

# RUL-QMoE: Multiple Non-crossing Quantile Mixture-of-Experts for Probabilistic Remaining Useful Life Predictions of Varying Battery Materials

Sel Ly, Rufan Yang, Ninad Dixit, Hung Dinh Nguyen

Nanyang Technological University, Singapore 639798  
 {sel.ly, hunghtd}@ntu.edu.sg, {rufan001, ninad001}@e.ntu.edu.sg

## Abstract

Lithium-ion (Li-ion) batteries are the major type of battery used in a variety of everyday applications, including electric vehicles (EVs), mobile devices, and energy storage systems. Predicting the Remaining Useful Life (RUL) of Li-ion batteries is crucial for ensuring their reliability, safety, and cost-effectiveness in battery-powered systems. The materials used for the battery cathodes and their designs play a significant role in determining the degradation rates and RUL, as they lead to distinct electrochemical reactions. Unfortunately, RUL prediction models often overlook the cathode materials and designs to simplify the model-building process, ignoring the effects of these electrochemical reactions. Other reasons are that specifications related to battery materials may not always be readily available, and a battery might consist of a mix of different materials. As a result, the predictive models that are developed often lack generalizability. To tackle these challenges, this paper proposes a novel material-based Mixture-of-Experts (MoE) approach for predicting the RUL of batteries, specifically addressing the complexities associated with heterogeneous battery chemistries. The MoE is integrated into a probabilistic framework, called *Multiple Non-crossing Quantile Mixture-of-Experts for Probabilistic Prediction (RUL-QMoE)*, which accommodates battery operational conditions and enables uncertainty quantification. The RUL-QMoE model integrates specialized expert networks for five battery types: LFP, NCA, NMC, LCO, and NMC-LCO, within a gating mechanism that dynamically assigns relevance based on the battery's input features. Furthermore, by leveraging non-crossing quantile regression, the proposed RUL-QMoE produces coherent and interpretable predictive distributions of the battery's RUL, enabling robust uncertainty quantification in the battery's RUL prediction. Trained on seven real-world datasets, the proposed RUL-QMoE achieves strong predictive performance across all battery types, with MAE = 65 (cycles), MAPE = 9.59%, RMSE = 100 (cycles), and  $R^2 = 96.84\%$ . Compared to traditional models like XGBoost, Random Forest, CNN, and LSTM, the proposed RUL-QMoE model consistently delivers lower RMSE and superior probabilistic insights, including survival probabilities and prediction intervals. The model has been integrated into our Battery AI platform in collaboration with Toyota Motor Engineering & Manufacturing North America, Inc., as part of a broader Battery Foundation Model initiative. This RUL-QMoE model will serve future Toyota EVs' users

and battery system designers.

## 1 Introduction

Lithium-ion batteries underpin a wide spectrum of modern technologies, from electric vehicles and renewable energy storage to portable electronics (Xie et al. 2024). However, their performance inevitably declines due to degradation mechanisms such as capacity fade, abnormal heat generation, and internal resistance growth (Yang and Nguyen 2025). These phenomena not only reduce operational efficiency but also raise safety concerns, particularly as batteries approach end-of-life (Xie, Weng, and Nguyen 2022). Accurate prediction of Remaining Useful Life (RUL) is therefore essential for ensuring safety, reliability, and economic efficiency. By forecasting the time or cycles before a battery's performance falls below acceptable thresholds, RUL prediction supports proactive maintenance scheduling, mitigates the risk of sudden failures, and optimizes end-of-life management for both economic and environmental benefits.

Traditional approaches to RUL prediction fall into three categories: physics-based models, data-driven models, and hybrid models (Hasib et al. 2021). Physics-based methods capture electrochemical and thermal processes in detail but are hindered by computational complexity and high sensitivity to parameter uncertainty (Wang et al. 2023). Data-driven methods, powered by machine learning (ML), bypass explicit modelling by learning patterns directly from historical and real-time data. They offer flexibility and scalability but demand large, high-quality datasets and often fail to generalize across battery chemistries due to dataset heterogeneity and feature engineering limitations (Zhang et al. 2024). Hybrid approaches seek to merge the interpretability of physics-based models with the flexibility of ML, yet integration challenges and inadequate uncertainty quantification remain unresolved. What is more, existing RUL prediction methods failed to address two problems: 1. Material-specific degradation: Existing models are typically tailored to a single type of Lithium-ion battery. However, different types of Lithium-ion batteries have different degradation and operational features. For instance, the LFP batteries exhibit flat voltage plateaus, while NMC shows steep capacity drops. 2. Uncertainty propagation: Existing methods provide point estimates of the battery RUL without confidence intervals, which are inaccurate since the factors to RUL are uncertain

during the operational life of the targeted Li-ion battery.

Artificial intelligence (AI) is now reshaping the RUL prediction landscape. Advances in deep learning, probabilistic modelling, and mixture-of-experts (MoE) architectures enable AI systems to learn degradation patterns from diverse datasets, adapt to varying chemistries and usage profiles, and provide probabilistic forecasts that account for uncertainty. Such capabilities are critical for real-world applications, where battery chemistries, such as Lithium Iron Phosphate (LFP), Nickel Cobalt Aluminum (NCA), Nickel Manganese Cobalt (NMC), and Lithium Cobalt Oxide (LCO) - exhibit distinct degradation signatures. Moreover, AI-based frameworks can be designed for continuous learning, improving prediction accuracy as new data becomes available.

In this study, we propose RUL-QMoE, a Multiple Non-crossing Quantile Mixture-of-Experts framework for probabilistic battery RUL prediction. Each expert model is specialized for a specific battery chemistry, while a gating network dynamically assigns relevance based on input features. By predicting multiple quantiles of the RUL distribution, RUL-QMoE offers a comprehensive view of potential outcomes rather than single-point estimates, enabling robust uncertainty quantification. To ensure coherent probabilistic outputs, the model enforces monotonicity across quantiles through a non-crossing constraint. Leveraging seven diverse real-world datasets, including CALCE (Xing et al. 2013; He et al. 2011), MATR (Severson et al. 2019; Hong et al. 2020), HUST (Ma et al. 2022), HNEI (Devie, Baure, and Dubarry 2018), RWTH (Li et al. 2021), SNL (Preger et al. 2020) and UR-PUR (Juarez-Robles, Jeevarajan, and Mukherjee 2020; Juarez-Robles et al. 2021) - our framework demonstrates high adaptability, scalability, and accuracy across chemistries and operational profiles.

By integrating AI-based tools, probabilistic forecasting, and chemistry-specific expertise, RUL-QMoE advances the state-of-the-art in battery health management. The RUL-QMoE model is currently being deployed as part of the Battery Foundation Model initiative in collaboration with Toyota Motor Engineering & Manufacturing North America, Inc., where it powers Battery AI agents for predictive maintenance in EVs. Also, it will lay the foundation for intelligent, data-adaptive, and risk-aware predictive maintenance systems that can scale with evolving battery technologies, accelerating the transition toward safer, more reliable, and more sustainable energy storage solutions.

## 2 Related Work

Recent works in data-driven RUL prediction for lithium-ion batteries leverage historical or real-time datasets to learn degradation behaviours without relying on explicit physical models. For instance, (Catelani et al. 2021) proposed a hybrid approach that combines a physics-based single exponential degradation model with a deep echo state network, optimized via genetic algorithms, and validated using NASA’s lithium-ion datasets. Similarly, (Yang 2021) developed a hybrid 3D-2D convolutional neural network with feature and multi-scale cycle attention mechanisms to predict battery cycle life and RUL, utilizing voltage, current, and temperature curves from charge cycles, and achieving

test errors of 1.1% and 3.6% under diverse operating conditions. While these models demonstrate the potential of deep learning to enhance prediction accuracy, their architectures are often monolithic, limiting adaptability to heterogeneous datasets and diverse degradation patterns.

To address these limitations, recent advancements have explored Mixture-of-Experts (MoE) architectures, which are particularly effective for battery forecasting tasks due to their ability to partition the learning process into specialized sub-models, each capturing distinct aspects of degradation. In (Lei and Yu 2025), the authors introduced *MSPMLP*, a multi-scale patch-based MLP combined with sparsely-gated MoE to forecast battery capacity, effectively modelling both long-term degradation and short-term recovery, and achieving a 41.8% improvement in MAPE over prior methods. Likewise, (Chen and Zhou 2024) proposed *AttMoE*, integrating transformer-based attention with MoE to enhance faded capacity prediction under noisy sensor conditions, yielding 10–20% improvements in relative error. Furthermore, (Zhao et al. 2024) presented *MMMe*, a hybrid model combining bidirectional GRUs, attention modules, and MLP-Mixer blocks with a 32-expert MoE layer, excelling in modelling both gradual degradation and sudden capacity drops with a RMSE of 0.0515 on the NASA dataset. These studies collectively underscore the versatility of MoE frameworks in handling heterogeneous data and capturing complex degradation behaviours. However, most existing MoE models focus on single-point predictions or are tailored to specific chemistries. In contrast, our proposed RUL-QMoE advances the field by delivering probabilistic RUL predictions across diverse battery types, ensuring non-crossing quantile outputs and robust uncertainty quantification.

## 3 Problem Formulation

In this section, we briefly introduce the task of constructing a predictive model for the conditional quantile of a response variable  $Y$ , given an input  $\mathbf{X} = \mathbf{x}$ , denoted as  $Q_{Y|\mathbf{X}=\mathbf{x}}(\tau)$ . Here,  $Y$  represents the remaining useful life (RUL) of a battery, defined as the number of cycles the battery can continue operating before its capacity drops below a specified threshold (e.g., 80% of its nominal capacity). The input vector  $\mathbf{X}$  comprises features that influence the battery’s RUL.

A basic approach to this problem is the quantile regression model proposed in (Koenker and Bassett Jr 1978). The quantile  $Q_Y(\tau)$  at level  $\tau \in (0, 1)$  is defined as the generalized inverse of the cumulative distribution function (CDF)  $F_Y(y) := \mathbb{P}(Y \leq y)$ , see Figure 1:

$$Q_Y(\tau) := F_Y^{-1}(\tau) = \inf\{y \in \mathbb{R} : \mathbb{P}(Y \leq y) \geq \tau\}. \quad (1)$$

Accordingly, the conditional quantile of  $Y$  given  $\mathbf{X} = \mathbf{x}$  is expressed as:

$$Q_{Y|\mathbf{X}=\mathbf{x}}(\tau) = \inf\{y \in \mathbb{R} : \mathbb{P}(Y \leq y | \mathbf{X} = \mathbf{x}) \geq \tau\}. \quad (2)$$

When  $\tau = 0.5$ , this corresponds to the conditional median of  $Y$ . By selecting multiple quantile levels, such as  $\tau \in \{0.05, 0.1, 0.2, \dots, 0.9, 0.95\}$ , we can generate a range of quantile estimates for  $Y$ , thereby enabling the construction of  $F_{Y|\mathbf{X}=\mathbf{x}}(y)$  for probabilistic forecasting.




Figure 1: Quantile definition:  $Q_Y(\tau) = F_Y^{-1}(\tau)$ ,  $\tau \in (0, 1)$ .

In this work, we aim to develop a general Mixture-of-Experts model for estimating the conditional distribution of a battery’s RUL, applicable across various chemistries such as LFP, NCA, NMC, LCO, and mixed NMC-LCO types. This is significantly important because each battery chemistry exhibits distinct RUL distributions, and conventional machine learning models typically capture only a single distribution per battery type.

Moreover, when trained on multiple quantile levels, standard models often violate the non-decreasing property of quantiles, leading to the quantile crossing problem. To address this, we design our MoE framework to ensure monotonicity across quantile outputs. Specifically, for any two quantile levels  $\tau_1 < \tau_2$ , the model guarantees that the predicted quantiles satisfy  $Q_{Y|X=x}(\tau_1) < Q_{Y|X=x}(\tau_2)$ . This constraint preserves the natural ordering of quantiles and contributes to more accurate and coherent predictive distributions.

## 4 Proposed Model

As aforementioned, we propose a novel MoE model that integrates five specialized Expert networks, each corresponding to a specific battery chemistry type: LFP, NCA, NMC, LC, and mixed NMC-LCO. The second core of the MoE model is a Gating network, which learns to assign dynamic, input-dependent weights to each expert, see Figure 2 for an overview.

Unlike traditional MoE models that provide single-point predictions using conditional mean values, we propose a multi-quantile MoE approach. This proposed model can output several quantiles of the battery’s RUL, denoted as  $Q_{Y|X=x} = (Q_{Y|X=x}(\tau_1), \dots, Q_{Y|X=x}(\tau_K))^T$ , where  $K$  is the number of quantile levels. Thereby, it enables us to construct the entire conditional distribution of the battery’s RUL for a given input  $X = x$ . Hence, the proposed model offers a comprehensive uncertainty quantification for RUL predictions. For your reference, we detail the architectures of the proposed Expert and Gating Networks, along with the methods used to construct the predictive battery’s RUL distributions in the appendices.

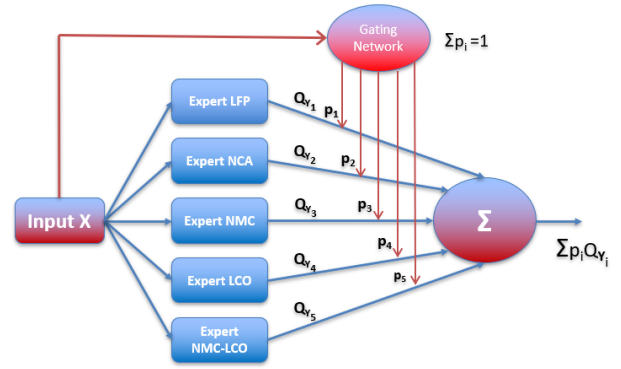


Figure 2: Proposed QMoE architecture for varying types of batteries, where  $Q_{Y_i} = Q_{Y_i|X=x}$  for short.

## 5 Experiments

### 5.1 Datasets

To develop and evaluate the proposed RUL-QMoE model, we integrated seven real-world datasets, as summarized in Table 1. These datasets were utilized for feature extraction, model training, and performance evaluation. For a more detailed description of each dataset, we refer readers to the Appendix.

### 5.2 Feature Engineering

The battery datasets described above are highly heterogeneous, with numerous factors influencing lifespan, such as chemistry, capacity, voltage range, and cycling conditions. While prior research has proposed various feature extraction strategies, these methods are often tailored to specific datasets and lack generalizability across chemistries and operational profiles (Zhang et al. 2024).

In this work, we unify seven publicly available datasets to construct a comprehensive mixed dataset. For raw data preprocessing, we adopt the standardized pipeline provided by (Zhang et al. 2024). To construct the input feature vector  $X$  for model training, we include: one-hot encoded cathode material (covering five types), nominal capacity, minimum and maximum voltage, charge and discharge C-rates, and a sequence of discharge voltage-capacity curves from cycle 10 to cycle 100, denoted as  $V_{d,t}(Q_{d,t})$  for  $t = 10, \dots, 100$ .

To ensure consistency across batteries, we interpolate each discharge curve to 1000 fixed capacity points. This results in a feature vector  $X$  of dimension 1010. Figure 3 illustrates the mean interpolated voltage-capacity curves for each cathode material, highlighting distinct degradation signatures across chemistries.

### 5.3 Implementation of Model Training

We train the proposed RUL-QMoE model in two stages to ensure both specialization and generalization across battery chemistries.

In **Stage 1**, we independently train five Expert Networks, each corresponding to a specific cathode material type: LFP, NCA, NMC, LCO, and NMC-LCO. Each expert is optimized using the Adam optimizer with a learning rate of

Dataset	Electrode Chemistry	Rated Capacity (Ah)	Voltage Range (V)	RUL Distr. (Cycles)	Cell Count
CALCE	LCO/graphite	1.1	2.7 - 4.2	$566 \pm 106$	13
MATR	LFP/graphite	1.1	2.0 - 3.6	$823 \pm 368$	180
HUST	LFP/graphite	1.1	2.0 - 3.6	$1899 \pm 389$	77
HNEI	NMC LCO/graphite	2.8	3.0 - 4.3	$248 \pm 15$	14
RWTH	NMC/carbon	1.11	3.5 - 3.9	$658 \pm 64$	48
SNL	NCA, NMC, LFP/graphite	1.1	2.0 - 3.6	$1256 \pm 1321$	61
UL-PUR	NCA/graphite	3.4	2.7 - 4.2	$209 \pm 50$	10

Table 1: Battery Datasets (Zhang et al. 2024).




Figure 3: Means of discharge voltage-capacity curves at cycle 100 for varying types of batteries.

0.001, a batch size of 256, and 256 hidden units. The dataset is split into 70% for training and 30% for testing, with a fixed random seed of 42 to ensure reproducibility. To prevent overfitting, we employ an Early Stopping strategy with a patience of 20 epochs, halting training when validation performance ceases to improve.

In **Stage 2**, the pre-trained experts are used as warm-start components, and the Gating Network is trained jointly with fine-tuning of the expert parameters. This stage enables the model to learn optimal expert combinations for each input while refining the experts’ predictions. Note that we conducted all training on a free T4 GPU-enabled environment using Google Colab to accelerate computation.

#### 5.4 Evaluation Metrics

To evaluate the performance of the trained RUL-QMoE model, we use the conditional median prediction, denoted as  $Q_{Y|X=x}(0.5)$ , as the representative single-point estimate of the battery’s RUL. We assess prediction accuracy using standard regression metrics, including Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), Root Mean Square Error (RMSE), and the Coefficient of Determination ( $R^2$ ). These metrics provide a comprehensive view of both absolute and relative error, as well as the model’s goodness-of-fit across diverse battery chemistries and datasets.

#### 5.5 Model Performance

The proposed RUL-QMoE model demonstrates strong predictive performance across a range of battery chemistries,

Battery Type	MAE	MAPE	RMSE	$R^2$
LFP	62	7.01%	116	97.85%
NCA	36	13.80%	57	92.26%
NMC	47	8.83%	61	83.83%
LCO	51	8.96%	82	33.22%
NMC-LCO	10	5.20%	12	83.43%
ALL	55	7.82%	98	98.34%

Table 2: Performance of the proposed RUL-QMoE model for varying types of batteries.

with particularly high accuracy observed for NMC-LCO and LFP types as shown in Table 2. The lowest MAE is achieved on the mixed NMC-LCO battery type (MAE = 10, MAPE = 5.2%), indicating highly precise single-point predictions. Moreover, the model achieves its highest coefficient of determination ( $R^2 = 98.34\%$ ) on the ALL dataset (containing all battery’s cathode materials), highlighting its ability to generalize effectively across diverse battery types. In contrast, the model’s performance on LCO batteries is comparatively weaker, with the lowest  $R^2$  value of 33.22%. This discrepancy could be due to a small sample size of the LCO dataset, and thus, insufficient feature representation for this chemistry. However, the model maintains robust performance across most battery types, making it a promising tool for RUL prediction in heterogeneous battery applications.

As mentioned, one notable advantage of our proposed model is its ability to provide uncertainty quantification, such as prediction intervals (PI) and survival probabilities. For instance, Figure 4 illustrates the probabilistic prediction capabilities for a specific battery, with a predicted median RUL of 1632 cycles, and a 90% PI ranging from 1565 to 1681 cycles, which covers the ground truth RUL of 1599 cycles. Furthermore, we can also perform survival analysis at a specific cycle, for instance, at cycle 100. If the model estimates a survival probability of  $\mathbb{P}(\text{RUL} > 1600) = 0.8$ , this indicates that the battery has an 80% chance of operating for at least an additional 1600 cycles beyond cycle 100.

#### 5.6 Comparisons with Other Models

To benchmark the performance of the proposed RUL-QMoE model, we shall test and compare our proposed model against several leading ML models in (Zhang et al. 2024) across eight battery datasets using cycling data at cycle 100 only. As can be seen from Table 3, the best performance is observed on the HUST dataset, which achieves the lowest MAE of 47 (cycles) and MAPE of 2.56%. Similarly, the




Figure 4: Probabilistic battery’s RUL Predictions.

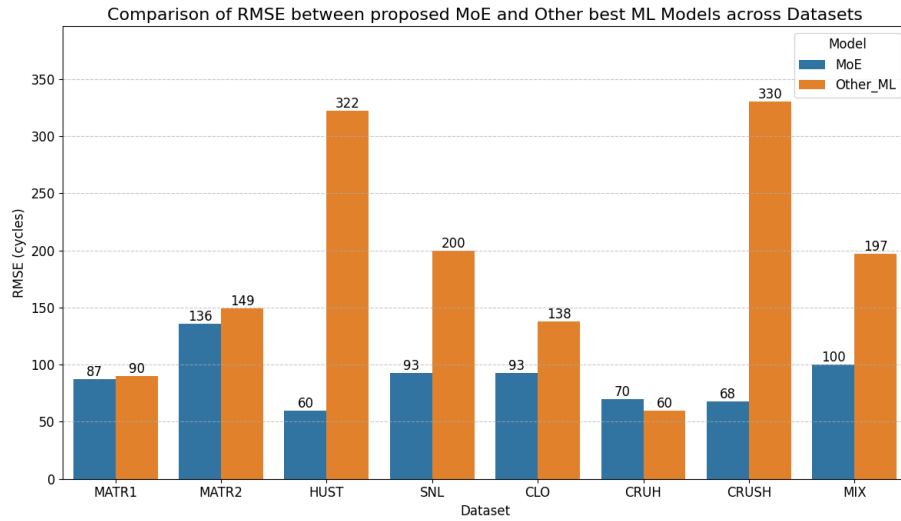


Figure 5: Performance comparison between the proposed RUL-QMoE and other best ML models in (Zhang et al. 2024) for RUL predictions across many testing datasets.

CRUSH dataset shows excellent results with a low MAE of 49 (cycles). In contrast, the MATR2 dataset presents the greatest challenge, with the highest MAE of 98 (cycles), RMSE of 136 (cycles), and the lowest  $R^2$  of 80.13%. This suggests that the dataset may contain greater variability or noise, making it more difficult for the model to learn reliable degradation patterns. For the MIX dataset, which combines all battery datasets, the model maintains strong performance, with MAE = 65 (cycles), RMSE = 100 (cycles), and  $R^2 = 96.84\%$ , validating its scalability and robustness across chemistries and operational conditions. These results confirm that the proposed RUL-QMoE is capable of delivering accurate RUL predictions from early-cycle data, making it a valuable tool for early-stage battery health assessment and predictive maintenance.

As shown in Figure 5, the proposed RUL-QMoE consistently achieves lower RMSE values, compared to best ML models implemented in (Zhang et al. 2024). Specifically, on the HUST dataset, the proposed RUL-QMoE delivers a dramatic improvement with an RMSE of 60 cycles,

Battery Data	MAE	MAPE	RMSE	$R^2$
MATR1	53	8.05%	87	94.93%
MATR2	98	10.19%	136	80.13%
HUST	47	2.56%	60	97.93%
SNL	67	9.94%	93	99.44%
CLO	65	8.92%	93	92.12%
CRUH	53	12.64%	70	84.55%
CRUSH	49	10.61%	68	99.55%
MIX	65	9.59%	100	96.84%

Table 3: Performance of the proposed RUL-QMoE model for several datasets of batteries using data at cycle 100 only.

compared to 322 cycles from the best ML model, namely, "Discharge model" in (Zhang et al. 2024). For CRUSH, the model achieves an RMSE of 100 cycles, significantly outperforming the XGBoost model in (Zhang et al. 2024) which yields an RMSE of 330 cycles. On CLO, RUL-QMoE reduces RMSE to 70 cycles versus 138 cycles from competing

“Full” model. For the MIX dataset, which combine all battery datasets, the RUL-QMoE model achieves an RMSE of 100 cycles - nearly halving the error compared to 197 cycles from the Random Forest model. A similar trend is observed on the SNL dataset, where RUL-QMoE reduces the RMSE to 93 cycles, significantly outperforming the Principal Component Regression (PCR) model, which records an RMSE of 200 cycles. Also, moderate improvements are also observed on MATR1, and MATR2, with RMSE reductions of 3 and 15 cycles, respectively. However, on the CRUH dataset, Partial Least Squares Regression (PLSR) slightly outperforms the proposed RUL-QMoE model, achieving a lower RMSE of 60 cycles compared to 70 cycles.

These results underscore the effectiveness of the RUL-QMoE framework in delivering accurate and generalizable RUL predictions across diverse battery chemistries and operational conditions. Besides, its consistent performance advantage over existing ML models validates the strength of its probabilistic, expert-driven architecture.

## 6 Path to Deployment

The proposed RUL-QMoE framework presents a significant advancement in probabilistic battery health and prognostics, with clear pathways for real-world deployment across applications reliant on lithium-ion batteries. Its architecture, performance, and uncertainty quantification capabilities align well with the operational needs in existing Li-ion battery EV systems. We outline the key deployment considerations and strategies below:

- **Seamless Integration with Battery Management Systems (BMS)**

For EVs and portable devices which has limited computational resources for BMS, RUL-QMoE can be compressed and deployed on edge devices and use early-cycle data from the BMS to evaluate the model and minimize the computational latency, enabling real-time and accurate RUL updates during operation. What is more, in large-scale storage systems such as grid storage, lightweight gating networks can be trained locally to assign experts, while computationally-intensive tasks like quantile predictions are offloaded to the cloud.

- **Uncertainty-Aware Maintenance and Operations**

The probabilistic outputs of RUL-QMoE, including survival probabilities, non-crossing prediction intervals, and full conditional distributions - transform uncertainty quantification into actionable intelligence for risk-informed maintenance. By dynamically estimating the survival probability of a battery at any cycle, the model enables adaptive replacement triggers based on user-defined risk thresholds, further supports conservative maintenance scheduling, ensuring replacements occur before the lower prediction interval bound to prevent failures while optimizing spare-parts logistics. Integration with industrial IoT platforms automates alerting and work-order generation in Computerized Maintenance Management Systems when survival probabilities reach operational limits or observed degradation deviates beyond prediction intervals.

- **Adaptability to Emerging Chemistries and Data Streams**

The MoE architecture supports seamless addition of new expert networks for emerging battery chemistries. Incremental fine-tuning with federated learning can preserve data privacy across manufacturers. Also, the pretrained experts can accelerate adaptation to new datasets with limited samples, reducing data acquisition costs.

RUL-QMoE bridges the gap between academic research and industrial deployment. Its chemistry-specific expertise, probabilistic outputs, and adaptability position it as a cornerstone for next-generation battery health management. Pilot deployments in collaboration with automotive and energy partners will be the critical next step toward commercialization.

## 7 Conclusion and Future Work

We presented RUL-QMoE, a novel probabilistic Mixture-of-Experts framework for estimating the Remaining Useful Life (RUL) of lithium-ion batteries across diverse chemistries and operational conditions. Our approach integrates specialized expert networks aligned with distinct cathode materials and enforces non-crossing quantile constraints to produce coherent, interpretable probabilistic forecasts. A dynamic gating mechanism enables adaptive relevance assignment, allowing the model to generalize effectively across heterogeneous datasets. Comprehensive evaluations on seven publicly available battery datasets demonstrate that RUL-QMoE consistently surpasses state-of-the-art baselines in RMSE, MAE, and  $R^2$ , with notable gains on mixed-chemistry and early-cycle data. Beyond predictive accuracy, the model supports survival analysis and prediction interval estimation - critical capabilities for risk-aware battery management and maintenance planning.

RUL-QMoE offers a scalable, data-efficient foundation for battery health prognostics, with strong potential for real-world deployment in electric vehicles, grid storage systems, and portable electronics. It serves as a core component of the Battery Foundation Model currently under development in collaboration with Toyota Motor Engineering & Manufacturing North America, Inc., and has been integrated into our Battery AI platform during Stage 1 of the project.

In Stage 2, we aim to extend the framework to monitor additional battery states, including State-of-Health (SoH) and State-of-Charge (SoC), and to support downstream applications such as estimating EV driving range under varying conditions (e.g., weather, traffic, battery specifications, vehicle brand, RUL, SoH, SoC, etc.). These enhancements will further position RUL-QMoE as a key enabler of intelligent, context-aware energy systems, advancing the deployment of AI in sustainable mobility and energy infrastructure.

## Acknowledgments

This research is supported by the Agency for Science, Technology and Research (A\*STAR), Singapore under its project M23M6c0114, and Toyota Motor Engineering & Manufacturing North America, Inc.

## References

- Catelani, M.; Ciani, L.; Fantacci, R.; Patrizi, G.; and Picano, B. 2021. Remaining useful life estimation for prognostics of lithium-ion batteries based on recurrent neural network. *IEEE Transactions on Instrumentation and Measurement*, 70: 1–11.
- Chen, D.; and Zhou, X. 2024. AttMoE: Attention with Mixture of Experts for remaining useful life prediction of lithium-ion batteries. *Journal of Energy Storage*, 84: 110780.
- Devie, A.; Baure, G.; and Dubarry, M. 2018. Intrinsic variability in the degradation of a batch of commercial 18650 lithium-ion cells. *Energies*, 11(5): 1031.
- Hasib, S. A.; Islam, S.; Chakraborty, R. K.; Ryan, M. J.; Saha, D. K.; Ahamed, M. H.; Moyeen, S. I.; Das, S. K.; Ali, M. F.; Islam, M. R.; et al. 2021. A comprehensive review of available battery datasets, RUL prediction approaches, and advanced battery management. *Ieee Access*, 9: 86166–86193.
- He, W.; Williard, N.; Osterman, M.; and Pecht, M. 2011. Prognostics of lithium-ion batteries based on Dempster-Shafer theory and the Bayesian Monte Carlo method. *Journal of Power Sources*, 196(23): 10314–10321.
- Hong, J.; Lee, D.; Jeong, E.-R.; and Yi, Y. 2020. Towards the swift prediction of the remaining useful life of lithium-ion batteries with end-to-end deep learning. *Applied energy*, 278: 115646.
- Juarez-Robles, D.; Azam, S.; Jeevarajan, J. A.; and Mukherjee, P. P. 2021. Degradation-safety analytics in lithium-ion cells and modules: Part III. Aging and safety of pouch format cells. *Journal of The Electrochemical Society*, 168(11): 110501.
- Juarez-Robles, D.; Jeevarajan, J. A.; and Mukherjee, P. P. 2020. Degradation-safety analytics in lithium-ion cells: Part I. Aging under charge/discharge cycling. *Journal of The Electrochemical Society*, 167(16): 160510.
- Koenker, R.; and Bassett Jr, G. 1978. Regression quantiles. *Econometrica: journal of the Econometric Society*, 33–50.
- Lei, Y.; and Yu, G. 2025. A multi-scale lithium-ion battery capacity prediction using mixture of experts and patch-based MLP. *arXiv preprint arXiv:2504.03706*.
- Li, W.; Sengupta, N.; Dechent, P.; Howey, D.; Annaswamy, A.; and Sauer, D. U. 2021. One-shot battery degradation trajectory prediction with deep learning. *Journal of Power Sources*, 506: 230024.
- Ma, G.; Xu, S.; Jiang, B.; Cheng, C.; Yang, X.; Shen, Y.; Yang, T.; Huang, Y.; Ding, H.; and Yuan, Y. 2022. Real-time personalized health status prediction of lithium-ion batteries using deep transfer learning. *Energy & Environmental Science*, 15(10): 4083–4094.
- Preger, Y.; Barkholtz, H. M.; Fresquez, A.; Campbell, D. L.; Juba, B. W.; Romàn-Kustas, J.; Ferreira, S. R.; and Chalamala, B. 2020. Degradation of commercial lithium-ion cells as a function of chemistry and cycling conditions. *Journal of The Electrochemical Society*, 167(12): 120532.
- Severson, K. A.; Attia, P. M.; Jin, N.; Perkins, N.; Jiang, B.; Yang, Z.; Chen, M. H.; Aykol, M.; Herring, P. K.; Fragedakis, D.; et al. 2019. Data-driven prediction of battery cycle life before capacity degradation. *Nature Energy*, 4(5): 383–391.
- Wang, X.; Ye, P.; Liu, S.; Zhu, Y.; Deng, Y.; Yuan, Y.; and Ni, H. 2023. Research progress of battery life prediction methods based on physical model. *Energies*, 16(9): 3858.
- Xie, J.; Weng, Y.; and Nguyen, H. D. 2022. Health-informed lifespan-oriented circular economic operation of li-ion batteries. *IEEE Transactions on Industrial Informatics*, 19(3): 2749–2760.
- Xie, J.; Yang, R.; Hui, S.-Y. R.; and Nguyen, H. D. 2024. Dual digital twin: Cloud–edge collaboration with Lyapunov-based incremental learning in EV batteries. *Applied Energy*, 355: 122237.
- Xing, Y.; Ma, E. W.; Tsui, K.-L.; and Pecht, M. 2013. An ensemble model for predicting the remaining useful performance of lithium-ion batteries. *Microelectronics reliability*, 53(6): 811–820.
- Yang, R.; and Nguyen, H. D. 2025. Temperature distribution learning of Li-ion batteries using knowledge distillation and self-adaptive models. *Applied Energy*, 382: 125196.
- Yang, Y. 2021. A machine-learning prediction method of lithium-ion battery life based on charge process for different applications. *Applied Energy*, 292: 116897.
- Zhang, H.; Gui, X.; Zheng, S.; Lu, Z.; Li, Y.; and Bian, J. 2024. BatteryML: An Open-source Platform for Machine Learning on Battery Degradation. In *The Twelfth International Conference on Learning Representations*.
- Zhao, L.; Song, S.; Wang, P.; Wang, C.; Wang, J.; and Guo, M. 2024. A MLP-Mixer and mixture of expert model for remaining useful life prediction of lithium-ion batteries. *Frontiers of Computer Science*, 18(5): 185329.