specifying the hidden-state transition.

2.2. *Degradation graph and liquid state evolution*

For each input window, a degradation graph $\mathcal{G}_t = (\mathcal{V}_t, \mathcal{E}_t)$ is constructed. Each local node corresponds to one observation step and is initialized by

$$\mathbf{h}_i^{(0)} = \phi(x_i), \qquad i = t - L + 1, \ldots, t, \quad (4)$$

where $\phi(\cdot)$ is a learnable projection function. In addition, a global degradation node is introduced to summarize the overall ageing condition of the current window using statistical descriptors such as the mean value, fluctuation level, dynamic range, and difference information. The graph contains temporal adjacency edges, cross-step dependency edges, and bidirectional connections between local nodes and the global node, allowing local evolution patterns and global degradation context to be modeled jointly.

Based on the constructed graph, node interactions are first aggregated from the neighborhood, and the current state, local input, and neighborhood message are fused to obtain a candidate hidden state:

$$\tilde{\mathbf{h}}_i(t) = \tanh\Big(\mathbf{W}_h \mathbf{h}_i(t) + \mathbf{W}_x \mathbf{x}_i + \mathbf{W}_n \mathbf{m}_i(t) + \mathbf{b}\Big), \quad (5)$$

where $\mathbf{m}_i(t)$ denotes the aggregated neighborhood message and $\mathbf{W}_h$, $\mathbf{W}_x$, $\mathbf{W}_n$, and $\mathbf{b}$ are trainable parameters.

To capture adaptive continuous-time state evolution, the hidden dynamics are formulated as

$$\frac{d\mathbf{h}_i(t)}{dt} = -\frac{1}{\tau_i(t)}\mathbf{h}_i(t) + \frac{1}{\tau_i(t)}\tilde{\mathbf{h}}_i(t), \quad (6)$$

where $\tau_i(t) > 0$ is an adaptive time constant determined by the current node state, input, and neighborhood message. This equation describes a relaxation process in which the hidden state evolves continuously toward the candidate state, with the update speed adjusted according to the current degradation condition. Using a forward Euler discretization, the continuous-time evolution is written in the gated form

$$\mathbf{h}_i^{(k+1)} = \left(1 - \beta_i^{(k)}\right)\mathbf{h}_i^{(k)} + \beta_i^{(k)}\tilde{\mathbf{h}}_i^{(k)}, \quad (7)$$

where $\beta_i^{(k)} = \Delta t / \tau_i^{(k)}$. Therefore, the gated update adopted here can be interpreted as the discrete implementation of the continuous-time liquid dynamics.

After K rounds of graph evolution, the final local and global node states are aggregated into

a graph representation $\mathbf{z}_t$, from which the next degradation increment $\hat{\Delta y}_{t+1}$ is predicted by a regression head, followed by capacity reconstruction using Eq. (2).

2.3. *Training objective*

To incorporate the stochastic degradation prior, a physics-informed regularization term is imposed on the predicted increments:

$$\mathcal{L}_{\text{phys}} = \frac{1}{T-1}\sum_{t=2}^{T}\Big[\Big(\hat{\Delta y}_t - \mu_t \Delta t\Big)^2 + \lambda_\sigma\left(\hat{v}_t - \sigma_t^2 \Delta t\right)^2\Big] \tag{8}$$

where $\hat{v}_t$ denotes the empirical local variance of predicted increments within a sliding window, and λ_σ is a balancing coefficient. The first term constrains the local degradation tendency, while the second term aligns the fluctuation level of the prediction with the diffusion intensity implied by the Wiener prior.

To preserve the monotonic fading property of battery capacity, a monotonicity penalty is further introduced:

$$\mathcal{L}_{\text{mono}} = \frac{1}{T-1}\sum_{t=2}^{T}\max\Big(\hat{\Delta y}_t, 0\Big). \tag{9}$$

The supervised loss is defined on the reconstructed capacity sequence,

$$\mathcal{L}_{\text{sup}} = \frac{1}{T-1}\sum_{t=2}^{T}(\hat{y}_t - y_t)^2, \tag{10}$$

and the final objective is

$$\mathcal{L} = \mathcal{L}_{\text{sup}} + \lambda_{\text{phys}}\mathcal{L}_{\text{phys}} + \lambda_{\text{mono}}\mathcal{L}_{\text{mono}}. \tag{11}$$

Through joint optimization of these terms, the model learns the observed degradation trajectory while maintaining stochastic consistency and physically plausible monotonic evolution.

3. Datasets and Degradation Characteristics

3.1. *Degradation Characteristics of the XJU Dataset*

The XJU dataset consists of 55 cylindrical lithium-ion cells cycled at room temperature under strictly identical voltage limits of 4.2 V and 2.5 V. The corresponding capacity degradation trajectories are presented in Fig 1, where pronounced cell-to-cell variability can be observed despite the highly controlled operating conditions. A typical

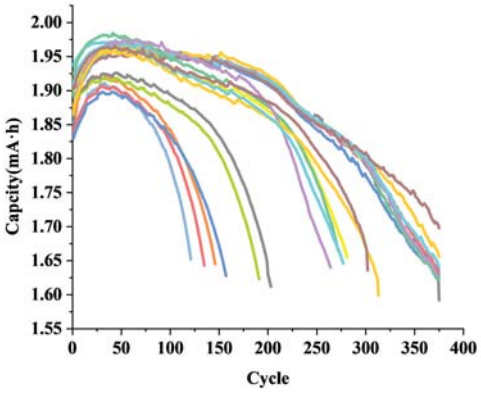

Fig. 1. Capacity degradation trajectories of the XJU dataset under fixed cycling conditions

ageing process emerges, including a short initial stabilization stage, a quasi-linear degradation phase, and a rapid capacity decline near end-of-life. However, both the degradation rates and the transition points between stages vary significantly among cells. Moreover, local fluctuations and short-term irregularities are frequently superimposed on the long-term fading trends, as illustrated in Fig 1.

Since all cells share the same operating protocol, this variability primarily reflects intrinsic stochasticity arising from manufacturing tolerances, microstructural differences, and non-uniform SEI formation. As a result, the XJU dataset is characterized by relatively shorter lifetimes and stronger noise, making it particularly suitable for evaluating model robustness to cell-level uncertainty.

3.2. *Degradation Characteristics of the HUST Dataset*

The HUST dataset includes 77 cells tested at 30 °C with identical charging conditions and diverse multi-stage discharge strategies. The resulting long-term capacity evolution is shown in Fig 2. Compared with XJU, the cells in this dataset

exhibit substantially longer cycle lives and overall smoother degradation trajectories. Although the

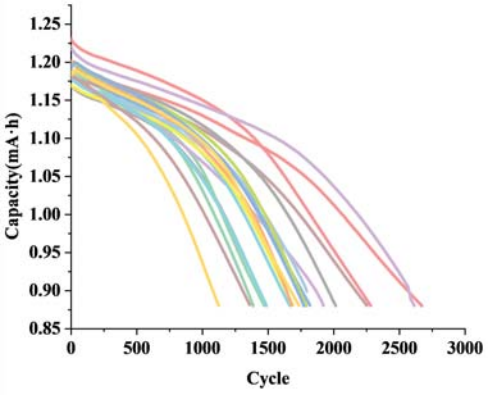

Fig. 2. Capacity degradation trajectories of the HUST dataset with extended cycle life

discharge protocols vary across cells, the capacity fading remains gradual and stable over extended operating horizons, with lifetimes ranging from several hundred to nearly 3000 cycles, as reflected in Fig 2. Short-term fluctuations are less pronounced, and the degradation curves generally preserve well-structured long-term trends.

This combination of long lifespan and smooth degradation places stringent requirements on long-horizon prediction stability. Prognostic models must sustain accurate and physically consistent extrapolation over thousands of cycles while adapting to protocol-dependent degradation rates, making the HUST dataset an effective benchmark for evaluating long-term prognostic performance.

Overall, Fig 1 and Fig 2 illustrate two complementary degradation scenarios: noise-dominated, short-lifetime ageing under fixed conditions and smooth, long-term degradation under protocol variability, respectively.

4. Experimental Validation and Performance Evaluation

4.1. *Results on the Short-Life, High-Fluctuation Dataset (XJU)*

The XJU dataset is characterized by relatively short cycle life and noticeable cycle-to-cycle capacity variability, posing a stringent test for long-horizon prognostics under stochastic disturbances. As shown in Fig 3, the proposed stochastic physics-informed graph–liquid model yields predictions that closely overlap with the measured capacity trajectories for the representative cells. The model not only matches the early-stage quasi-linear fading region, but also preserves the curvature evolution as the degradation rate increases. Importantly, no progressive drift is observed as cycling proceeds, indicating that the learned dynamics remain stable even when local fluctuations are present. This behavior is consistent with the role of the embedded Wiener-type stochastic prior, which regularizes the long-term evolution and prevents the predictor from being over-driven by short-term perturbations. The cycle-wise error

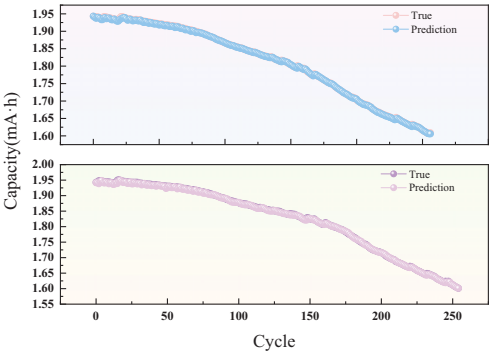

Fig. 3. Comparison between predicted and measured capacity trajectories for representative XJU cells (XJU3 and XJU4)

distributions in Fig 4 provide a complementary view of reliability. Across XJU1–XJU4, the error mass is strongly concentrated near zero, implying that most cycles are predicted with small deviations. A long-tail behavior is visible for XJU1, suggesting the existence of sporadic outlier cycles where prediction error increases sharply. However, the central width of the violin remains narrow, indicating that these larger errors are rare rather than persistent. This is further supported by Table 1, where all cells achieve R2 values close to unity, demonstrating that the global degradation shape is accurately captured. Meanwhile, the

Table 1. Prediction performance on the XJU dataset.

Cell	MSE	RMSE	R^2
XJU1	0.0000117	0.0034234	0.9987541
XJU2	0.0000041	0.0020234	0.9996153
XJU3	0.0000043	0.0020667	0.9996183
XJU4	0.0000070	0.0026445	0.9993795
Average	**0.0000068**	**0.0025395**	**0.9993418**

RMSE remains at the 10^{-3} scale for all cells, with XJU1 showing the largest value, consistent with its heavier-tailed error distribution. Overall, the XJU results verify that the proposed framework can maintain trajectory-level fidelity while being resilient to irregular, non-stationary fluctuations.

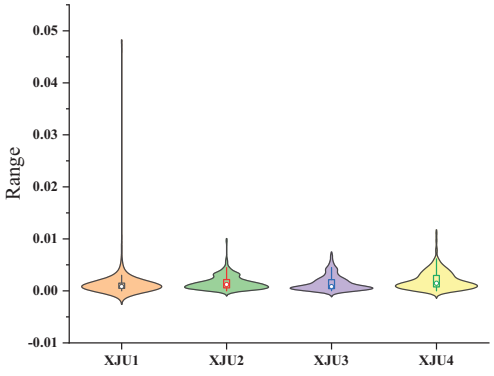

Fig. 4. Violin plots of cycle-wise prediction errors for all XJU cells

4.2. *Results on the Long-Life, Stable Dataset (HUST)*

The HUST dataset exhibits substantially longer cycle life and smoother degradation trajectories, enabling evaluation of whether the proposed model can sustain consistency over extended horizons. As illustrated in Fig 5, the predicted curves for the representative cells almost entirely overlap with the measured capacities across thousands of cycles. The model captures both the slow early-stage decay and the gradual curvature change toward later cycles, without introducing oscillations or bias accumulation. This indicates that the continuous-time liquid dynamics provide sufficient flexibility to represent long-term gradual ageing, while the stochastic prior acts as a stability constraint that discourages unphysical deviations when extrapolating far into the future.

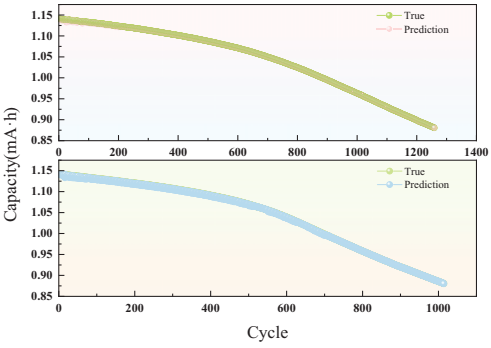

Fig. 5. Comparison between predicted and measured capacity trajectories for representative HUST cells (HUST3 and HUST4)

The error distributions in Fig 6 are notably more compact than those in Fig 4, with shorter tails and smaller spread for all cells, reflecting the reduced volatility of this dataset. Table 2 further confirms this improvement quantitatively, where the average RMSE decreases relative to the short-life dataset and the average R^2 increases to an even higher level. The consistency across HUST1–HUST4 suggests that the proposed model is not tuned to a single degradation profile, but generalizes across cells with similar operating conditions and lifetime characteristics. Together with the XJU results, these findings demonstrate scalability across fundamentally different ageing regimes, from strongly fluctuating short-life cells to long-life, smooth degradation processes.

4.3. *Comparison with Deep Sequence Baselines*

To ensure a fair comparison, all competing methods were evaluated under the same protocol. Specifically, the same training–test split, sliding-window setting, prediction procedure, and metric definitions were used for all models on each dataset. All results were calculated on the original capacity scale. For multi-step forecasting,

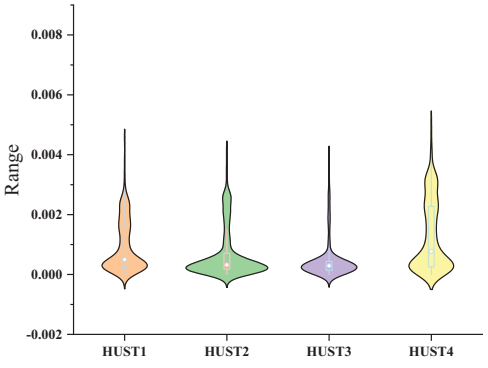

Fig. 6. Violin plots of cycle-wise prediction errors for all HUST cells

Table 2. Prediction performance on the HUST dataset.

Cell	MSE	RMSE	R^2
HUST1	0.0000016	0.0012807	0.9997095
HUST2	0.0000012	0.0011094	0.9998004
HUST3	0.0000008	0.0009106	0.9998623
HUST4	0.0000031	0.0017572	0.9995192
Average	**0.0000017**	**0.0012645**	**0.9997229**

the same recursive prediction setting was applied consistently across all methods. In addition, the baseline models were implemented according to their standard formulations, and their main hyperparameters were selected based on the original references or validation performance under the same data conditions.

Fig. 7 and Table 3 present the comparison with representative deep sequence baselines, including Bi-LSTM and Transformer. The proposed method achieves the lowest RMSE on all evaluated cells in both the XJU and HUST datasets. In the XJU dataset, the errors of the baseline models vary substantially across cells, while the proposed method remains consistently low. Combined with the prediction curves, this result suggests that, under abnormal fluctuations, the baseline sequence models are more likely to be disturbed by local outliers and then produce overly smooth or nearly horizontal long-horizon predictions, leading to pronounced deviation from the true degradation trajectory. By contrast, the proposed method maintains stable tracking performance across XJU1–XJU4, which is consistent with the role of the Wiener-based stochastic degradation constraint and the adaptive state update in stabilizing the long horizon evolution.

Table 3. RMSE comparison between the proposed method and baseline models.

Cell	Proposed	Bi-LSTM	Transformer
XJU1	**0.003423**	0.441970	0.372117
XJU2	**0.002023**	0.445234	0.202919
XJU3	**0.002067**	0.429428	0.068047
XJU4	**0.002645**	0.390955	0.068466
Average	**0.002540**	0.426897	0.177887
HUST1	**0.001281**	0.391089	0.108757
HUST2	**0.001109**	0.421444	0.167379
HUST3	**0.000911**	0.368368	0.108765
HUST4	**0.001757**	0.373731	0.292421
Average	**0.001265**	0.388658	0.169331

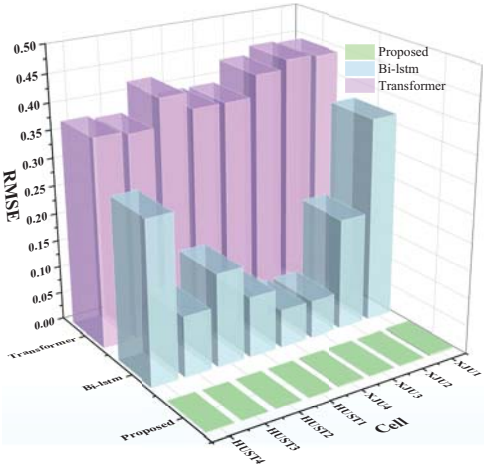

Fig. 7. RMSE comparison between the proposed method and deep sequence baselines on the XJU and HUST datasets.

A similar pattern is observed for the HUST dataset. Although the degradation trajectories in HUST are relatively smoother, the two baseline models still show consistently higher errors than the proposed method across all tested cells. This indicates that the advantage of the proposed

framework is not limited to strongly fluctuating cases, but is maintained under different degradation characteristics. Overall, the comparison shows that the proposed method provides more accurate and more stable long-horizon prediction than the two representative deep sequence baselines under the same evaluation setting.

5. Conclusion

A stochastic physics-informed graph–liquid neural framework is developed for lithium-ion battery RUL prediction by integrating degradation-increment modeling and Wiener-driven stochastic regularization. The proposed architecture captures nonstationary ageing via joint local–global spatiotemporal learning and enables robust online prognostics under time-varying conditions.

Future studies will extend the framework to multi-modal health indicators, operating-dependent stochastic dynamics, and uncertainty-aware RUL intervals, with broader applicability to physics-informed prognostics of complex energy systems.

Acknowledgement

This study was sponsored by the National Natural Science Foundation of China (52275079), the Youth Beijing Scholars program, the Guangxi Science and Technology Major Project (AA23062031).

References

Hasani, R., M. Lechner, A. Amini, D. Rus, and R. Grosu (2021). Liquid time-constant networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Volume 35, pp. 7657–7666.

Li, F., S. Zhou, and R. Liu (2024). Parameter estimation for the hammerstein-wiener nonlinear system and application in lithium-ion batteries. *Journal of Energy Storage 102*, 114265.

Li, X., Y. Wang, J. Xing, and Y. Wang (2025). Causal graph inference with adaptive dynamic structure learning for mechanism-oriented fault diagnosis in dynamic industrial systems. *Reliability Engineering & System Safety*, 111865.

Li, X., C. Yuan, and Z. Wang (2020). Multi-time-scale framework for prognostic health condition of lithium battery using modified gaussian process regression and nonlinear regression. *Journal of Power Sources 467*, 228358.

Lyu, D., B. Zhang, E. Zio, and J. Xiang (2024). Battery cumulative lifetime prognostics to bridge laboratory and real-life scenarios. *Cell Reports Physical Science 5*(9).

Mehraj, N., C. Mateu, H. Bastida, Y. Li, Y. Ding, A. Sciacovelli, and L. F. Cabeza (2025). Artificial intelligence in state of charge estimation: Pioneering approaches across energy storage systems. *Energy*, 138166.

Shahjalal, M., P. K. Roy, T. Shams, A. Fly, J. I. Chowdhury, M. R. Ahmed, and K. Liu (2022). A review on second-life of li-ion batteries: Prospects, challenges, and issues. *Energy 241*, 122881.

Wang, F., Z. Zhai, Z. Zhao, Y. Di, and X. Chen (2024). Physics-informed neural network for lithium-ion battery degradation stable modeling and prognosis. *Nature Communications 15*(1), 4332.

Wang, Y., S. Guo, Y. Cui, L. Deng, L. Zhao, J. Li, and Z. Wang (2025). A comprehensive review of machine learning-based state of health estimation for lithium-ion batteries: data, features, algorithms, and future challenges. *Renewable and Sustainable Energy Reviews 224*, 116125.

Wang, Y., J. Zhu, L. Cao, B. Gopaluni, and Y. Cao (2023). Long short-term memory network with transfer learning for lithium-ion battery capacity fade and cycle life prediction. *Applied Energy 350*, 121660.

Zhang, J., D. Lyu, and J. Xiang (2024). A model-data-fusion method for real-time continuous remaining useful life prediction of lithium batteries. *Measurement 238*, 115312.

Zhang, Y., X. Feng, M. Zhao, and R. Xiong (2023). In-situ battery life prognostics amid mixed operation conditions using physics-driven machine learning. *Journal of Power Sources 577*, 233246.