IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 16 August 2026, accepted 18 September 2026, date of publication 25 September 2026,
date of current version 22 October 2026.

Digital Object Identifier 10.1109/ACCESS.2026.3737729

RESEARCH ARTICLE

Cross-Chemistry Meta-Learning With Frequency-Domain Augmentation for Robust Capacity Retention Prediction in Vanadium Redox Flow Batteries

YOUNG-DAE SHIM[1], JUNGHWAN YIM[1], NIAN LIU[2], CHEONG HOON KWON[3], HAE-JIN CHOI[4], AND SEUNG-KYUM CHOI[1]

[1]George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA
[2]School of Chemical and Biomolecular Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA
[3]Department of Energy Resources and Chemical Engineering, Kangwon National University, Samcheok 25913, Republic of Korea
[4]School of Mechanical Engineering, Chung-Ang University, Seoul 06974, Republic of Korea

Corresponding authors: Seung-Kyum Choi (schoi@me.gatech.edu) and Hae-Jin Choi (hjchoi@cau.ac.kr)

This research was funded by the Ministry of Trade, Industry and Energy (MOTIE), Korea, through the Global Industrial Technology Cooperation Center (GITCC) Program, supervised by the Korea Institute for Advancement of Technology (KIAT) (Task No. P246800264).

ABSTRACT Vanadium redox flow batteries (VRFBs) are promising for long-duration energy storage, but reliable capacity-retention forecasting is difficult because available VRFB degradation data are scarce and operating protocols vary widely. To address this domain-mismatch problem, this article proposes a cross-chemistry meta-learning framework that meta-trains on lithium-ion battery degradation data and adapts to VRFB capacity trajectories using only early-cycle observations. A cycle-based frequency-domain data augmentation approach based on CEEMDAN decomposition is introduced to increase trajectory diversity while preserving the global degradation trend; the four NASA lithium-ion cell datasets are augmented to 4,000 synthetic trajectories for meta-training. The target domain consists of VRFB datasets collected from prior studies under heterogeneous current densities (100, 120, 160, and 200 mA/cm^2), and the same augmentation procedure is applied. A weighted ensemble of five MLP predictors, trained with Bayesian hyperparameter optimization, is integrated into the meta-learning pipeline to further enhance robustness. Across all target VRFB datasets, the proposed method (Meta-Ensemble-Aug) achieves strong generalization, outperforming conventional models trained on single-dataset baselines.

INDEX TERMS Battery capacity, meta-learning, predictive maintenance, vanadium-redox flow battery.

I. INTRODUCTION

In recent years, the rapid growth of artificial intelligence (AI) and data centers has led to an explosive increase in electricity consumption [1], [2], [3], [4]. To meet this rising demand, fossil-fuel-based and nuclear power generation remain widely used; however, growing environmental concerns about CO_2 emissions and radioactive waste have drawn global attention to renewable energy sources. Nevertheless, solar and wind power are inherently intermittent because their output fluctuates with weather conditions and the time of day. This mismatch between actual electricity consumption patterns and the timing of power generation underscores the increasing importance of energy storage systems (ESSs), which can store surplus energy and supply it when needed [5], [6], [7].

Most commercially available ESSs are currently based on lithium-ion batteries [8], [9], [10]. Lithium-ion batteries offer high energy density and fast response characteristics, and advances in mass-production technologies have reduced their cost, enabling widespread deployment in applications such as industrial backup power and peak-load shaving.

The associate editor coordinating the review of this manuscript and approving it for publication was Arman Oshnoei.

© 2026 The Authors. This work is licensed under a Creative Commons Attribution 4.0 License.
For more information, see https://creativecommons.org/licenses/by/4.0/

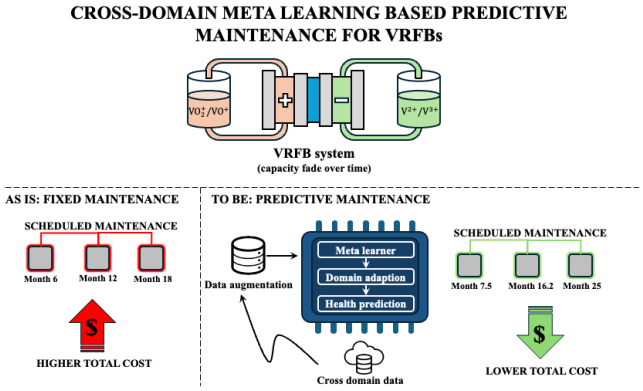


FIGURE 1. Meta-learning framework for VRFB capacity retention rate prediction predictive maintenance.

However, because both the electrochemical reactions and energy storage occur within the electrodes, it is difficult to scale energy and power independently. In addition, thermal runaway and electrode degradation can occur during long-term operation, limiting their suitability for long-duration energy storage [11], [12], [13], [14]. To address these limitations, vanadium redox flow batteries (VRFBs), which use redox reactions of vanadium ions in the electrolyte, have emerged as promising next-generation technologies for large-scale, long-duration energy storage [15], [16], [17].

The VRFB is an electrochemical energy storage system that stores and releases energy through the redox reactions of vanadium ions (V^{2+}/V^{3+} and VO^{2+}/VO_2^+) in the electrolyte. Unlike lithium-ion batteries, in which electrochemical reactions occur within solid electrodes, VRFBs exchange energy through redox reactions in liquid electrolytes. Because both positive and negative electrolytes are based on the same metal ion (vanadium), chemical instability caused by crossover through the membrane is reduced, the risk of thermal runaway is mitigated, and stable long-term operation is enabled. However, during extended cycling, side reactions associated with electrolyte imbalance, ion crossover, and the hydrogen evolution reaction (HER) accumulate, leading to a gradual decrease in usable capacity. Electrolyte imbalance arises from asymmetric proton (H^+) transport during overcharging or repeated charge–discharge cycles, which establishes concentration gradients and drives vanadium ions to diffuse across the membrane, where a fraction become electrochemically inactive. These phenomena ultimately result in reduced coulombic efficiency and cause long-term capacity fade [18], [19].

Because this capacity fade arises primarily from chemical imbalances in the electrolyte rather than from structural degradation of the electrodes, it can be partially recovered through electrolyte rebalancing or by adding makeup electrolyte. However, these procedures require system shutdown and incur additional costs; therefore, optimizing the timing and frequency of rebalancing is a critical challenge for ensuring reliable and cost-effective long-term operation. In current practice, rebalancing is often performed at fixed intervals (e.g., every ∼100 cycles), which can result in unnecessary maintenance when no significant electrolyte imbalance has developed. Accordingly, AI-based real-time state prediction is needed to quantitatively forecast capacity changes, determine the optimal timing for electrolyte rebalancing and makeup electrolyte supplementation, and support maintenance decision-making.

Recently, early prediction of battery capacity and lifetime based on signals from only a few initial cycles has become a standard paradigm for full-lifecycle prognosis [20]. Building on conventional machine learning methods and architectures such as long short-term memory (LSTM), gated recurrent unit (GRU), and convolutional neural networks (CNNs), transformer-based time-series models and transfer learning have rapidly gained traction [21], [22], [23], [24]. To address challenges arising from limited labels and domain heterogeneity, self-supervised and semi-supervised learning approaches have been introduced [25], while physics-informed modeling and uncertainty quantification have been emphasized to support operational decision-making in real-world systems [26], [27]. In signal based prognostics, electrochemical impedance spectroscopy (EIS) and spectral measurements have been successfully used to estimate capacity under varying conditions and to support SOC/SOH estimation [28], [29], and, in the context of VRFBs, studies using model-based state estimation and rebalancing optimization have also been reported [30], [31]. However, many of these state-of-the-art models exhibit strong predictive performance only for single cells under fixed cycling protocols, while their accuracy degrades significantly when the chemistry, protocol, temperature, or manufacturing lot changes, revealing clear limitations in generalization. To overcome this limitation, this study propose a meta-learning approach that uses lithium-ion batteries as the source domain and predicts future VRFB capacity rentention trajectories by fine-tuning on early-cycle VFRB data.

Despite recent advances in battery prognostics, developing a generalized prediction model for VRFB capacity retention remains challenging because publicly available VRFB degradation datasets are extremely limited and exhibit substantial heterogeneity in electrochemical systems and operating conditions. This work addresses the Different-Component, Different-Working-Condition (DCDW) scenario, where both the battery chemistry and operating protocols differ between the source and target domains. Under this setting, conventional deep-learning and transfer-learning methods often fail to generalize. Therefore, the key technical challenge is to develop a learning framework capable of transferring degradation knowledge across fundamentally different battery chemistries while maintaining robust prediction performance using only limited early-cycle VRFB observations.

In this study, we develop a cross-chemistry meta-learning framework to predict the capacity evolution of VRFBs. The framework is designed to meta-train on lithium-ion battery data as the source domain and then predict future

capacity-retention trajectories of VRFBs by fine-tuning on their early-cycle data. For meta-training, datasets from four NASA 18650 cells tested under room-temperature conditions are used, and data diversity is enhanced through frequency-domain augmentation. In the training stage, five individual models are pretrained on the source domain, and their weights are optimally combined using an ensemble strategy. Each model is based on a three-layer Multi-Layer Perceptron (MLP) architecture, and its hyperparameters are tuned using a Bayesian optimization. The framework is then fine-tuned on early cycle VRFB data following the same procedure as in the source domain. The final model is used to predict future capacity-retention trajectories of VFRBs and validate the generalization performance.

The main contributions of this work are summarized as follows:

(1) Propose a cross-chemistry meta-learning framework for transferring degradation knowledge from lithium-ion batteries to VRFBs under the DCDW scenario.

(2) Introduce a CEEMDAN-based frequency-domain augmentation strategy that increases trajectory diversity while preserving global degradation trends.

(3) Integrate a jointly trained softmax-gated ensemble into the meta-learning framework to improve robustness across heterogeneous datasets.

(4) Validate the framework on four independent VRFB datasets and demonstrate improved cross-dataset generalization compared with conventional deep learning approaches.

This paper is organized as follows. Section II describes the mechanisms and causes of capacity fade in vanadium redox flow batteries. Section III presents the overall meta-learning framework. Sections IV and V detail the experimental setup and discuss the corresponding results, respectively. Finally, Section VI concludes the paper.

II. RELATED WORK

A. DEEP-LEARNING-BASED BATTERY CAPACITY AND RUL PREDICTION

Early-cycle prediction of battery remaining useful life (RUL) has been dominated by recurrent and convolutional architectures. Severson et al. [20] established a landmark result by predicting Li-ion lifetime from the first 100 cycles using elastic-net regression on discharge-voltage features. Subsequent studies have replaced hand-crafted features with sequence models: LSTM and GRU networks have been applied to capacity-fade curves [32], [33], and CNN-based approaches have extracted degradation-related patterns directly from raw voltage–current trajectories [34], [35]. More recently, Transformer-based time-series models [36], [37] have achieved strong performance on battery prognostics benchmarks. However, these models exhibit severe accuracy degradation when the chemistry, protocol, temperature, or manufacturing lot differs from those seen during training, motivating transfer-oriented approaches.

B. META-LEARNING AND FEW-SHOT ADAPTATION

Meta-learning, or "learning to learn," aims to produce models that adapt rapidly to new tasks from limited data. Model-Agnostic Meta-Learning (MAML) [38] achieves this by back-propagating through a small number of inner-loop adaptation steps, whereas Reptile [39] achieves a comparable effect with a first-order update that simply moves the meta-parameters toward the parameters obtained after task-specific adaptation, avoiding back-propagation through the inner loop.

Applications of meta-learning to battery health have begun to appear, using MAML-style adaptation to estimate state of health from small samples and to transfer across operating conditions or measurement domains [40], [41]. These studies confirm that few-shot adaptation is effective for degradation data, but the tasks over which they meta-train are drawn from a single electrochemistry — typically variants of lithium-ion cells differing in temperature, protocol or manufacturing lot. The task distribution shares a common degradation process, and the adaptation problem can therefore be viewed as a covariate shift. However, the source and target differ in the active species, the degradation mechanism, and the reversibility of capacity loss. The common structure available for meta-learning is thus not tied to a specific physical mechanism, but rather to the form of the degradation trajectory. Whether meta-learning can effectively transfer across these differences remains, to our knowledge, an open question that we investigate in this work.

C. CAPACITY FADE AND HEALTH PREDICTION IN VANADIUM REDOX FLOW BATTERIES

Research on VRFB degradation has been dominated by mechanistic modelling, reflecting the fact that capacity fade in a flow battery is a property of the electrolyte rather than of the electrode structure. Loktionov et al. [19] quantified the separate contributions of faradaic imbalance and vanadium-ion crossover to capacity loss, establishing that the two mechanisms act on different timescales. Ma et al. [42] evaluated the effect of the hydrogen evolution reaction on cell performance, and Huang et al. [43] experimentally validated the role of side reactions in capacity fade, together identifying the oxidation-state imbalance that this paper takes as its physical premise. Zou et al. [18] developed a capacity-fade prediction model for long-term VRFB operation on this mechanistic basis. A parallel body of work attacks the problem at the materials level: Toja et al. [44] mitigated capacity decay and volume variation through modified electrolyte preparation, and Wang et al. [45] suppressed decay using an electrolyte with elevated average vanadium valence, while Morozov et al. [46] optimized the flow factor as a control variable. These studies explain and mitigate fade, but they require species-level parameters — crossover coefficients, per-species concentrations, membrane transport properties — that are rarely observable in a deployed stack, and they do

not provide a forecast of capacity retention from early-cycle measurements.

Data-driven VRFB studies are comparatively few. Zhao et al. [30] estimated state of charge and capacity in real time using an unscented Kalman filter with a forgetting factor, and Zhao et al. [31] improved SOC estimation through multi-source data fusion. Both are state-estimation methods: they infer the present condition of the cell from current measurements rather than forecasting its future trajectory. The closest prior work to ours is Zhao and Jung [47], who forecast long-term VRFB capacity fade by combining Gaussian process regression with an Informer sequence model, and reported low average errors across several current densities. That study establishes that data-driven long-horizon forecasting is feasible for VRFBs, but it is trained on VRFB data alone. Consequently its accuracy is bounded by the size of the VRFB corpus, a model must be re-fitted for each operating condition, and the far larger body of Li-ion degradation data is left unused. The present work is complementary in exactly this respect: it asks whether degradation priors learned from an abundant, mechanistically different source chemistry can reduce the target-domain data requirement for VRFB forecasting.

D. TRANSFER LEARNING FOR CROSS-CHEMISTRY BATTERY MODELING META-LEARNING AND FEW-SHOT ADAPTATION

Transfer learning has been applied to bridge related battery chemistries. Pretrain-and-fine-tune approaches initialize weights on a large source dataset and adapt them to a target dataset [48], whereas distribution-matching methods, such as MMD and CORAL align source and target feature statistics [49], [50]. Within battery health estimation, these techniques have been used to transfer models across ambient temperatures, charge-discharge protocols and cathode materials [51], [52].

However, most existing studies focus on homogeneous lithium-ion datasets and do not address electrochemical transfer between fundamentally different systems such as Li-ion batteries and VRFBs. The unrelated degradation processes of those two batteries forces statistical alignment between them, which risks negative transfer. Also, pre-train and fine-tune optimizes the source objective only; nothing in the procedure encourages an initialization from which adaptation is fast. Meta-learning optimizes precisely that property. This is tested empirically and reported in Section VI, a conventional pre-train and fine-tune baseline recovers only part of the improvement obtained by the full meta-learning formulation, and the ablation in the Appendix isolates the meta-learning contribution from those of augmentation and ensembling.

E. RESEARCH GAP AND CONTRIBUTION

Table 1 summarizes the studies discussed above that matters for the present problem. The literatures can be divided into three groups, which none addresses the target setting. The first trains battery capacity on single chemistry and achieves high accuracy within same working condition, but degrades sharply when the protocol or the testing environment changes. The second models VRFB degradation directly, either based on physical domain knowledge or data driven method. The third provides adaption method — meta-learning and domain adaption — but has been demonstrated within single chemistry.

To our knowledge, no prior work combines (i) meta-learning for rapid few-shot adaptation, (ii) frequency-domain augmentation that preserves the global degradation trend, and (iii) cross-chemistry transfer from Li-ion batteries to VRFBs. This combination is the central contribution of this paper.

In summary, existing work either (a) trains on large single-chemistry datasets but generalizes poorly across chemistries and operating conditions, (b) uses meta-learning within a single chemistry, or (c) uses classical transfer learning, which, as shown empirically in Section IV, recovers only part of the benefit of the full meta-learning formulation. The proposed cross-chemistry Reptile framework, augmented by CEEMDAN-based frequency-domain data augmentation and a softmax-gated MLP ensemble, addresses this gap and demonstrates that Li-ion degradation priors can be transferred to VRFB capacity-retention forecasting under heterogeneous current densities.

III. MECHANISM AND DEGRADATION OF VRFBS

A VRFB is an energy storage system that uses the redox reactions of vanadium ions to store and release electrical energy. The overall system architecture, including the electrolyte tanks, pumps, and cell stack, is illustrated in Fig. 1. The electrolyte is stored in external tanks and circulated through the cell stack by pumps. Because the electrochemical reactions occur in the liquid phase and the electrolyte volume and the cell-stack size can be designed independently, the energy capacity and power rating can be decoupled. Owing to this structural advantage, VRFBs are well suited for large-scale, long-duration energy storage applications and can provide the extended storage times and high operational reliability required in renewable-energy-based power grids.

In a VRFB, both the charging and discharging processes proceed through changes in the oxidation states of vanadium ions. The overall electrochemical reactions are given as follows [15], [46].

$$VO_2^+ + 2H^+ + e^- \leftrightarrow VO^{2+} + H_2O \tag{1a}$$

$$V^{3+} + e^- \leftrightarrow V^{2+} \tag{1b}$$

Because the same metal-ion species is used in both half-cells, the risk of chemical contamination from ion crossover is significantly lower than in lithium-based batteries, and the absence of thermal runaway enables inherently high operational safety. In addition, the external electrolyte storage configuration enables complete discharge without damage, and the system capacity can be partially recovered during

TABLE 1. Summary of related work and positioning of the proposed framework.

Category	Studies	Chemistry	Method	Target	Limitation
Feature-based early-life prognosis	Severson et al. [20]	Li-ion (LFP, 124 cells)	Elastic-net on early-cycle discharge-voltage features	Cycle life	Requires a large, protocol-consistent cell population; hand-crafted features do not transfer to a different electrochemistry
Recurrent sequence models	Zhang et al. [32]; Li et al.[33]; Liu et al.[34]; Chang et al. [35]	Li-ion (NASA, CALCE)	LSTM / variant-LSTM / GRU on capacity-fade sequences	SOH, RUL	Trained and tested within one dataset; accuracy collapses when the protocol or chemistry changes
Hybrid and decomposition-assisted models	Fan et al.[22]; Wang et al.[23]; Huang et al. [24]	Li-ion	Signal decomposition (EMD / CEEMDAN family) combined with GRU, TCA or SVR	RUL	Decomposition is used only as denoising pre-processing on a single trajectory; it does not enlarge the effective sample pool
Attention and Transformer models	Zhao and Behdad[21]; Suh et al.[37]; Zhao et al. [36]	Li-ion (EV, NASA)	Transformer, spatio-temporal attention, MLP-Mixer with mixture-of-experts	SOH, RUL	Data-hungry; the parameter count is incompatible with the tens of trajectories available for VRFBs
Physics-informed and hybrid-physics models	Navidi et al. [26]; Wang et al.[27]; Li et al. [28]	Li-ion	PINN and physics-guided diagnostics with degradation-mode constraints	Degradation modes, lifetime	The embedded physics is Li-ion-specific (SEI growth, Li plating) and does not describe electrolyte-imbalance fade in VRFBs
Limited-data Li-ion frameworks	Che et al.[25]; Sudarshan et al. [29]	Li-ion	Self-supervised pre-training; scalable diagnosis from limited data	SOH	Source and target remain the same chemistry; the domain gap is one of scale, not of degradation mechanism
VRFB mechanistic capacity-fade modelling	Zou et al.[18]; Loktionov et al.[19]; Ma et al.[42]; Huang et al. [43]	VRFB	Zero-dimensional / electrochemical models of crossover, faradaic imbalance and HER	Capacity fade	Requires species-level parameters that are rarely measurable in a deployed stack; not a forecasting model from early-cycle data
VRFB state estimation	Zhao et al.[30]; Zhao et al. [31]	VRFB	Unscented Kalman filter with forgetting factor; multi-sensor data fusion	SOC, capacity	Estimates the present state rather than forecasting future retention; tuned per cell
VRFB data-driven fade forecasting	Zhao and Jung [47]	VRFB	Gaussian process regression combined with an Informer model	Long-term capacity fade	Fitted per dataset on VRFB data alone; no mechanism to exploit the far larger Li-ion degradation corpus
Meta-learning	Finn et al.[38]; Nichol and Schulman [39]	Generic (vision, RL)	MAML; first-order Reptile	Fast adaptation	Not instantiated for battery degradation; no treatment of source-domain scarcity
Transfer learning / domain adaptation	Tan et al.[48]; Long et al.[49]; Sun and Saenko [50]	Generic; Li-ion variants	Pre-train and fine-tune; JAN (MMD); Deep CORAL	Feature alignment	Assumes a shared feature-generating process; degrades when the source and target degradation physics differ, as between Li-ion and VRFB
This work	Proposed framework	Li-ion source (NASA) to VRFB target	Reptile meta-learning + CEEMDAN frequency-domain augmentation + softmax-gated MLP ensemble	Capacity retention	Addresses cross-chemistry transfer under target-domain scarcity across four current densities

long-term operation through electrolyte rebalancing or by adding make-up electrolyte.

However, the capacity of VRFBs gradually decreases with repeated charge–discharge cycling because of changes in electrolyte composition. This degradation is mainly caused by ionic imbalance within the electrolytes, vanadium-ion crossover through the membrane, and the hydrogen evolution reaction. During cycling, accumulated proton transport coupled with changes in the oxidation states of vanadium ions, alters the redox-state distribution between the positive and negative electrolytes, thereby reducing the total amount of electrochemically active species. In addition, when vanadium ions cross the ion-exchange membrane into the opposite electrolyte, irreversible reactions among oxidation states can occur, forming inactive species and progressively exacerbating the concentration imbalance during long-term operation.

On the negative side, when charging continues after V^{2+} ions have been sufficiently reduced, the HER becomes dominant. This process can be expressed as follows [42], [43]:

$$2H^+ + 2e^- \rightarrow H_2. \tag{2}$$

This reaction diverts a portion of the charging current to hydrogen gas generation, thereby reducing coulombic efficiency and altering the electrolyte composition, which accelerates capacity loss. Consequently, unlike lithium-ion batteries, in which capacity fade is primarily driven by electrode material degradation, capacity decay in VRFBs is predominantly governed by changes in electrolyte composition and irreversible ion-transport processes.

Considering these characteristics, the capacity of VRFBs can be partially restored through electrolyte rebalancing or by adding makeup electrolyte to correct electrolyte imbalances, which is important for ensuring long-term system reliability and economic viability. However, because rebalancing requires system downtime and incurs additional costs, accurately diagnosing capacity fade and determining the optimal timing of such interventions are essential. Therefore, for the long-term optimal operation of VRFBs, a data-driven approach capable of precisely predicting and managing electrolyte-state evolution and capacity degradation is required. This study proposes a meta-learning-based capacity-prediction framework to address this need.

IV. METHODOLOGY

Unlike conventional transfer learning, which directly fine-tunes a pretrained model on the target domain, the proposed framework is specifically designed for cross-chemistry adaptation under severe domain mismatch. The methodological novelty lies in jointly combining Reptile-based meta-learning, CEEMDAN-based frequency-domain augmentation, and a learnable weighted ensemble into a unified framework for VRFB capacity-retention prediction.

TABLE 2. Summary of notation.

Symbol	Meaning
Meta-learning (Sec. III-B)	
τ_i	i-th battery task
τ^{sup}, τ^{qry}	Support and query subsets of a task
θ_0	Meta-parameters (initialization learned on source)
$\theta^{(k)}$	Parameters of the k-th ensemble MLP
θ'	Task-adapted parameters after inner loop
T	Number of inner-loop gradient steps
S	Support-set size (number of cycles)
α	Inner-loop (task adaptation) learning rate
β	Outer-loop (meta-update) learning rate
B	Meta-batch size (tasks per outer step)
E	Number of meta-epochs
$\mathcal{L}$	Loss function (MSE on retention)
Model & Ensemble (Sec. III-C)	
$x \in R^6$	Input feature vector per cycle
k	Cycle index
μ_r, σ_r	Rolling mean and standard deviation of retention
$min_r, \max_r$	Rolling min and max of retention
Δ_r	First-difference (slope) of retention
W	Rolling window size ($=5$)
K	Number of ensemble members ($=5$)
$w \in R^K$	Learnable gating parameters
α_k	Softmax-normalized ensemble weight
$\hat{y}, y$	Predicted and true retention (at given cycle)
CEEMDAN Augmentation (Sec. III-A)	
$x(k)$	Original retention signal in cycle domain
$\mathcal{E}$	CEEMDAN noise amplitude
n	Number of noise-assisted realizations
IMF_j	j-th intrinsic mode function
α_j, β_j	Amplitude / frequency perturbation factors for IMF_j
CEEMDAN Augmentation (Sec. IV-C)	
Q	Held-out query (final 20%) segment
$\bar{y}_Q$	Mean retention on the query segment

A. CEEMDAN DECOMPOSITION AND AUGMENTATION OF SOURCE DOMAIN DATA

For meta-learning, this study uses NASA's 18650 cell lifetime datasets measured at room temperature as the source domain and several VRFB datasets from prior literature [44], [45], [47]. However, because only four cells are available, the diversity of degradation trajectories is limited for robust model training. To address this constraint while preserving the intrinsic degradation characteristics and long-term trends of the original signals, we adopt a time-series decomposition and augmentation strategy based on the Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEMDAN) [22]. The battery capacity degradation signal is defined in the cycle domain, with the charge-discharge cycle index, rather than time, serving as the independent variable, and it exhibits a nonlinear decay pattern as cycles accumulate. Accordingly, we denote the original degradation signal as $x(c)$ and decompose it into multiple intrinsic mode functions (IMFs) and a residual component using CEEMDAN, enabling separate treatment of oscillatory components and long-term degradation trends for data augmentation.

CEEMDAN provides stable decomposition of nonstationary signals while effectively mitigating mode mixing, making it well suited for augmentation that preserves underlying degradation patterns. The decomposed signal can be expressed as

$$x_k(c) = x\,(c) + \epsilon \cdot \omega_k(c), \tag{3a}$$
$$m\,(c) = (e_{\max}\,(c) + e_{\min}\,(c))/2, \tag{3b}$$
$$h\,(c) = x_k\,(c) - m(c), \tag{3c}$$
$$h_{k+1}\,(c) = h_k\,(c) - m_k(c). \tag{3d}$$

In the above formulation, $x_k(c)$ denotes the signal obtained by adding Gaussian noise $\omega_k(c)$ to the original degradation signal, where ϵ is a scaling parameter that controls the noise amplitude. Injecting noise into the original signal helps alleviate mode mixing during the decomposition process and enables the extraction of more stable IMFs. For each noisy realization $x_k(c)$, $e_{\max}(c)$ and $e_{\min}(c)$ represent the upper and lower envelope signals interpolated from its local maxima and minima, respectively, and their mean is defined as $m(c)$. The candidate component $h(c)$ is then obtained by subtracting this mean envelope from the noisy signal. For $h(c)$ to be regarded as an IMF, it must satisfy the following conditions.

$$|\text{Zero crossings} - \text{Extrema}| \leq 1 \tag{4a}$$
$$m\,(c) \approx 0 \tag{4b}$$

These conditions ensure that the IMF contains only oscillatory components without any residual trend. If the conditions are not satisfied, the sifting process in (3b) and (3c) is repeated iteratively. Once an IMF that meets the criteria is obtained, the final IMF is determined as the ensemble mean over n noise-assisted realizations, as given by the following expression.

$$\text{IMF}_k(c) = \frac{1}{n}\sum_{i=1}^{n} \text{IMF}_k^i(c) \tag{5}$$

Each IMF captures a distinct frequency component: the first modes predominantly represent high-frequency noise, the intermediate modes reflect degradation-related patterns, and the final residual term describes the global trend. As described above, we augment the source-domain data by applying frequency-domain perturbations. Specifically, for each IMF, its frequency-domain representation is modified by adjusting the amplitudes and selected frequency bands as follows.

$$\tilde{\text{IMF}}_k(c) = (1+\delta_{A,k})\text{IMF}_k((1+\delta_{\omega,k})c), \tag{6a}$$
$$x\,(c) = \sum_{k=1}^{3} \tilde{\text{IMF}}_k(c) + r(c). \tag{6b}$$

In the above equation, $\alpha_{A,k}$ and $\alpha_{\omega,k}$ denote the amplitude and frequency perturbation factors for the k-th IMF,

respectively. Each IMF is perturbed according to (6a), and the augmented degradation signal is reconstructed via (6b). The ranges of amplitude and frequency perturbations are specified in Table 3, and the perturbation values are randomly sampled from uniform distributions within these ranges. Different perturbation magnitudes are applied to each IMF. For IMF1, which predominantly represents measurement noise, excessive frequency perturbation is avoided to prevent unrealistic fluctuations in the reconstructed signal. IMF2 and IMF3 are associated with medium-scale degradation mechanisms influenced by operating conditions and temperature; thus, their amplitude perturbations are chosen between those applied to IMF1 and the residual component. For the residual, which captures the global degradation trend, only small amplitude perturbations are introduced without modifying its frequency characteristics. Using this strategy, we generate diverse yet physically consistent degradation trajectories while preserving the overall capacity fade behavior. In this study, the original four NASA cell datasets are augmented to 4,000 synthetic trajectories, which are employed as the source-domain data for meta-learning.

Standard PyEMD (EMD-signal) implementation with the following settings was used: noise amplitude $\varepsilon = 5 \times 10^{-3}$ (relative to the signal standard deviation, applied with noise scale =1.0 Gaussian), number of noise-assisted realizations $n = 100$, range threshold $\theta_{\text{range}} = 0.01$ and total-power threshold $\theta_{\text{power}} = 0.05$ as the sifting stopping criteria (both PyEMD defaults). No explicit cap is placed on the maximum IMF count; in practice, the algorithm converges to $K \in [4, 6]$ IMFs for the NASA and VRFB retention trajectories used in this study, which are then grouped into IMF1 (high-frequency measurement noise), IMF2-IMF4 (degradation-related medium frequencies), and the residue (global trend).

TABLE 3. Perturbation range of amplitude and frequency.

IMFs	Physical meaning	$\delta_{A,k}$	$\delta_{\omega,k}$
IMF_1	Measurement noise	$\mathcal{U}[-0.05,0.05]$	$\mathcal{U}[-0.005,0.005]$
IMF_2	Medium scale degradation	$\mathcal{U}[-0.03,0.03]$	$\mathcal{U}[-0.01,0.01]$
IMF_3	Medium scale degradation	$\mathcal{U}[-0.02,0.02]$	$\mathcal{U}[-0.01,0.01]$
IMF_4	Global degradation	$\mathcal{U}[-0.01,0.01]$	$\mathcal{U}[-0.0,0.0]$

B. META-LEARNING FRAMEWORK FOR CROSS-DOMAIN PROBLEMS

To address the limited availability of training data and the variability of degradation patterns under different operating conditions, while remaining independent of specific electrochemical mechanisms, we adopt a meta-learning framework. In this framework, the datasets are partitioned into a source domain and a target domain to update hierarchical model parameters. Typically, model parameters are first trained on a large amount of source-domain data, after which fine-tuning is performed using a relatively small number of target-domain samples to maximize generalization performance. A key advantage of meta-learning is its ability to maintain strong predictive performance even when the similarity between the source and target domains is low. As summarized in Table 4, the source–target configurations can be categorized into four cases depending on whether the electrochemical system and operating conditions are matched or mismatched. In this study, lithium-ion battery data are used as the source domain and VRFB data are used as the target domain, with differing operating conditions; thus, the problem setting corresponds to the different-component, different-working-condition (DCDW) case.

Among various meta-learning approaches, this study adopts the Reptile algorithm [39], a first-order meta-learning method related to MAML [38] that avoids back-propagating through the inner adaptation steps. In Reptile, the dataset for each task is divided into a support set and a query set. The model parameters are first updated on the support set via inner-loop gradient descent, and the resulting task-adapted parameters are compared with the meta-parameters to form a first-order update; unlike full MAML, no second-order derivative is required, which reduces both memory cost and gradient-estimation variance in the few-shot regime considered in this paper.

For the lithium-ion battery dataset, the support set is partitioned into mini-batches of a predefined size, and the inner-loop optimization updates the model parameters using the average loss over each mini-batch, which contributes to the meta-parameter update.

$$\mathcal{L}_{\text{M}}(\theta^t) = \frac{1}{k}\sum_{i=1}^{K} \ell_S(f_{\theta^t}\left(x_i^s\right), y_i^s)$$
$$\theta^{(t+1)} = \theta^{(t)} - \alpha\nabla_\theta \mathcal{L}_M(\theta^{(t)}) \tag{7a}$$

where,

$$t = 0, 1, \ldots N \text{ and } \theta' = \theta^{(N+1)}. \tag{7b}$$

In this formulation, θ denotes the initial meta-parameters, and f_{θ^t} represents the model parameterized by θ^t. Here, t is the mini-batch index, K is the number of samples in each mini-batch, and ℓ_S is the loss function computed on the support set. Because this study addresses a regression problem, the root mean squared error (RMSE) is employed as the loss function. In each inner loop, the average loss over the K samples in a mini-batch is calculated, and the meta-parameters are updated via gradient descent using the predefined learning rate α. This inner-loop optimization is repeated over all mini-batches in the support set.

In the outer loop of meta-learning, the meta-parameters updated using the support set are evaluated on the query set.

The following expressions define the outer-loop loss and the corresponding meta-parameter update procedure as

$$\Delta_i = \theta_i' - \theta \text{ for each task imeta-batch,} \tag{8a}$$
$$\theta \leftarrow \theta + \beta \cdot \frac{1}{|B|}\sum_{i\in B} \Delta_i. \tag{8b}$$

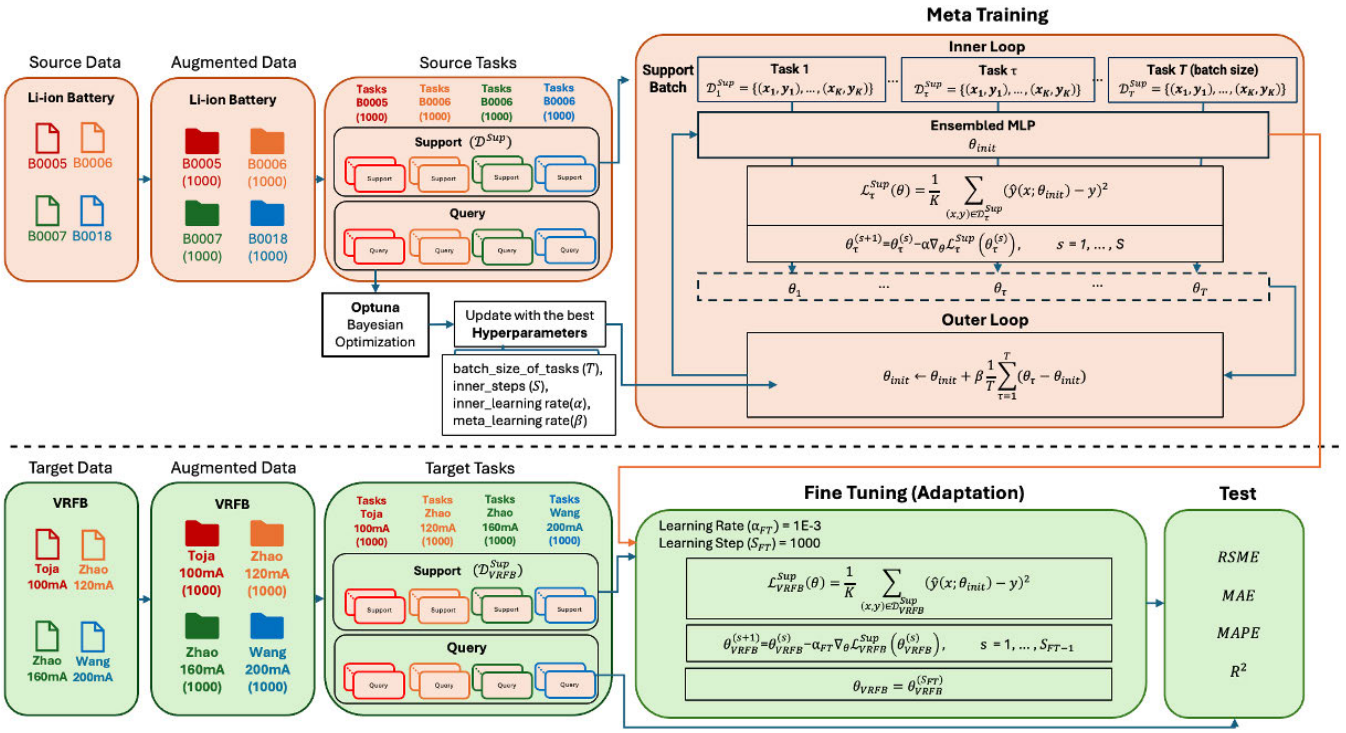


FIGURE 2. Meta-learning framework for VRFB capacity retention rate prediction.

In these equations, θ_i denotes the task-adapted parameters obtained after (T) inner-loop gradient steps on task (i)'s support set (7b), B is the mini-batch of tasks sampled at each meta-epoch, and β is the outer-loop learning rate. The outer update (8b) moves the meta-parameters (θ) in the averaged direction of task-specific adaptation without computing $\nabla_\theta \theta_i'$; this is the Reptile approximation to the full MAML meta-gradient, and in practice, produces comparable few-shot performance at a fraction of the compute. The loss itself is still evaluated on the query set during Optuna hyperparameter search (leave-one-cell-out), but it is not back-propagated into θ. By iteratively performing the inner and outer loops across all source-domain tasks, the meta-learning process converges to an initial set of meta-parameters that generalize effectively to unseen target-domain tasks.

Although the number of target-domain samples is typically limited, initializing with well-generalized meta-parameters from the source domain enables efficient adaptation with only a small amount of data. In the target domain, parameter updates based on the support set are carried out as

$$\theta^{(t+1)} = \theta^{(t)} - \alpha_T \nabla_\theta \mathcal{L}_T(\theta) \quad \text{where,}$$
$$t = 0, 1, \ldots n \text{ and } \theta^* = \theta^{(n+1)}. \tag{9}$$

In this formulation, $\mathcal{L}_T$ denotes the loss incurred during fine-tuning on the target domain, and α_T is the learning rate used for this process. As in the source domain, parameter optimization is performed via gradient descent. The initial model parameters in the target domain are set to the meta-parameters θ obtained from source-domain training, and through fine-tuning on the target support set, they are updated to the final adapted parameters θ^*.

The meta-learning framework operates as described above and enables effective generalization when data from the target domain are scarce by leveraging abundant data from a related source domain. Using this framework, the present study predicts the capacity-degradation trajectory of VFRBs based on their early-cycle capacity data and estimates the point at which the capacity falls below a predefined threshold.

TABLE 4. Definition for meta-learning scenarios.

Scenario	Source Domain	Target Domain
Same Component Same Working Condition (SCSW)	Lithium Ion Battery	Lithium Ion Battery
Same Component Different Working Condition (SCDW)	Lithium Ion Battery	Lithium Ion Battery
Different Component Same Working Condition (DCSW)	Lithium Ion Battery	Vanadium Redox Flow Battery
Different Component Different Working Condition (DCDW)	Lithium Ion Battery	Vanadium Redox Flow Battery

C. MLP ALGORITHM WITH ENSEMBLING

Within the meta-learning framework, data-driven modeling is implemented using a MLP algorithm. The MLP is a fully connected neural network composed of an input layer, one or

TABLE 5. Final hyperparameters selected by Optuna (TPE) with leave-one-cell-out validation on the Li-ion source domain.

Symbol	Description	Value
T	Inner-loop steps	5
α	Inner-loop learning rate	1.4 × 10−3
β	Outer-loop (meta) learning rate	4.87 × 10−2
E	Meta-epochs	623
|B|	Meta-batch size (tasks per step)	3
S	Support-set size (cycles per task)	80
α_T	Target fine-tuning learning rate (9)	1.0 × 10−3
n	Target fine-tuning steps	8000

Algorithm 1 Reptile Meta-Training Used in This Work

Require: Source tasks {τi}Mi=1 with support/query split; inner steps T; inner LR α; meta LR β; meta-batch size |B|; meta-epochs E
1: Initialize meta-parameters θ randomly
2: for epoch e = 1 to E do
3: Sample mini-batch of tasks B ⊆ {τ_i}
4: $\Delta \leftarrow 0$
5: for each task $\tau \in$ B do
6: $\theta\prime \leftarrow \theta$ {reset to meta-parameters}
7: for t = 1 to T do
8: $\theta' \leftarrow \theta' - \alpha \cdot \nabla\theta' \mathcal{L}M(\theta'; \tau^{\text{sup}})$
9: end for
10: $\Delta \leftarrow \Delta + (\theta\prime - \theta)$
11: end for
12: $\theta \leftarrow \theta + \beta \cdot (\Delta/|B|)$ {Reptile outer update}
13: end for
14: Return θ {initialization for target fine-tuning (9)}

more hidden layers, and an output layer. At each layer, affine transformations with learnable weights followed by nonlinear activation functions are applied, enabling the network to approximate complex nonlinear mappings between the input features and the target values.

The MLP can be mathematically formulated as follows.

$$\mathbf{z}_l = \mathbf{W}_l \mathbf{h}_{l-1} + \mathbf{b}_l \tag{10a}$$

$$\mathbf{h}_l = \sigma(\mathbf{z}_l) \tag{10b}$$

$$\hat{y} = \mathbf{w}_{\text{out}} \mathbf{h}_L + b_{\text{out}} \quad \text{where} \quad l = 1, 2 \ldots L \tag{10c}$$

In the above equations, $\mathbf{W}_l$ is the weight matrix of the l-th layer, $\mathbf{b}_l$ is the corresponding bias vector, and $\mathbf{z}_l$ denotes the linearly transformed input to that layer. The layer output is represented by $\mathbf{h}_l$, with $\mathbf{h}_0 = \mathbf{x}$ being the input feature vector. The input feature dimension is six, corresponding to the statistical characteristics of the given window size. The selected features for each window are cycle number, mean, standard deviation, maximum value, minimum value, and slope. In the algorithm, the window slides along the input data and collects the features used as inputs to the MLP. Because the proposed model solves a regression problem for predicting battery-capacity prediction, the parameter set θ is optimized over the

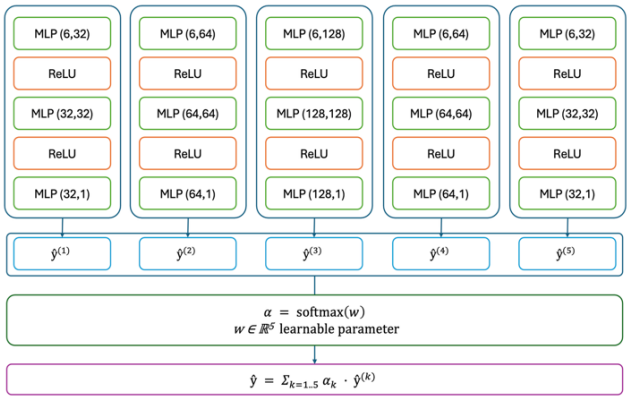


FIGURE 3. Proposed framework for capacity-fade prediction of redox-flow batteries, comprising (i) CEEMDAN-based frequency-domain augmentation of the Li-ion source domain, (ii) Reptile meta-learning, and (iii) a jointly-trained, softmax-gated MLP ensemble of K=5 experts. The softmax block normalises only the learnable gating vector $w \in R^K$); individual MLPs produce regresion scalar outputs directly.

training dataset $\{(\mathbf{x}_i, y_i)\}_{i=1}^N$ by minimizing the mean squared error (MSE), as follows.

$$\mathcal{L}(\theta) = \frac{1}{N} \sum_{i=1}^{N} (y_i - f(\mathbf{x}_i;\theta))^2 . \tag{11}$$

In the above equation, N represents for the length of the total input data, and subscript i denotes for the data index. The loss is minimized by optimizing the MLP model parameter θ.

To enhance the stability and generalization performance of the proposed framework, a weighted ensemble strategy was adopted instead of relying on a single MLP model. As shown in Fig. 3, $K = 5$ MLPs with heterogeneous hidden widths [32, 64, 128, 64, 32] and per-model input-dropout rates [0.1, 0.2, 0.3, 0. 2, 0.1] to promote diversity are trained jointly, and their scalar regression outputs are linearly combined through a softmax-normalized learnable gating vector $w \in R^K$. The gating parameters w are part of the model and are updated simultaneously with the backbone parameters θ^k during both the Reptile inner-loop (7b) and the outer-loop (8b). No post-hoc validation-based reweighting is performed. This approach allows the contribution of each MLP to be adjusted according to its per-target predictive reliability, thereby improving accuracy and robustness compared with uniform averaging.

For a given input vector $\mathbf{x}$, the k-th MLP model is defined as follows.

$$\hat{y}^{(k)} = f^{(k)}(\mathbf{x};\theta^{(k)}), \tag{12a}$$

$$\hat{y} = \sum_{k=1}^{5} \alpha_k \hat{y}^{(k)}, \tag{12b}$$

$$\alpha_k = \exp(w_k)/\Sigma\exp(w_j), \quad w \in \mathbb{R}\hat{}5\text{learnable} \tag{12c}$$

where α_k is the softmax-normalized mixing coefficient for the k-th MLP, obtained from the learnable gating vector w. At initialization, w is set to 0 so that $\alpha^0 = (1/K)1$ (uniform averaging), and during meta-training both w and the backbone parameters θ^k are updated jointly via Reptile. Because

only the gating vector, not the individual predictions $\hat{y}^k$ passes through a softmax, the ensemble remains a weighted combination of regression outputs rather than a classification-style mixture. This learnable gating allows the ensemble to assign greater weight to MLPs with superior per-target predictive reliability, reducing prediction errors relative to uniform averaging and providing a principled uncertainty estimate through the cross-member standard deviation. A schematic illustration of the MLP-based weighted ensemble model is provided in Fig. 3.

D. HYPERPARAMETER TUNING WITH BAYESIAN OPTIMIZATION

Bayesian optimization is an efficient technique for exploring hyperparameters in machine learning and deep learning models. Traditional methods, such as Grid Search and Random Search, evaluate performance by sampling parameters either randomly or at fixed intervals within a specified range. However, because these methods do not use prior exploration results to inform the selection of the next parameter, they are inefficient because of repeated, unnecessary searches, resulting in significant computational resource consumption. In particular, for deep learning models that require substantial training time, Bayesian optimization can significantly accelerate the optimization process by using a computationally inexpensive surrogate model to probabilistically estimate the objective function.

Bayesian optimization trains the surrogate model using objective-function values (RMSE in this study) obtained from randomly sampled initial parameter values. In this study, the Tree-structured Parzen Estimator (TPE) algorithm is employed as the surrogate model. TPE classifies observed data into a superior group and an inferior group based on performance, independently estimating the probability density function for each group. The number of data points classified into the superior group, denoted as n_{good}, is determined based on the total number of observed data points (n) as follows:

$$n_{\text{good}} = \min([0.1 \times n, 25]). \tag{13}$$

As shown in (13), TPE selects the top 10% of data as the superior group by default. However, as data accumulate, if the count exceeds 25, only the top 25 data points are used to model the probability density function. Through this process, the distribution of the superior group, $l(x)$, and the distribution of the inferior group, $g(x)$, are modeled, respectively. Subsequently, the next parameter candidate to be explored is determined via an Acquisition Function. In this study, the Expected Improvement (EI) is used.

$$\text{EI}(x) = \frac{l(x)}{g(x)}, \tag{14}$$

where the EI function calculates a score that maximizes the probability of belonging to the superior-group distribution $l(x)$ while minimizing the probability of belonging to the inferior-group distribution $g(x)$. Specifically, within the Optuna implementation, 24 hyperparameter samples are randomly extracted from the parameter space at each optimization iteration. The candidate with the highest EI value is selected for actual model training, and the resulting performance is fed back to update the surrogate model iteratively.

V. EXPERIMENTAL DESIGN

In this study, NASA-provided lithium-ion battery datasets are used as the source domain. Among the available datasets, four cells are selected, labeled as B0005, B0006, B0007, and B0018, and tested under room-temperature conditions. These cells are commercial 18650-type lithium-ion batteries. The experimental protocol is as follows: during charging, a constant current (CC) of 1.5 A is applied until the cell voltage reaches 4.2 V, followed by a constant-voltage (CV) step at 4.2 V until the current decreases below 20 mA. During discharge, each cell is operated under constant-current conditions down to cutoff voltages of 2.7 V, 2.5 V, 2.2 V, and 2.5 V for B0005, B0006, B0007, and B0018, respectively. The charge–discharge voltage and current profiles are shown in Fig. 4, and the corresponding capacity-fade trajectories are presented in Fig. 5. The four cells exhibit distinct fade patterns: B0005 and B0007 show gradual monotonic decay, whereas B0006 and B0018 display more pronounced mid-life inflection points, which motivates the CEEMDAN-based frequency-domain augmentation in Section IV-A to expand the diversity of degradation trajectories available for meta-training. The capacity-retention rate is normalized to the initial value to use the data for source-domain training.

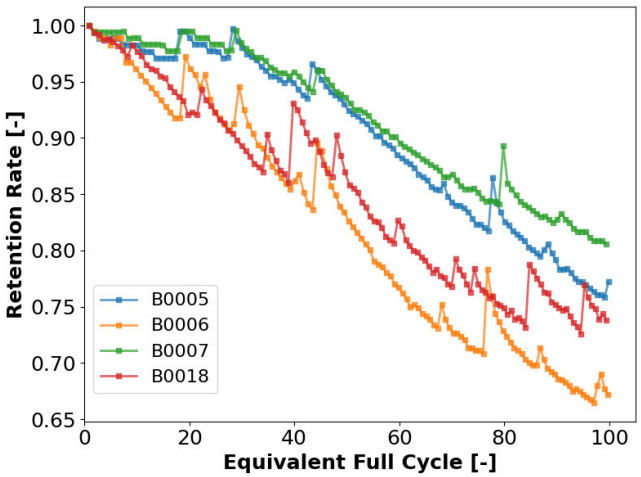


FIGURE 4. Capacity retention rate of NASA battery dataset.

To effectively use the NASA datasets as the source domain, data augmentation is required. Augmentation is performed using the CEEMDAN-based decomposition and frequency-domain perturbation approach described earlier. When applying CEEMDAN to the source-domain data, the number of decomposition modes must be selected carefully. Too few modes can lead to information loss, whereas too many modes may result in overlap among the frequency bands of individual components. Preliminary experiments showed

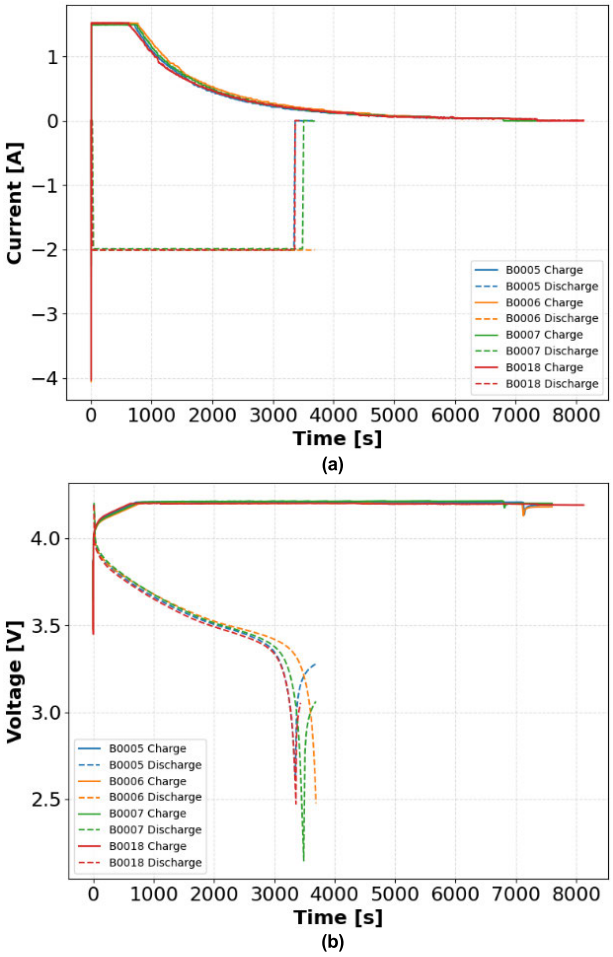


FIGURE 5. Experimental data of NASA dataset: (a) Current trend by time and (b) voltage curves.

that using five modes provides a suitable balance, avoiding significant band overlap while preserving the essential features of the degradation signals. Accordingly, the capacity trajectories of the lithium-ion cells are decomposed into high-frequency noise (IMF1), intermediate-scale degradation components (IMF2, IMF3, IMF4), and a global trend (residual). The resulting decomposed signals are shown in Fig. 6. As illustrated, the early IMFs contain high-frequency components, whereas the frequency content gradually decreases for higher-order IMFs, and the residual captures a smooth global capacity-fade trend.

To augment the original degradation trajectories, a frequency-domain perturbation is applied to all IMFs except the residual, thereby preserving the global trend. For each IMF, the dominant frequencies and amplitudes are modified within the predefined ranges specified in Table 3 by randomly sampling perturbation factors. The corresponding cyclic spectra for each IMF are shown in Fig. 7. As observed, the main frequency components differ across IMFs and exhibit minimal overlap, confirming that each IMF captures distinct information and that the decomposed modes provide a suitable basis for controlled and physically consistent data augmentation.

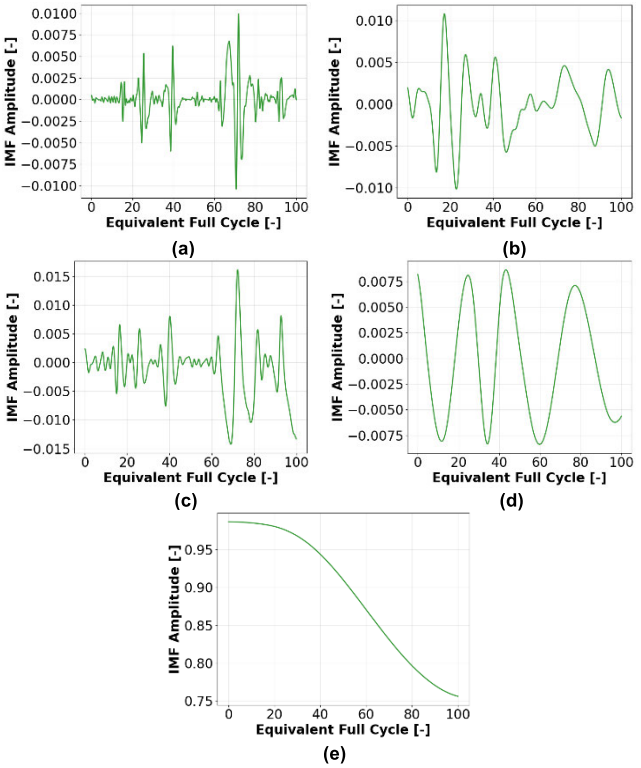


FIGURE 6. IMF signals of B0005 capacity retention: (a) IMF1 (b) IMF2 (c) IMF3 (d) IMF4 and (e) IMF5.

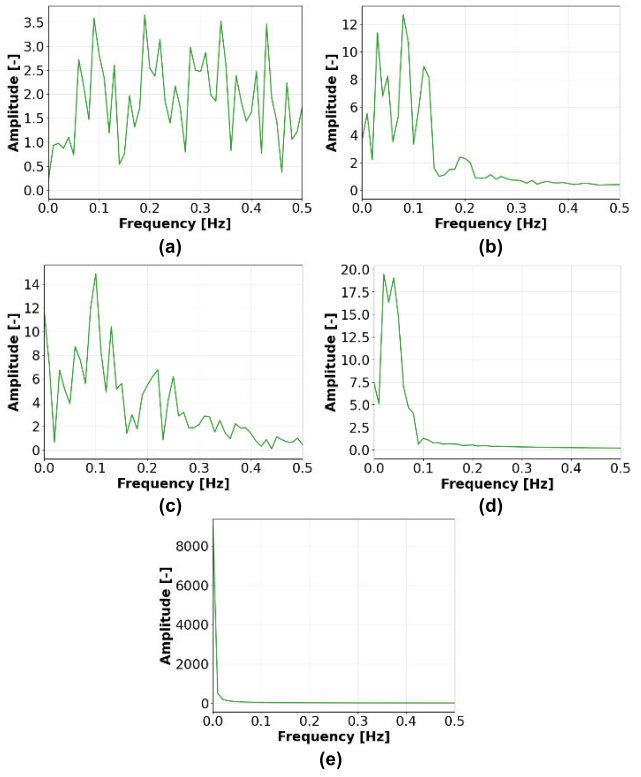


FIGURE 7. Cyclic domain spectra for each IMFs: (a) IMF1 (b) IMF2 (c) IMF3 (d) IMF4 and (e) IMF5.

The final perturbed signal for B0005 is shown in Fig. 8. The figure presents five example perturbed signals that

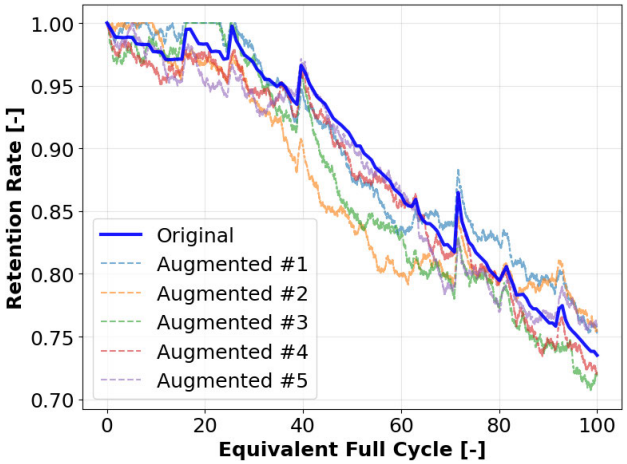


FIGURE 8. Regenerated capacity fade Li-ion battery signals based on CEEMDAN decomposition and frequency domain perturbation.

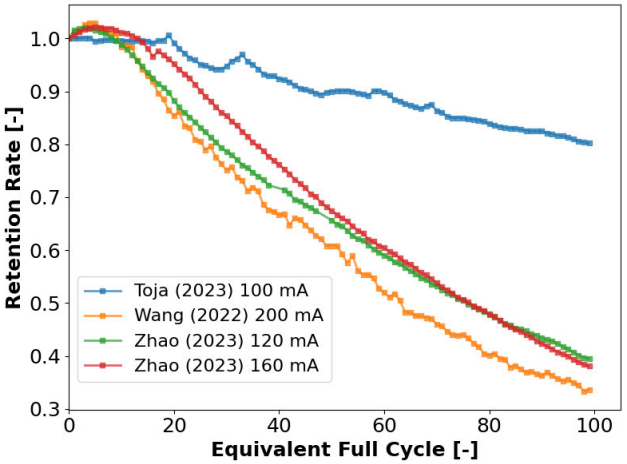


FIGURE 9. Capacity retention rate of VRFB used for target domain.

exhibit different local variations relative to the original capacity-retention data while preserving the global trend. In total, 1,000 perturbations are generated for each capacity profile, augmenting the NASA dataset to 4,000 samples.

The capacity-fade datasets of VRFBs used for the target domain were obtained from previously published studies [44], [45], [47]. Because the datasets differ in initial capacity and test protocols, they correspond to the DCDW setting in meta-learning scenarios. The test conditions and specifications of each dataset are summarized in Table 6. Because the test cycle count varies with operating conditions, the cycle index is normalized to the cycle count. Similarly, the capacity values are normalized as a retention rates and used as target-domain data for model training. Individual VRFB datasets are shown in Fig. 9.

TABLE 6. Specification of VRFB capacity data for the Target domain.

Reference	Current density [mA/cm^2]	Flow rate [mL/min]	Electrolyte volume per tank [mL]	Equivalent full cycle [-]
Wang (2023)	200	40	30	30
Toja (2025)	100	40	50	227
Zhao (2023)	120	25	50	80
Zhao (2023)	160	25	50	81

The VRFB datasets are augmented using the same procedure applied to the NASA dataset. All feature-extraction and normalization statistics are computed exclusively from the first 80% of each trajectory. Because the objective of capacity forecasting is to predict the remaining 20% of the capacity trajectory using the initial 80% of the equivalent-cycle data, only the first 80% of each sequence is augmented and provided to the model as input. After decomposing each signal into components using the CEEMDAN algorithm, we perturb the frequency- and amplitude-related features in the cycle domain according to the rules summarized in Table 3. For each capacity-fade profile, 1,000 augmented samples are generated, resulting in 4,000 target-domain samples in total.

The meta-learning framework is implemented following the workflow shown in Fig. 2. In the source domain, lithium-ion battery data are used to determine the optimal initial model parameters and hyperparameters. Model parameters are optimized via gradient-based learning; however, to improve generalization, we employ an ensemble strategy that aggregates (i.e., averages) the model parameters obtained from multiple mini-batches. Hyperparameters are optimized using Bayesian optimization. These parameters are then transferred to the target domain, where adaptation is performed using the VRFB data. The transferred generalizable parameters are further fine-tuned to better suit VRFB capacity prediction. Finally, the generalized MLP model is evaluated using the query set.

The proposed framework is compared with conventional deep learning models, including MLP, CNN, and LSTM. Each baseline model is trained using the first 80% of the capacity trajectory and was evaluated by forecasting the remaining 20%. A key difference from the baselines is that the proposed framework fine-tunes once and then predicts across all target datasets, whereas conventional models must be trained separately for each dataset. This evaluation setting enables a direct assessment of generalization performance across datasets.

To evaluate the performance of the proposed capacity-prediction model for VRFBs, four metrics are employed: the coefficient of determination (R^2), root mean squared error (RMSE), mean absolute error (MAE), and mean absolute percentage error (MAPE). Values of R^2 and 1-MAPE closer to 1 indicate higher predictive accuracy, whereas lower values of RMSE, MAE indicate that the predicted capacities are in better agreement with the measured values, reflecting higher model reliability. The definitions of these metrics are

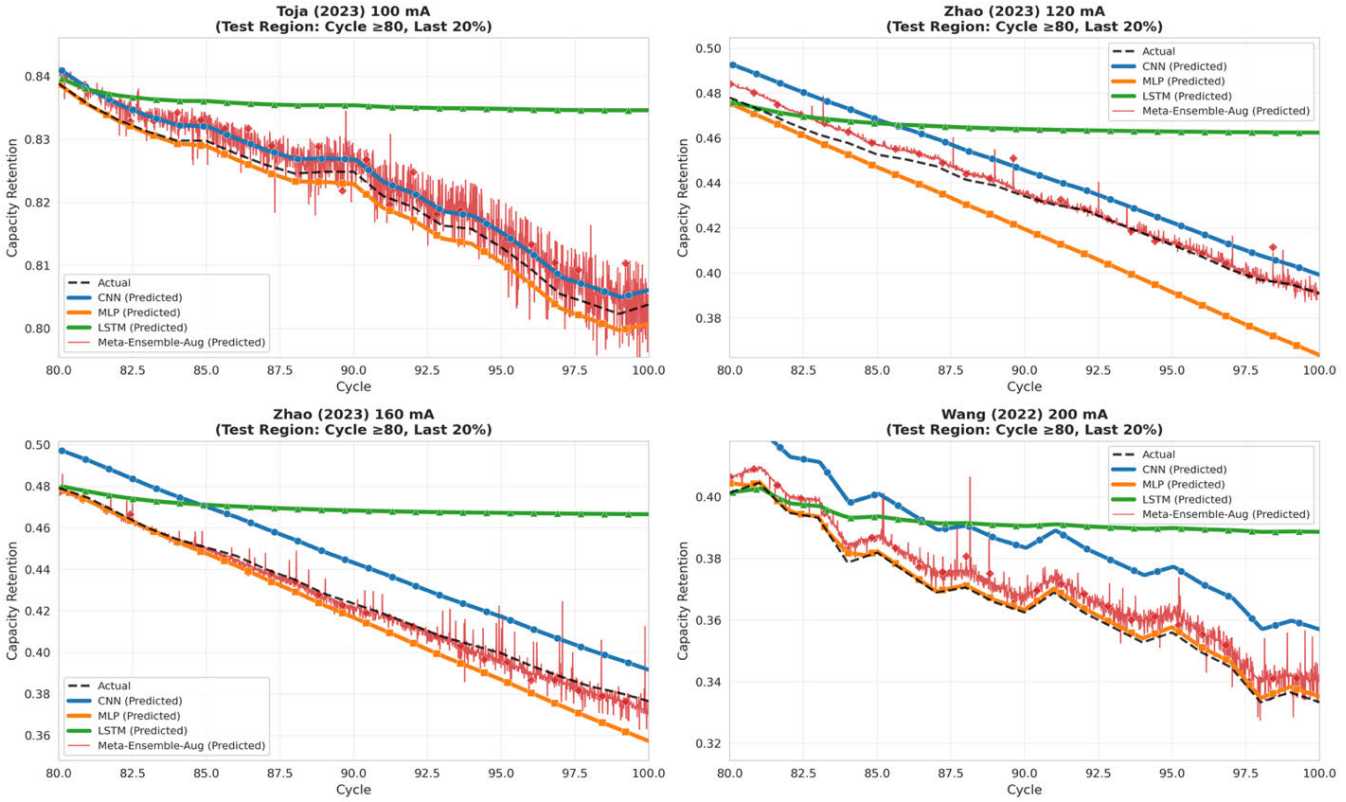


FIGURE 10. Predicted vs. measured capacity-retention trajectories for each dataset (T, W, Z-12, Z-16) comparing the proposed method and representative baselines.

given as

$$R^2 = 1 - \frac{\sum_{i=1}^{n} (y_i - \hat{y}_i)^2}{\sum_{i=1}^{n} (y_i - \bar{y}_i)^2}, \tag{15a}$$

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^{n} (\hat{y}_i - y_i)^2}, \tag{15b}$$

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^{n} |\hat{y}_i - y_i|, \tag{15c}$$

$$1 - \text{MAPE} = 1 - \frac{1}{n} \sum_{i=1}^{n} \left| \frac{\hat{y}_i - y_i}{y_i} \right|, \tag{15d}$$

where $\hat{y}_i$ is the predicted capacity-retention rate from the meta-learning framework and y_i is the true value from the datasets. The variable n denotes the number of cycles for each dataset and subscript i is the number for the cycle index.

VI. EXPERIMENTAL RESULTS

This section evaluates the proposed framework on target-domain VRFB datasets. We compare the generalization performance of the proposed method that of conventional deep learning baselines.

Across all target VRFB datasets, the proposed framework consistently achieves high predictive accuracy and strong generalization performance. Among conventional baselines, the MLP performs strongly on datasets with degradation trajectories that exhibit consistent local curvature, indicating its effectiveness when the underlying geometric pattern remains stable throughout the cycle. However, its performance degrades markedly when the degradation behavior in the early cycles differs from that in the later cycles. This limitation is evident in the Zhao datasets at 120 mA and 160 mA, which show predominantly linear degradation trends with minor local fluctuations but exhibit a distinct change in slope between the early and late stages of the retention trajectory. In contrast, the Tajo 100 mA and Wang 200 mA datasets maintain strong and consistent local curvature from the initial to the final cycles. Although the MLP achieves competitive results on these uniformly curved trajectories, it shows weak learning dynamics and significantly lower prediction accuracy on the Zhao datasets, where the overall trend is approximately linear yet characterized by nonuniform slopes across cycles, despite its generally strong performance on other datasets.

From a model-centric perspective, the CNN and LSTM baselines also fail to robustly capture meaningful degradation trajectories across the evaluated datasets. The LSTM typically fits only the initial level of the sequence and subsequently collapses to an almost constant prediction, reflecting its inability to model long-range global degradation trends. The CNN captures the overall trend only partially

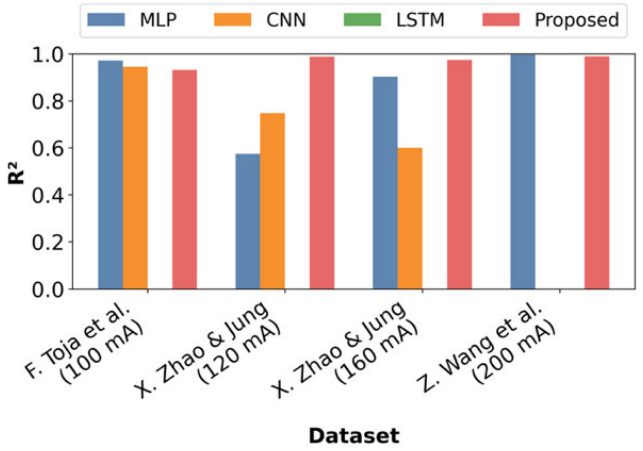


FIGURE 11. Bar chart of R^2 for MLP, CNN, LSTM, and proposed model on each dataset.

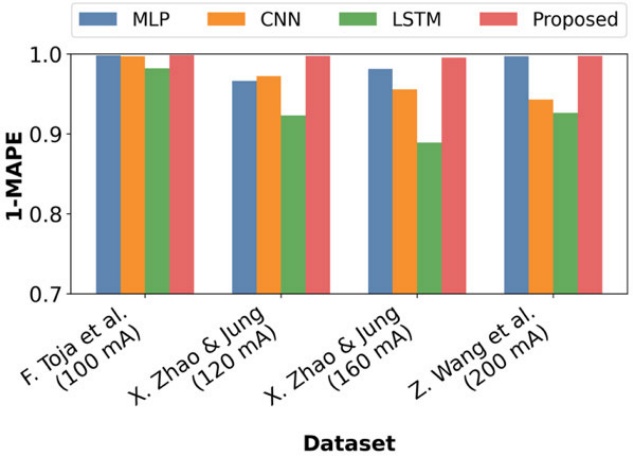


FIGURE 12. Bar chart of 1-MAPE for MLP, CNN, LSTM, and proposed model on each dataset.

and exhibits poor alignment at the beginning of the prediction interval. In the case of the MLP, predictions often rely heavily on extrapolating the early-stage slope, resulting in progressively increasing error as the prediction horizon lengthens. Although the MLP can outperform the proposed framework under limited conditions, particularly when a similar degradation pattern persists throughout the entire lifecycle, as in the Tajo or Wang datasets, it performs poorly on the Zhao trajectories, highlighting its limited robustness to near-linear degradation patterns with differing early- and late-stage slopes.

Despite these dataset-specific behaviors, the proposed method provides the best overall performance when averaged across all target datasets, consistent with the goal of developing a generalized RUL prediction model.

The improved robustness arises from cross-domain meta-learning, which leverages degradation priors learned from the augmented source domain and enables efficient adaptation to each target dataset via fine-tuning. Importantly, the framework trains a single model across multiple datasets rather than fitting separate models for each dataset, promoting generalization under distribution shifts. The results indicate that this strategy does not sacrifice performance; instead, performance is largely preserved and often improved. Specifically, the proposed method achieves the best $1-\text{MAPE}$ metric across all target datasets and attains the highest R^2 on the Zhao datasets, while remaining competitive on the Wang and Toja datasets. Overall, combining cross-domain meta-learning with ensemble learning and augmentation yields robust learning and accurate forecasting under challenging target-domain shifts.

Overall, the proposed method (Meta-Ensemble-Aug) achieves the best mean performance across the target datasets, delivering the lowest average errors (RMSE $=$ 0.003057; MAE $=$ 0.001334) and the strongest fit-oriented metrics (mean $1-\text{MAPE} = 0.9971$; mean $R^2 = 0.9704$). On the Zhao datasets, the conventional single-model baselines degrade substantially, whereas Meta-Ensemble-Aug remains consistently strong in both $1-\text{MAPE}$ and R^2. Under some conditions, such as the Wang and Toja datasets, the MLP attains the best RMSE or R^2, as shown in Tables 7 and 8; however, these isolated gains do not translate into reliable cross-dataset generalization. Taken together, the results indicate that Meta-Ensemble-Aug is the most robust choice for developing a generalized RUL prediction model.

TABLE 7. Model prediction results in terms of RMSE.

	Toja (2023) 100 mA (↓)	Zhao (2023) 120 mA (↓)	Zhao (2023) 160 mA (↓)	Wang (2022) 200 mA (↓)	Average (↓)
MLP	**0.001758**	0.016116	0.009356	**0.001259**	0.007122
LSTM	0.017671	0.038681	0.052456	0.030633	0.034860
CNN	0.002417	0.012329	0.018848	0.020829	0.013606
Ours	0.002774	**0.002573**	**0.004720**	0.002162	**0.003057**

TABLE 8. Model prediction results in terms of MAE.

	Toja (2023) 100 mA (↓)	Zhao (2023) 120 mA (↓)	Zhao (2023) 160 mA (↓)	Wang (2022) 200 mA (↓)	Average (↓)
MLP	0.001527	0.014138	0.007659	0.001113	0.006109
LSTM	0.014905	0.032033	0.045061	0.026162	0.029540
CNN	0.002410	0.012074	0.018808	0.020756	0.013512
Ours	**0.001357**	**0.001162**	**0.001937**	**0.000878**	**0.001334**

TABLE 9. Model prediction results in terms of 1-MAPE.

	Toja (2023) 100 mA (↑)	Zhao (2023) 120 mA (↑)	Zhao (2023) 160 mA (↑)	Wang (2022) 200 mA (↑)	Average (↑)
MLP	0.9981	0.9663	0.9810	0.9969	0.9856
LSTM	0.9817	0.9230	0.8889	0.9260	0.9299
CNN	0.9971	0.9723	0.9556	0.9429	0.9670
Ours	**0.9983**	**0.9972**	**0.9952**	**0.9975**	**0.9971**

TABLE 10. Model prediction results in terms of R^2.

	Toja (2023) 100 mA (↑)	Zhao (2023) 120 mA (↑)	Zhao (2023) 160 mA (↑)	Wang (2022) 200 mA (↑)	Average (↑)
MLP	**0.9713**	0.5736	0.9025	**0.9959**	0.8608
LSTM	-1.9267	-1.4844	-2.1000	-1.4825	-1.7484
CNN	0.9452	0.7476	0.5998	-0.1477	0.5362
Ours	0.9309	**0.9876**	**0.9741**	0.9888	**0.9704**

The proposed method (Meta-Ensemble-Aug) achieves the best mean performance across the target datasets, yielding the lowest average errors and the strongest fit-oriented metrics. On the Zhao datasets, the conventional single-model baselines degrade, whereas Meta-Ensemble-Aug remains consistently strong in both $1-\text{MAPE}$ and R^2, as shown in Figs. 11 and 12. These results indicate that Meta-Ensemble-Aug is the most robust approach for developing a generalized RUL prediction model.

VII. LIMITATIONS AND FUTURE WORK

This study demonstrates the feasibility of cross-chemistry meta-learning for VRFB capacity-retention prediction using limited target-domain data. By transferring degradation priors from augmented Li-ion trajectories to early-cycle VRFB observations, the proposed framework improves robustness under heterogeneous current-density conditions. Nevertheless, the present work should be regarded as an initial demonstration rather than a fully general battery-health prediction framework.

Future work will first focus on expanding dataset diversity. The source domain in this study was constructed from four NASA Li-ion cells and enriched through CEEMDAN-based frequency-domain augmentation, while the target domain consisted of four VRFB capacity-retention trajectories from prior studies. Although this strategy enables effective learning under data scarcity, synthetic augmentation cannot fully replace genuine variability across cells, protocols, and systems. Larger Li-ion, non-Li-ion, and VRFB datasets will therefore be needed to systematically evaluate transferability across chemistries and degradation patterns.

In addition, broader operating conditions should be considered for practical deployment. The current VRFB datasets cover different current densities, but factors such as electrolyte composition, flow rate, temperature, stack configuration, and rebalancing history were not systematically varied. Incorporating such operating-condition metadata, uncertainty-aware prediction, and mechanism-aware features may further improve robustness and support maintenance decisions such as electrolyte rebalancing or makeup-electrolyte scheduling. This work suggests that degradation knowledge learned from data-rich battery chemistries can be successfully transferred to data-scarce electrochemical systems. We believe this cross-chemistry learning paradigm provides a promising direction for developing generalized battery prognostics when target-domain data are limited.

VIII. CONCLUSION

This study proposes a cross-chemistry meta-learning framework for robust capacity-retention prediction in VRFBs. By meta-training on an augmented lithium-ion source domain and fine-tuning using limited early-cycle VRFB data, the proposed approach addresses the key challenge of generalization under severe domain mismatch arising from differences in battery chemistry and operating conditions. The main findings and practical implications can be summarized as follows.

- The meta-learned initialization enables efficient adaptation to VRFB datasets using only early-cycle data while maintaining high predictive accuracy over the late-cycle region.
- Compared with conventional deep learning baselines trained per dataset (MLP, CNN, and LSTM), the proposed framework achieves substantially better cross-dataset generalization, including challenging cases in which single-trajectory training fails.
- Incorporating a weighted ensemble and frequency-domain augmentation further improves robustness and accuracy by increasing the diversity of plausible degradation patterns while preserving the global trend.
- The ablation under the multi-seed protocol (Table 11) confirms strong component synergy: all three components contribute positively in the marginal aggregate ($\Delta\text{Meta} = +0.0037$, $\Delta\text{Ens} = +0.0019$, $\Delta\text{Aug} \approx 0$ on RMSE), and row 8 achieves the best performance on every metric within that table. The most informative controlled pair (row 4 vs row 8, only Meta differs) shows a 47% RMSE reduction attributable to meta-learning in the full stack, complementing the single-run baseline comparisons in Tables 7–10
- The proposed approach moves beyond within-condition prediction by supporting capacity-retention forecasting across heterogeneous current conditions (100, 120, 160, and 200 mA).
- The successful transfer of degradation priors from lithium-ion data to VRFB trajectories suggests that the methodology may extend to other battery chemistries with distinct degradation mechanisms.

APPENDIX

To isolate the contribution of Reptile meta-learning (Meta), CEEMDAN frequency-domain augmentation (Aug), and the learnable softmax-gated five-MLP ensemble (Ens), we evaluated all $2^3 = 8$ on/off combinations on the four VRFB targets under the same paper-recipe protocol with five random seeds. Table 11 reports RMSE, MAE, $1-\text{MAPE}$, and R^2 as the mean $\pm$ standard deviation across 20 (target $\times$ seed) observations per row.

Table 11 reports mean $\pm$ standard deviation over five random seeds and four VRFB targets under a fixed reproduction protocol. The ablation isolates relative component contributions under the multi-seed protocol and is not intended to reproduce the absolute levels in Tables 7–10. For example, the full stack (row 8 in Table 11) yields average

RMSE=0.0116 ± 0.0062 and R^2 =0.4610 ± 0.708, whereas Table 7 reports 'Ours' RMSE=0.0031 and Table 10 reports 'Ours' R^2 =0.9704 from the single representative run. Because R^2 exhibits high seed-to-seed variance under this protocol, we emphasize RMSE, MAE, and 1−MAPE for component attribution.

The full configuration (Meta+Aug+Ens, row 8) achieves the best performance on every metric within Table 11 (RMSE = 0.0116 ± 0.0062, MAE = 0.0087 ± 0.0053, 1−MAPE = 0.9798 ± 0.0162, R^2 = 0.4610 ± 0.7081), improving over the closest partial configuration (Meta+Ens without Aug, row 6) by 10% in RMSE.

Holding augmentation and ensemble fixed and enabling meta-learning (row 4 → row 8) reduces RMSE from 0.0220 to 0.0116 (47% relative decrease) and raises mean R^2 from −1.03 to +0.46. Toggling the ensemble with Meta and Aug fixed (row 7 → row 8) reduces RMSE from 0.0232 to 0.0116 (50% decrease). Toggling Meta in the only-Ens branch (row 2 → row 6) gives a 33% RMSE decrease, confirming Meta's contribution even without augmentation.

All three components contribute positively on the marginal axis (ΔMeta = +0.0037, ΔEns = +0.0019, ΔAug ≈ 0 on RMSE), with the largest gains appearing when components are combined — demonstrating that the proposed components are jointly synergistic.

TABLE 11. Component ablation results.

Meta	Aug	Ens	RMSE (↓)	MAE (↓)	1−MAPE (↑)	R^2 (↑)
–	–	–	0.0222 ± 0.0205	0.0202 ± 0.0196	0.9535 ± 0.0493	−0.8607 ± 2.3178
–	–	✓	0.0191 ± 0.0107	0.0162 ± 0.0108	0.9633 ± 0.0285	−0.1069 ± 0.9695
–	✓	–	0.0133 ± 0.0104	0.0114 ± 0.0089	0.9739 ± 0.0211	0.0641 ± 1.7548
–	✓	✓	0.0220 ± 0.0138	0.0188 ± 0.0138	0.9596 ± 0.0319	−1.0268 ± 3.4701
✓	–	–	0.0142 ± 0.0097	0.0122 ± 0.0084	0.9721 ± 0.0236	−0.1413 ± 1.8698
✓	–	✓	0.0128 ± 0.0068	0.0095 ± 0.0058	0.9769 ± 0.0183	0.4402 ± 0.5788
✓	✓	–	0.0232 ± 0.0186	0.0202 ± 0.0166	0.9525 ± 0.0433	−0.8339 ± 2.0171
✓	✓	✓	0.0116 ± 0.0062	0.0087 ± 0.0053	0.9798 ± 0.0162	0.4610 ± 0.7081

REFERENCES

[1] V. Bolón-Canedo, L. Morán-Fernández, B. Cancela, and A. Alonso-Betanzos, "A review of green artificial intelligence: Towards a more sustainable future," *Neurocomputing*, vol. 599, Sep. 2024, Art. no. 128096, doi: 10.1016/j.neucom.2024.128096.

[2] G. Tamburrini, "The AI carbon footprint and responsibilities of AI scientists," *Philosophies*, vol. 7, no. 1, p. 4, Jan. 2022, doi: 10.3390/philosophies7010004.

[3] L. Lannelongue, J. Grealey, and M. Inouye, "Green algorithms: Quantifying the carbon footprint of computation," *Adv. Sci.*, vol. 8, no. 12, Jun. 2021, Art. no. 2100707, doi: 10.1002/advs.202100707.

[4] K. M. U. Ahmed, M. H. J. Bollen, and M. Alvarez, "A review of data centers energy consumption and reliability modeling," *IEEE Access*, vol. 9, pp. 152536–152563, 2021, doi: 10.1109/ACCESS.2021.3125092.

[5] M. Khalid, "Smart grids and renewable energy systems: Perspectives and grid integration challenges," *Energy Strategy Rev.*, vol. 51, Jan. 2024, Art. no. 101299, doi: 10.1016/j.esr.2024.101299.

[6] J. Zhao, F. Li, and Q. Zhang, "Impacts of renewable energy resources on the weather vulnerability of power systems," *Nature Energy*, vol. 9, no. 11, pp. 1407–1414, Oct. 2024, doi: 10.1038/s41560-024-01652-1.

[7] L. Xu, K. Feng, N. Lin, A. T. D. Perera, H. V. Poor, L. Xie, C. Ji, X. A. Sun, Q. Guo, and M. O'Malley, "Resilience of renewable power systems under climate risks," *Nature Rev. Electr. Eng.*, vol. 1, no. 1, pp. 53–66, Jan. 2024, doi: 10.1038/s44287-023-00003-8.

[8] S. B. Wali, M. A. Hannan, P. J. Ker, S. A. Rahman, K. N. Le, R. A. Begum, S. K. Tiong, and T. M. I. Mahlia, "Grid-connected lithium-ion battery energy storage system towards sustainable energy: A patent landscape analysis and technology updates," *J. Energy Storage*, vol. 77, Jan. 2024, Art. no. 109986, doi: 10.1016/j.est.2023.109986.

[9] M. Rouholamini, C. Wang, H. Nehrir, X. Hu, Z. Hu, H. Aki, B. Zhao, Z. Miao, and K. Strunz, "A review of modeling, management, and applications of grid-connected Li-ion battery storage systems," *IEEE Trans. Smart Grid*, vol. 13, no. 6, pp. 4505–4524, Nov. 2022, doi: 10.1109/TSG.2022.3188598.

[10] D. Choi, N. Shamim, A. Crawford, Q. Huang, C. K. Vartanian, V. V. Viswanathan, M. D. Paiss, M. J. E. Alam, D. M. Reed, and V. L. Sprenkle, "Li-ion battery technology for grid application," *J. Power Sources*, vol. 511, Nov. 2021, Art. no. 230419, doi: 10.1016/j.jpowsour.2021.230419.

[11] X. Cao, J. Du, C. Qu, J. Wang, and R. Tu, "An early diagnosis method for overcharging thermal runaway of energy storage lithium batteries," *J. Energy Storage*, vol. 75, Jan. 2024, Art. no. 109661, doi: 10.1016/j.est.2023.109661.

[12] Y. Jin, Z. Zhao, S. Miao, Q. Wang, L. Sun, and H. Lu, "Explosion hazards study of grid-scale lithium-ion battery energy storage station," *J. Energy Storage*, vol. 42, Oct. 2021, Art. no. 102987, doi: 10.1016/j.est.2021.102987.

[13] J. Guo, Y. Li, K. Pedersen, and D.-I. Stroe, "Lithium-ion battery operation, degradation, and aging mechanism in electric vehicles: An overview," *Energies*, vol. 14, no. 17, p. 5220, Aug. 2021, doi: 10.3390/en14175220.

[14] Z. Lin, D. Li, and Y. Zou, "Energy efficiency of lithium-ion batteries: Influential factors and long-term degradation," *J. Energy Storage*, vol. 74, Dec. 2023, Art. no. 109386, doi: 10.1016/j.est.2023.109386.

[15] A. Leba Akman, M. Z. Arslan, and M. Farsak, "The rise of vanadium redox flow batteries: A game-changer in energy storage," *J. Alloys Compounds*, vol. 1038, Aug. 2025, Art. no. 182869, doi: 10.1016/j.jallcom.2025.182869.

[16] A. Aluko and A. Knight, "A review on vanadium redox flow battery storage systems for large-scale power systems application," *IEEE Access*, vol. 11, pp. 13773–13793, 2023, doi: 10.1109/ACCESS.2023.3243800.

[17] B. Xiong, J. Tang, Y. Li, C. Xie, Z. Wang, X. Zhang, and H. B. Gooi, "Design of a two-stage control strategy of vanadium redox flow battery energy storage systems for grid application," *IEEE Trans. Sustain. Energy*, vol. 13, no. 4, pp. 2079–2091, Oct. 2022, doi: 10.1109/TSTE.2022.3181751.

[18] W.-J. Zou, Y.-B. Kim, and S. Jung, "Capacity fade prediction for vanadium redox flow batteries during long-term operations," *Appl. Energy*, vol. 356, Feb. 2024, Art. no. 122329, doi: 10.1016/j.apenergy.2023.122329.

[19] P. Loktionov, A. Pustovalova, R. Pichugov, D. Konev, and A. Antipov, "Quantifying effect of Faradaic imbalance and crossover on capacity fade of vanadium redox flow battery," *Electrochimica Acta*, vol. 485, May 2024, Art. no. 144047, doi: 10.1016/j.electacta.2024.144047.

[20] K. A. Severson, P. M. Attia, N. Jin, N. Perkins, B. Jiang, Z. Yang, M. H. Chen, M. Aykol, P. K. Herring, D. Fraggedakis, M. Z. Bazant, S. J. Harris, W. C. Chueh, and R. D. Braatz, "Data-driven prediction of battery cycle life before capacity degradation," *Nature Energy*, vol. 4, no. 5, pp. 383–391, Mar. 2019, doi: 10.1038/s41560-019-0356-8.

[21] Y. Zhao and S. Behdad, "State of health estimation of electric vehicle batteries using transformer-based neural network," *J. Energy Resour. Technol.*, vol. 146, no. 10, Oct. 2024, Art. no. 101703, doi: 10.1115/1.4065762.

[22] Y. Fan, Z. Lin, F. Wang, and J. Zhang, "A hybrid approach for lithium-ion battery remaining useful life prediction using signal decomposition and machine learning," *Sci. Rep.*, vol. 15, no. 1, p. 8161, Mar. 2025, doi: 10.1038/s41598-025-92262-8.

[23] X. Wang, K. Dai, M. Hu, and N. Ni, "Lithium-ion battery health state and remaining useful life prediction based on hybrid model MFE-GRU-TCA," *J. Energy Storage*, vol. 95, Aug. 2024, Art. no. 112442, doi: 10.1016/j.est.2024.112442.

[24] X. Huang, Y. Mao, L. Li, and Y. Li, "A hybrid-driven method for predicting the remaining useful life of lithium-ion batteries," *Meas. Sci. Technol.*, vol. 36, no. 5, May 2025, Art. no. 056122, doi: 10.1088/1361-6501/add042.

[25] Y. Che, Y. Zheng, X. Sui, and R. Teodorescu, "Boosting battery state of health estimation based on self-supervised learning," *J. Energy Chem.*, vol. 84, pp. 335–346, Sep. 2023, doi: 10.1016/j.jechem.2023.05.034.

[26] S. Navidi, A. Thelen, T. Li, and C. Hu, "Physics-informed machine learning for battery degradation diagnostics: A comparison of state-of-the-art methods," *Energy Storage Mater.*, vol. 68, Apr. 2024, Art. no. 103343, doi: 10.1016/j.ensm.2024.103343.

[27] F. Wang, Z. Zhai, Z. Zhao, Y. Di, and X. Chen, "Physics-informed neural network for lithium-ion battery degradation stable modeling and prognosis," *Nature Commun.*, vol. 15, no. 1, p. 4332, May 2024, doi: 10.1038/s41467-024-48779-z.

[28] T. Li, Z. Zhou, A. Thelen, D. A. Howey, and C. Hu, "Predicting battery lifetime under varying usage conditions from early aging data," *Cell Rep. Phys. Sci.*, vol. 5, no. 4, Apr. 2024, Art. no. 101891, doi: 10.1016/j.xcrp.2024.101891.

[29] M. Sudarshan, J. V. R. Vajja, and V. Tomar, "DegradAI: A scalable framework for early battery health diagnosis from limited data," *npj Clean Energy*, vol. 1, no. 1, p. 8, Aug. 2025, doi: 10.1038/s44406-025-00008-2.

[30] X. Zhao, J. Nam, H.-Y. Jung, and S. Jung, "Real-time state of charge and capacity estimations of vanadium redox flow battery based on unscented Kalman filter with a forgetting factor," *J. Energy Storage*, vol. 74, Dec. 2023, Art. no. 109146, doi: 10.1016/j.est.2023.109146.

[31] X. Zhao, K. Kim, and S. Jung, "State-of-charge estimation using data fusion for vanadium redox flow battery," *J. Energy Storage*, vol. 52, Aug. 2022, Art. no. 104852, doi: 10.1016/j.est.2022.104852.

[32] Y. Zhang, R. Xiong, H. He, and M. G. Pecht, "Long short-term memory recurrent neural network for remaining useful life prediction of lithium-ion batteries," *IEEE Trans. Veh. Technol.*, vol. 67, no. 7, pp. 5695–5705, Jul. 2018, doi: 10.1109/TVT.2018.2805189.

[33] P. Li, Z. Zhang, Q. Xiong, B. Ding, J. Hou, D. Luo, Y. Rong, and S. Li, "State-of-health estimation and remaining useful life prediction for the lithium-ion battery based on a variant long short term memory neural network," *J. Power Sources*, vol. 459, May 2020, Art. no. 228069, doi: 10.1016/j.jpowsour.2020.228069.

[34] Y. Liu, G. Zhao, X. Peng, and C. Hu, "Lithium-ion battery remaining useful life prediction with long short-term memory recurrent neural network," in *Proc. Annu. Conf. PHM Soc.*, Oct. 2017, vol. 9, no. 1, pp. 1–7, doi: 10.36001/phmconf.2017.v9i1.2447.

[35] Y.-H. Chang, Y.-C. Hsieh, Y.-H. Chai, and H.-W. Lin, "Remaining-useful-life prediction for Li-ion batteries," *Energies*, vol. 16, no. 7, p. 3096, Mar. 2023, doi: 10.3390/en16073096.

[36] L. Zhao, S. Song, P. Wang, C. Wang, J. Wang, and M. Guo, "A MLP-mixer and mixture of expert model for remaining useful life prediction of lithium-ion batteries," *Frontiers Comput. Sci.*, vol. 18, no. 5, Oct. 2024, Art. no. 185329, doi: 10.1007/s11704-023-3277-4.

[37] S. Suh, D. A. Mittal, H. Bello, B. Zhou, M. S. Jha, and P. Lukowicz, "Remaining useful life prediction of lithium-ion batteries using spatio-temporal multimodal attention networks," *Heliyon*, vol. 10, no. 16, Aug. 2024, Art. no. e36236, doi: 10.1016/j.heliyon.2024.e36236.

[38] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, pp. 1126–1135.

[39] A. Nichol, J. Achiam, and J. Schulman, "On first-order meta-learning algorithms," 2018, *arXiv:1803.02999*.

[40] S. Zhang, Z. Liu, and H. Su, "State of health estimation for lithium-ion batteries on few-shot learning," *Energy*, vol. 268, Apr. 2023, Art. no. 126726, doi: 10.1016/j.energy.2023.126726.

[41] X. Zhao, Z. Wang, T. Han, W. Yang, F. Gu, and A. David Ball, "A meta-learning method for few-shot multidomain state-of-health estimation of lithium-ion batteries," *IEEE Trans. Transport. Electrific.*, vol. 11, no. 1, pp. 4830–4840, Feb. 2025, doi: 10.1109/TTE.2024.3470551.

[42] T. Ma, Z. Huang, X. Xie, and B. Li, "Evaluation of the effect of hydrogen evolution reaction on the performance of all-vanadium redox flow batteries," *Electrochimica Acta*, vol. 504, Nov. 2024, Art. no. 144895, doi: 10.1016/j.electacta.2024.144895.

[43] Z. Huang, Y. Liu, X. Xie, C. Huang, Q. Huang, Z. Guo, and Y. Liu, "Experimental validation of side reaction on capacity fade of vanadium redox flow battery," *J. Electrochemical Soc.*, vol. 171, no. 1, Jan. 2024, Art. no. 010521, doi: 10.1149/1945-7111/ad1ec8.

[44] F. Toja, L. Perlini, D. Facchi, A. Casalegno, and M. Zago, "Dramatic mitigation of capacity decay and volume variation in vanadium redox flow batteries through modified preparation of electrolytes," *Appl. Energy*, vol. 354, Jan. 2024, Art. no. 122262, doi: 10.1016/j.apenergy.2023.122262.

[45] Z. Wang, Z. Guo, J. Ren, Y. Li, B. Liu, X. Fan, and T. Zhao, "An electrolyte with elevated average valence for suppressing the capacity decay of vanadium redox flow batteries," *ACS Central Sci.*, vol. 9, no. 1, pp. 56–63, Jan. 2023, doi: 10.1021/acscentsci.2c01112.

[46] A. Morozov, M. Pugach, A. Polyakov, P. Osinenko, A. Bolychev, V. Terzija, and S. Parsegov, "Optimal flow factor determination in vanadium redox flow battery control," *IEEE Access*, vol. 12, pp. 19277–19284, 2024, doi: 10.1109/ACCESS.2024.3361830.

[47] X. Zhao and S. Jung, "Long-term capacity fade forecasting of vanadium redox flow battery with Gaussian process regression combined with informer model," *J. Power Sources*, vol. 585, Nov. 2023, Art. no. 233670, doi: 10.1016/j.jpowsour.2023.233670.

[48] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, and C. Liu, "A survey on deep transfer learning," in *Proc. Int. Conf. Artif. Neural Netw.*, 2018, pp. 270–279, doi: 10.1007/978-3-030-01424-7_27.

[49] M. Long, H. Zhu, J. Wang, and M. I. Jordan, "Deep transfer learning with joint adaptation networks," in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, pp. 2208–2217.

[50] B. Sun and K. Saenko, "Deep CORAL: Correlation alignment for deep domain adaptation," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 443–450, doi: 10.1007/978-3-319-49409-8_35.

[51] A. Giuliano, Y. Wu, J. Yawney, and S. Andrew Gadsden, "Transformer-based transfer learning for battery state-of-health estimation," *Energies*, vol. 18, no. 20, p. 5439, Oct. 2025, doi: 10.3390/en18205439.

[52] X. Li, M. Zhao, S. Zhong, J. Li, Z. Cui, S. Fu, and Z. Yan, "Deep transfer learning enabled Online state-of-health estimation of lithium-ion batteries under small samples across different cathode materials, ambient temperature and charge-discharge protocols," *J. Power Sources*, vol. 650, Sep. 2025, Art. no. 237503, doi: 10.1016/j.jpowsour.2025.237503.



YOUNG-DAE SHIM received the Ph.D. degree from the Department of Smart Fabrication Technology, Sungkyunkwan University, South Korea, in 2024. After completing his Ph.D., he is currently a Postdoctoral Researcher with Georgia Tech. His research interests include smart manufacturing, prognostics and health management, and electromagnetic sensors.



JUNGHWAN YIM received the B.S. and M.S. degrees in computational physics from the University at Buffalo, State University of New York. He is currently pursuing the Ph.D. degree with the Georgia Institute of Technology. His research focuses on simulation-driven digital twin systems for modeling and understanding complex physical environments, with an emphasis on physical AI and agentic AI for self-adaptive and intelligent decision-making systems. His interests include high-fidelity simulation, reinforcement learning, autonomous agents, and the integration of physics-based models with data-driven artificial intelligence.



NIAN LIU is an Associate Professor with the School of Chemical and Biomolecular Engineering, Georgia Institute of Technology (Georgia Tech), where he also holds the title of Robert G. Miller Faculty Fellow. His research focuses on high-energy, low-cost, and ultra-safe battery technologies, application-driven nanoscale materials design, and fundamental understanding of electrochemical processes through advanced microscopy and experimental methods. In his group, he integrates nanomaterials, electrochemistry, and light microscopy to address critical challenges in energy storage science and technology.



CHEONG HOON KWON received the Ph.D. degree in chemical and biological engineering from Korea University, in 2008. She is currently an Associate Professor with the Department of Energy Resources and Chemical Engineering, Kangwon National University, Samcheok, Republic of Korea. Her work integrates functional nanomaterials, conductive hydrogels, and interfacial engineering to develop high-performance and sustainable electrochemical systems. She has authored and co-authored numerous peer-reviewed publications in international journals. Her research focuses on electrochemical energy conversion and storage, as well as bioelectronics interfaces for enzyme-based biofuel cells and electrochemical biosensors.



HAE-JIN CHOI received the M.S. and Ph.D. degrees in mechanical engineering from Georgia Institute of Technology (Georgia Tech), Atlanta, GA, USA, in 2001 and 2005, respectively. He is currently a Professor with the School of Mechanical Engineering, Chung-Ang University, Seoul, Republic of Korea, where he has been a Faculty Member since 2010. Before joining Chung-Ang University, he was an Assistant Professor with Nanyang Technological University, Singapore. His primary research areas include additive manufacturing, smart manufacturing, robust design, metamodeling, and industrial artificial intelligence (AI) in engineering systems.



SEUNG-KYUM CHOI is currently an Associate Professor with Georgia Institute of Technology. His research interests include structural reliability, probabilistic mechanics, statistical approaches to design of structural systems, multidisciplinary design optimization, and deep learning methods for complex engineered systems. His current research focus is on the development of robust simulation and design-support tools to assist the digital transformation of complex engineered systems involving material structures design and multifunctional products. He is currently an Associate Editor of *ASME Journal of Computing and Information Science in Engineering* and *Journal of Computational Design and Engineering*.