Journal of Energy Storage 148 (2026) 119996

Contents lists available at ScienceDirect

Journal of Energy Storage

journal homepage: www.elsevier.com/locate/est

Research papers

An end-to-end dual-domain attention fusion network with evidential uncertainty for battery state of charge estimation

Zhou Wu [a,c,e], Baile Xu [a,b], Fuguang Wen [e], Jian Zhao [a,d], Furao Shen [a,b,*]

[a] National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing, 210023, PR China
[b] School of Artificial Intelligence, Nanjing University, Nanjing, 210023, PR China
[c] School of Computer Science, Nanjing University, Nanjing, 210023, PR China
[d] School of Electronic Science and Engineering, Nanjing University, Nanjing, 210023, PR China
[e] Nanjing Guodian Nanzi Power Grid Automation Co., Ltd, Nanjing, 210032, PR China

ARTICLE INFO

Keywords:
State-of-Charge estimation
Lithium-ion battery
Uncertainty quantification
Deep learning
Spectral analysis
Transferability
Battery management system
USFFNet

ABSTRACT

Accurate State-of-Charge (SoC) estimation of lithium-ion batteries (LIBs) is pivotal for electric vehicle safety and efficiency. However, current methods often struggle under variable operational conditions due to limitations in capturing complex battery dynamics. In this work, we propose a lightweight yet powerful dual-domain fusion architecture — Uncertainty-aware Spatial-Frequency Fusion Network (USFFNet) — which jointly leverages complementary temporal and spectral electrochemical information, significantly enhancing the robustness and precision of battery SoC estimation. USFFNet integrates learnable spectral decomposition with temporal convolutional encoders, adaptively focusing on electrochemically informative frequency components while maintaining real-time responsiveness. Importantly, a Bayesian evidential deep learning framework is incorporated, enabling principled quantification of **estimation uncertainty** via a Normal-Inverse-Gamma representation. Comprehensive validation on the LG-18650HG2 dataset, spanning harsh thermal conditions (−20 °C to 40 °C), demonstrates that USFFNet achieves substantial improvements over state-of-the-art methods, reducing Mean Absolute Error (MAE) by up to 68% at low temperatures and by a factor of 16.7 at high temperatures. Cross-domain experiments show sustained performance, maintaining over 85% accuracy when transferred across battery chemistries and operational scenarios. Extensive real-world deployment validation on a 1 MWh grid-scale storage system confirms its practical efficacy (MAE of 0.069–0.073) and exceptional **estimation reliability**, correctly identifying 94% of critical **estimation deviations** via uncertainty estimates. This study establishes USFFNet as a reliable framework for intelligent, uncertainty-aware battery management, providing both theoretical insights and practical tools essential for sustainable and efficient energy storage systems.

1. Introduction

Electrification has surged dramatically, with global lithium-ion battery (LIB) production capacity surpassing 1 TW h in 2024, projected to exceed 5 TW h by 2030, primarily driven by electric vehicles (EVs) and battery energy storage systems (BESS) [1,2]. LIBs offer attractive benefits, including high specific energy, power density, and cost-effectiveness [3]. However, these advantages come with operational fragilities such as thermal runaway [4], accelerated aging, and heterogeneous degradation, posing serious safety and economic risks [5]. Central to mitigating these risks is accurate state-of-charge (SoC) estimation, crucial for optimizing charging rates, predicting remaining capacity, and assessing battery health [6]. Small estimation errors (e.g., 2 % absolute SoC error) can lead to substantial operational inefficiencies or severe degradation events [7]. Historically, SoC estimation techniques fall into two main categories: physics-based and data-driven methods [8].

Physics-based models, such as Kalman Filters [9], Electrochemical Models [10], and Equivalent Circuit Models [11], offer interpretability but demand extensive parameter calibration, complex instrumentation, and suffer accuracy degradation over time due to model invalidation and sensor drift [12]. Conversely, data-driven methods, including deep neural networks (DNNs), excel in capturing nonlinear relationships from historical data, showcasing superior scalability and robustness under varying operational conditions [13].

* Correspondence to: Nanjing University Xianlin Campus, No. 163, Xianlin Avenue, Qixia District, Nanjing, Jiangsu Province, PR China.
E-mail addresses: zhouwu@smail.nju.edu.cn (Z. Wu), xubaile@nju.edu.cn (B. Xu), fuguangwen@sac-china.com (F. Wen), jianzhao@nju.edu.cn (J. Zhao), frshen@nju.edu.cn (F. Shen).

https://doi.org/10.1016/j.est.2025.119996
Received 21 August 2025; Received in revised form 16 November 2025; Accepted 19 December 2025
Available online 9 January 2026

List of Abbreviations

Abbreviation	Definition
BMS	Battery Management System
SoC	State of Charge
LIB	Lithium-Ion Battery
EDL	Evidential Deep Learning
NIG	Normal-Inverse-Gamma
FFT	Fast Fourier Transform
IFFT	Inverse Fast Fourier Transform
CNN	Convolutional Neural Network
LSTM	Long Short-Term Memory
GRU	Gated Recurrent Unit
TCN	Temporal Convolutional Network
MAE	Mean Absolute Error
RMSE	Root Mean Squared Error
CR	Calibration Reliability
CV	Coefficient of Variation

Yet, data-driven methods have become more and more central to battery state estimation, largely supplanting traditional physics-based models due to their superior ability to capture complex nonlinear dynamics without requiring extensive parameterization. A significant body of research focuses on time-domain approaches, utilizing recurrent neural network (RNN) architectures such as Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRU) to process sequences of voltage, current, and temperature measurements [14,15]. These models, often enhanced with filtering techniques like the Unscented Kalman Filter (UKF), have demonstrated high accuracy in controlled settings [16]. However, their reliance on direct time-series data can be a critical limitation; such methods often struggle to extract the deeper, more discriminative features that are crucial for robust performance under the variable and noisy conditions seen in real-world applications [15]. This dependency on raw temporal data can lead to performance degradation, as it may not fully capture the underlying electrochemical processes that manifest across different domains.

To overcome the limitations of purely time-domain analysis, recent research has shifted towards fusing temporal data with frequency-domain features, which offer a more direct representation of the battery's internal electrochemical state. For instance, methods incorporating Electrochemical Impedance Spectroscopy (EIS) and Distribution of Relaxation Times (DRT) have been used to create features that explicitly map to physical aging mechanisms like conductivity loss (CL) and loss of lithium inventory (LLI). Building on this, Zhao and Liu proposed a novel feature fusion method for State of Health (SOH) estimation that combines time-domain signals with multi-scale frequency-domain features extracted using Mel-frequency cepstral coefficients (MFCCs) [17]. Their work, which uses a transformer-based deep network, demonstrates that fusing these domains yields low-coupling, complementary features that significantly improve estimation accuracy. The success of such dual-domain fusion in the context of SOH strongly supports its potential for the more dynamic and safety-critical task of SoC estimation, motivating the development of architectures that can similarly leverage both temporal and spectral battery dynamics for enhanced robustness and precision.

Building upon these advances in SOH estimation, recent research on SoC prediction has increasingly adopted dual-domain learning strategies for modeling battery dynamics. Time–frequency convolutional frameworks based on spectrogram-derived 2D representations [18] demonstrate that incorporating spectral information improves estimation accuracy over purely temporal models, but they rely on handcrafted STFT/CWT parameters and deep 2D CNN backbones, resulting in high computational cost and limited adaptability to spectral shifts. Hybrid architectures that couple temporal Mamba blocks with signal decomposition modules [19] further enhance generalization by jointly capturing temporal and spectral patterns, yet the two branches are still loosely coupled and are usually fused via simple concatenation or late fusion. Lightweight designs integrating FFT operations with arrangement-transformation convolutional backbones [20] achieve more efficient temporal–spectral interaction and accurate SoC estimation across multiple chemistries and real-vehicle datasets, but typically operate on fixed FFT magnitude features with shallow spectral modeling and no uncertainty quantification. Explicit time–frequency fusion with adaptive spectral weighting [21] improves robustness to sensor noise and temperature variations, while dual-branch temporal–frequency networks [22] still incur considerable memory/latency overhead and remain sensitive to distribution shifts because their frequency-domain representations are obtained from non-learnable transforms.

Collectively, these studies confirm that explicit time–frequency feature fusion improves the robustness and transferability of SoC estimation under non-stationary conditions, but also reveal persistent limitations: dependence on rigid, non-trainable spectral preprocessing, sub-optimal or decoupled fusion of temporal and spectral cues, heavy backbones that hinder deployment on resource-constrained BMS hardware, and the absence of principled uncertainty modeling. Motivated by these gaps, the proposed USFFNet employs learnable spectral convolution on FFT-based magnitude–phase representations, attention-guided spatial–frequency fusion that enables deep interaction between temporal and spectral features, and an evidential uncertainty-aware prediction head that outputs calibrated SoC estimates together with confidence information. These designs allow USFFNet to provide reliable and interpretable SoC estimation under dynamic operating conditions while maintaining computational efficiency suitable for practical BMS implementations.

While the principle of fusing time- and frequency-domain data is promising, existing dual-domain architectures often exhibit significant limitations that hinder their practical deployment and performance. From an engineering perspective, many proposed fusion models carry substantial computational overhead. Architectures relying on deep and complex backbones — such as large Transformers [23], intricate Mamba-CNN hybrids [24], or sequential CNN-BiLSTM networks [25] — are often ill-suited for the memory and latency constraints of real-time embedded Battery Management Systems (BMS). This gap between algorithmic complexity and hardware reality renders such models impractical for on-device deployment.

Architecturally, a primary deficiency lies in the suboptimal and rigid nature of feature fusion. Many approaches rely on spectral features obtained via static, non-trainable preprocessing, which prevents the model from adaptively learning to prioritize the most salient patterns. For instance, methods may use a fixed Continuous Wavelet Transform (CWT) [25], Mel-frequency cepstral coefficient (MFCC) filter banks [17], or manually engineered features from Electrochemical Impedance Spectroscopy (EIS) and Incremental Capacity Analysis (ICA) [23]. This rigidity means the model cannot refine the frequency decomposition to focus on degradation-specific harmonics that may emerge over a battery's lifetime. Furthermore, the fusion strategy itself is often simplistic; lightweight models may employ a "late fusion" of these engineered features [17], while deeper models extract features in independent streams and combine them via naïve concatenation. This fails to foster a synergistic learning process and can lead to information conflict, where features are not effectively reconciled. These challenges — computational inefficiency, rigid spectral analysis, and simplistic fusion — underscore the need for a unified framework that facilitates adaptive, co-dependent learning across both spatial and spectral domains.

To address these challenges, we propose the *Uncertainty-aware Spatial-Frequency Fusion Network (USFFNet)*—a lightweight yet expressive framework that jointly learns temporal and spectral battery dynamics through a pair of dedicated reconstruction branches. Informed

by recent advances in branch-independent modular networks [26, 27], USFFNet integrates these representations via a learnable Score Block [28], enabling the model to selectively emphasize informative harmonics and localized temporal variations. Crucially, we embed Evidential Deep Learning (EDL) [29] to estimate calibrated uncertainty with minimal overhead, circumventing the computational burden of traditional Bayesian approaches [30] while ensuring interpretability via Normal–Inverse-Gamma distributions [31]. We further benchmark USFFNet on a scenario-driven dataset incorporating thermal drift, voltage anomalies, and sensor faults [32], demonstrating superior generalization in safety-critical operating regimes. In doing so, our approach not only reconciles the performance-efficiency trade-off but also bridges the gap between controlled laboratory validation and the variability of real-world deployment.

This paper proceeds as follows: Section 3 describes the methodology; Section 5 presents experimental results; Section 6 concludes with insights and future outlook.

2. Preliminaries

2.1. Problem formulation

Battery SoC prediction is fundamentally a supervised multivariate time-series regression task. Formally, let the windowed sensor data sequence be represented as $\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_T] \in \mathbb{R}^{N\times L\times T}$, where:

- N denotes the number of sensor channels (e.g., voltage, current, temperature),
- L is the look-back window length capturing local temporal patterns,
- T represents the prediction horizon.

At each time step i, a local data patch $\mathbf{X}_i \in \mathbb{R}^{N\times L}$ is constructed:

$$\mathbf{X}_i = \left[(x_1^{i-L+1}, \dots, x_N^{i-L+1})^\top, \dots, (x_1^i, \dots, x_N^i)^\top\right], \tag{1}$$

where x_n^j is the jth measurement of sensor channel n.

From a dynamical systems perspective, SoC acts as a cumulative state variable responding to complex power excitations. Its evolution is governed by:

$$\mathrm{SoC}(t) = \mathrm{SoC}_0 - \frac{1}{C}\int_0^t I(\tau)\,\mathrm{d}\tau, \tag{2}$$

where $I(t)$ denotes instantaneous current influenced by dynamic load conditions, temperature variations, and self-discharge effects, and C is the battery's rated capacity. Consequently, SoC exhibits inherent inertia, behaving as a low-pass filtered response to $I(t)$.

The objective is to learn a predictive mapping f_θ:

$$\hat{y}_i = f_\theta(\mathbf{X}_i), \quad \text{for } i = 1, \dots, T, \tag{3}$$

accommodating epistemic uncertainty (model-related) and aleatoric uncertainty (measurement noise), both critical for safety-critical applications.

2.2. Motivation for frequency-domain integration

Battery SoC dynamics inherently exhibit multi-scale frequency characteristics stemming from the underlying electrochemical processes, diffusion mechanisms, and thermal dynamics, each associated with distinct frequency bands. Operationally, batteries encounter a broad frequency spectrum:

- Low-frequency (< 0.1 Hz): Capturing long-term degradation, seasonal temperature variations, and daily operational cycles,
- Mid-frequency (0.1–10 Hz): Representing transient load fluctuations, electrochemical kinetics, and battery management interventions,
- High-frequency (> 10 Hz): Associated primarily with measurement noise, electromagnetic interference, and switching transients.

These frequency components propagate through the battery's equivalent circuit model (ECM), which can be represented by frequency-dependent impedance:

$$Z(\omega) = R_s + \sum_{i=1}^{n} \frac{R_i}{1 + \mathrm{j}\omega R_i C_i}, \tag{4}$$

where each RC element reflects specific electrochemical processes with characteristic time constants $\tau_i = R_i C_i$. Frequency-domain analysis thus enables explicit isolation and quantification of these processes.

Spectral decomposition of SoC signals provides three distinct advantages over purely temporal approaches:

- Mechanistic interpretability: Spectral features directly correspond to specific electrochemical and thermal processes, facilitating clearer physical insights.
- Long-range dependency preservation: Low-frequency spectral features efficiently encode aging trajectories and long-term capacity variations beyond local temporal windows, achievable via FFT operations with computational complexity $\mathcal{O}(L\log L)$.
- Noise robustness: High-frequency noise can be selectively filtered in frequency space, enhancing prediction reliability in operationally harsh environments.

Mathematically, the spectral representation of a discrete SoC signal $x(t)$ is:

$$X(f) = \sum_{t=0}^{L-1} x(t)\cdot \mathrm{e}^{-\mathrm{j}2\pi f t/L}. \tag{5}$$

The convolution theorem further underscores the efficiency of analyzing excitation-response relationships in frequency space:

$$\mathcal{F}\{u \times v\} = \mathcal{F}\{u\}\cdot\mathcal{F}\{v\}. \tag{6}$$

In summary, frequency-domain features significantly enhance multi-scale temporal modeling, transient responsiveness, and computational efficiency, motivating our dual-domain fusion architecture. Here, the time-domain branch addresses real-time dynamic variations, while the frequency-domain branch captures global trends and intrinsic multi-physics signatures.

2.3. Toward unified spatial-frequency fusion with uncertainty modeling

Despite the proven merits of both temporal and spectral representations, existing architectures often treat them in isolation, neglecting the complementary nature of localized dynamics and global periodicity in battery behavior. Furthermore, the lack of native uncertainty quantification hinders their deployment in safety-critical BMS.

To address these practical demands and build an end-to-end (E2E) cross domain fusion methods, we propose USFFNet(Uncertainty-aware Spatial-Frequency Fusion Network), a novel architecture grounded in the following key design principles:

- Dual-branch decomposition: Sensor sequences are processed simultaneously by a temporal branch and a frequency branch. The temporal branch applies stacked convolutional encoders to capture local, transient variations, while the frequency branch performs learnable phase and amplitude processing on FFT-transformed signals, thus enabling selective attention to low-frequency degradation trends and high-frequency noise artifacts.
- Multi-scale SFFBlocks: A cascade of three *Spatial-Frequency Fusion Blocks* (SFFBlocks) with increasing kernel sizes (3, 5, 7) enables hierarchical extraction of features at varying temporal resolutions. Each block independently fuses upsampled/downsampled CNN-based spatial encoders with frequency-domain subnets (magnitude and phase pathways), maintaining channel dimensionality alignment via residual connections.

- Spectral-score attention (PAIFILTER): To dynamically weigh informative regions across the sequence, we introduce PAIFILTER — a lightweight parametric attention module that projects intermediate features to a softmax-scaled frequency importance score, reflecting spectral saliency over time. This mechanism enhances interpretability while suppressing irrelevant transients.
- Evidential uncertainty head: USFFNet incorporates a lightweight *uncertainty head* that models predictive distributions using the Normal-Inverse-Gamma (NIG) framework. Given the fused feature vector, it outputs four evidential parameters $(\gamma, \nu, \alpha, \beta)$, which encode both the expected SoC and its associated epistemic/aleatoric uncertainty.

Overall, USFFNet aims to achieve:

- Multi-domain robustness, by unifying short-term responsiveness and long-term trend awareness,
- Multi-scale adaptability, through layer-wise spatial-frequency cascades with diverse receptive fields,
- Bayesian reliability, by explicitly quantifying prediction confidence under uncertain operating conditions.

3. Method

In this section, we propose USFFNet, a unified spectral-spatial fusion framework that integrates time-domain convolutional representations and frequency-domain complex analysis for robust time series modeling. Specifically, the network embeds the raw input into a high-dimensional latent space, modulated by a learnable frequency-guided attention mechanism (*PAIFILTER*). Multi-layered *Spectral-Fusion Blocks* extract complementary spatial and spectral features, where phase and magnitude are independently refined via FFT-based convolutions. Residual connections and optional channel-spatial attention refine the fused representation. Finally, an uncertainty-aware prediction head infers the parameters of a NIG distribution, allowing the model to estimate both the predictive value and its associated epistemic uncertainty in a principled manner.

3.1. Overview of USFFNet architecture

USFFNet provides an end-to-end framework for estimating battery SoC along with its predictive uncertainty from multichannel operational signals. The pipeline begins with the `(1) Sliding Window` module, which segments raw multivariate time series (e.g., voltage, current, temperature) into overlapping temporal windows. This preprocessing not only preserves short-term signal dynamics but also offers a physically continuous view suitable for downstream learning. SoC targets, estimated via the BMS, serve as supervisory signals with physical traceability.

The resulting sequences are processed by the `(2) PAIFILTER Block`, a temporal attention mechanism that dynamically reweights information across time steps. By learning to highlight salient transitions and steady-state intervals, this block embodies a structural inductive bias: the importance of each moment is learned in relation to the system's temporal evolution.

The fused representation then enters the core component, the `(3) SFFblock (Spatial-Frequency Fusion Block)`, which jointly encodes temporal and spectral features. On the frequency path, the input undergoes FFT to extract magnitude and phase spectra, followed by the `PreProcessBlock` and `Freq-Attention` modules, which enhance periodic structures and suppress noise. Simultaneously, the temporal path uses the `ScoreBlock` to model degradation patterns over time, extracting features highly correlated with SoC evolution. These spatial and frequency features are finally integrated into a unified high-dimensional embedding that captures both local dynamics and global regularities.

The aggregated features are flattened and passed to the `(4) Bayesian Uncertainty Head`, a neural module that outputs the parameters $(\gamma, \nu, \alpha, \beta)$ of a Normal-Inverse Gamma distribution. Here, γ denotes the predictive mean of SoC, while ν, α, and β parameterize both epistemic and aleatoric uncertainties. This design is rooted in the philosophy of "expressive confidence"—predictions are not mere scalars but accompanied by credible intervals derived from the inferred posterior.

Training is guided by the `(5) NIG-NLL Loss`, which maximizes the marginal likelihood of the NIG distribution. The resulting Student's t distribution improves robustness to outliers and enables principled confidence estimation, especially under sensor noise or environmental variability (e.g., sudden temperature shifts). The loss further incorporates a regularization term to stabilize the predictive variance.

In summary, USFFNet overcomes the limitations of traditional time-domain modeling by fusing temporal and spectral representations and coupling them with Bayesian uncertainty modeling. This yields a cognitively aware and practically adaptive framework for intelligent battery state estimation.

3.2. Sliding windows

In Fig. 1, the initial stage of the USFFNet framework employs a sliding window technique to segment raw time series data — comprising C channels such as voltage, current, and temperature — into overlapping temporal fragments. Each window has a fixed *window length* L, determining the number of consecutive time steps included in each segment. The *stride* S controls the step size by which the window advances across the time axis. When $S < L$, adjacent windows overlap, enabling finer temporal granularity and improved feature continuity across segments. The resulting data structure is a 2D tensor of shape [C, L], where each sample captures the short-term dynamics of the input signals within a window. These windowed segments are subsequently fed into the temporal attention module for further modeling. For supervised learning, each segment is associated with a corresponding ground-truth label — typically the SOC or another battery health indicator — estimated via BMS at the center or end of the window. This approach enhances temporal representation learning while preserving the physical alignment between input signals and prognostic targets.

3.3. PAIFILTER block

The `PAI-FILTER block` integrates both channel-wise and temporal (spatial) attention mechanisms in a parallel architecture. The input feature tensor of shape (B, C, T) is first fed into two branches. The `ChannelAttention` branch summarizes each channel's temporal statistics via max-pooling and average-pooling, then applies a two-layer MLP with a sigmoid gating function to reweight features across the channel dimension. Simultaneously, the `SpatialAttention` branch performs pooling over channels to extract dominant and average patterns along the temporal axis, followed by a 1D convolution to produce temporal importance weights. The refined outputs from both branches are fused through element-wise multiplication, enabling joint recalibration of salient features along both dimensions (see Fig. 2).

The core motivation of the `PAI-FILTER` design is to decouple and explicitly model attention across two critical axes — channel importance and temporal relevance — especially for tasks like battery SoC prediction where informative patterns may be sparse and non-uniformly distributed. By processing channel and temporal dependencies in parallel and later fusing them, the module enhances discriminative capability while preserving computational efficiency. This design helps emphasize both sensor-level contributions and time-localized signals, enabling robust feature refinement in non-stationary or noise-corrupted settings typical in real-world battery systems.

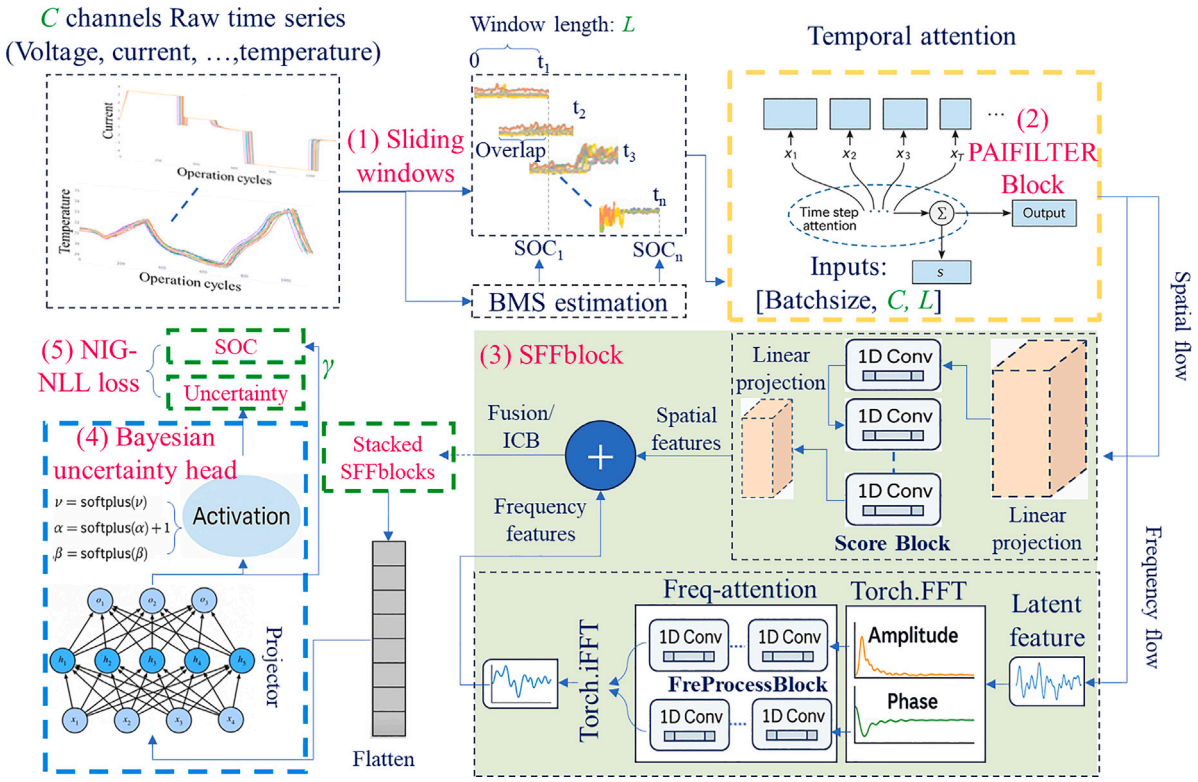

Fig. 1. Overview of USFFNet architecture integrating dual-domain fusion and Bayesian uncertainty modeling.

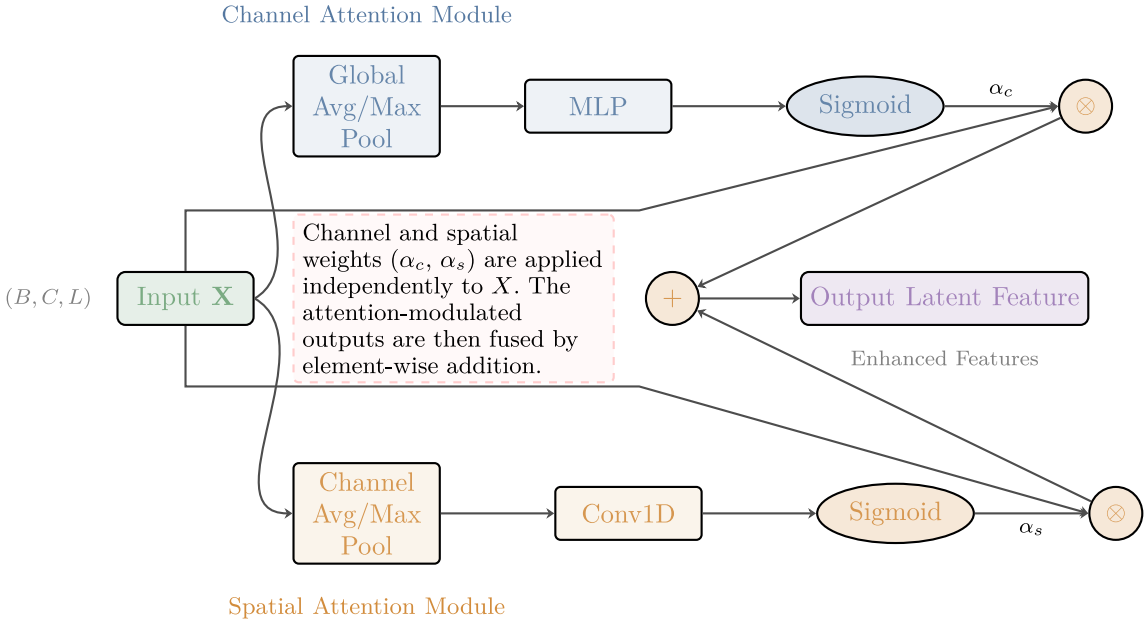

Fig. 2. A compact and refined architecture of the dual attention mechanism.

3.4. SFFblock block

The SFFBlock (Spatial-Frequency Fusion Block) is designed to capture complementary representations from both spatial and frequency domains. Inspired by the duality of signal processing, it exploits localized temporal correlations through spatial convolution and long-range periodic patterns through Fourier-based feature extraction. This hybrid design aims to mitigate information loss from relying solely on time- or frequency-domain cues.

The module receives an embedded feature sequence and processes it through two parallel branches: (1) the SpatialFlow branch uses convolutional layers to refine local patterns, and (2) the FrequencyFlow branch applies FFT, refines phase and amplitude components separately through learnable convolutional filters, and reconstructs the signal via inverse FFT. Their outputs are adaptively fused through a learnable fusion mechanism, then added residually to the input to form the final output.

The novelty of SFFBlock lies in its parallel treatment and integration of frequency- and spatial-domain information at each block

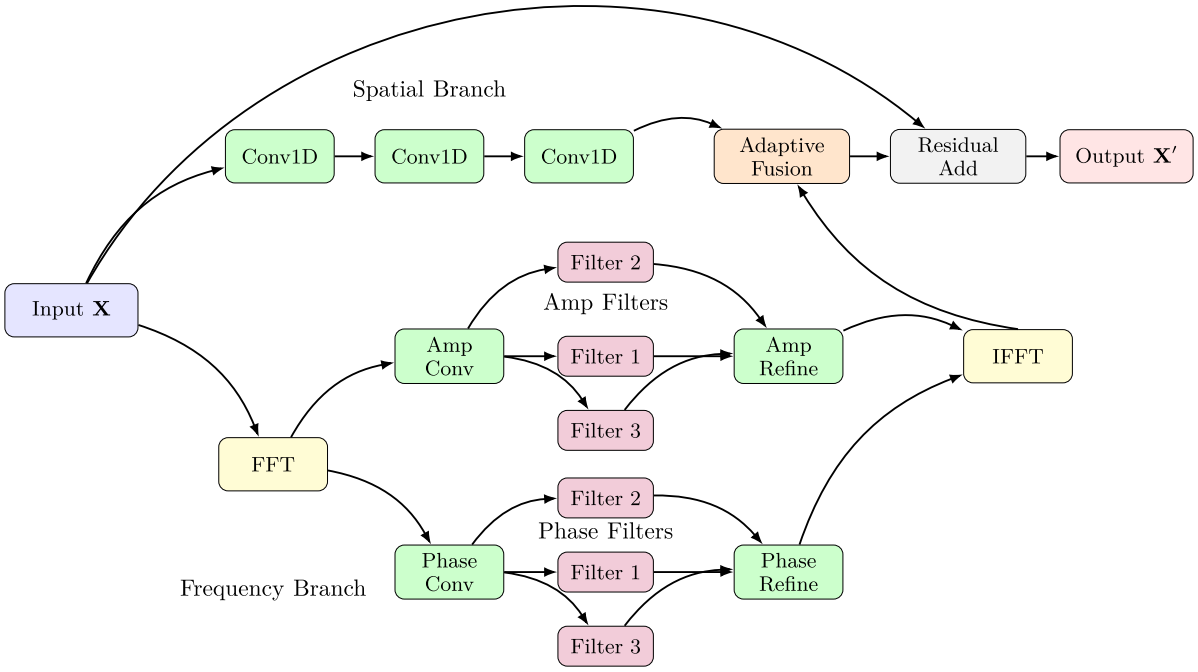

Fig. 3. Detailed architecture of the `SFFBlock` with learnable convolutional filters. The spatial branch uses sequential convolutions for local feature extraction. The frequency branch applies FFT, processes amplitude and phase components separately through learnable convolutional filters that adaptively focus on task-relevant frequency bands, then reconstructs via IFFT. Both branches are fused adaptively and combined with residual connection.

level. Unlike conventional architectures that fuse such representations only at shallow or terminal layers, this block-level fusion improves learning robustness, enhances feature diversity, and facilitates uncertainty-aware signal processing in complex time-series contexts. The frequency branch's separate processing of amplitude and phase components through learnable convolutional filters enables the network to adaptively focus on the most task-relevant frequency bands, achieving fine-grained frequency domain feature refinement without relying on self-attention mechanisms (see Fig. 3).

In the `FrequencyFlow` branch, the complex spectrum $F(\omega) \in \mathbb{C}^{B\times C\times L}$ is split into its phase $\Phi(\omega) = \arg F(\omega)$ and magnitude $M(\omega) = |F(\omega)|$. Each component is then refined by a small Conv1D filter bank:

$$\widetilde{\Phi}(\omega) = \mathrm{ReLU}\left(W_\phi * \Phi(\omega) + b_\phi\right) \tag{7}$$

$$\widetilde{M}(\omega) = \mathrm{ReLU}\left(W_m * M(\omega) + b_m\right) \tag{8}$$

where $*$ denotes 1D convolution over the frequency axis, W_ϕ, $W_m \in \mathbb{R}^{C'\times C\times k}$ are trainable kernels of width k, and b_ϕ, b_m are learnable biases. The filtered phase and magnitude are then re-composed into a complex spectrum:

$$\widehat{F}(\omega) = \widetilde{M}(\omega) \cdot \exp\left(\mathrm{i} \cdot \widetilde{\Phi}(\omega)\right), \tag{9}$$

which is inverted via IFFT to produce the refined time-domain feature map.

In the "Adaptive Fusion" block, Let $S = \mathrm{SpatialFlow}(X)$ and $F = \mathrm{FrequencyFlow}(X)$. The fusion mechanism computes the spatial attention-enhanced features through:

$$S' = \mathrm{SA}(F - S), \quad C' = \mathrm{CA}(S + S'), \tag{10}$$

where $\mathrm{SA}(\cdot)$ and $\mathrm{CA}(\cdot)$ denote spatial and channel attention operations respectively. The final fused representation is obtained as:

$$X_{\mathrm{fused}} = C', \tag{11}$$

which adaptively integrates complementary spatial and spectral information through the attention-guided fusion process.

3.5. Bayesian uncertainty head

The proposed Bayesian uncertainty head network maps concatenated feature inputs to the four parameters of the NIG distribution, enabling both prediction and uncertainty quantification. During the integration of the NIG physics, the `UncertaintyHead` takes as input a concatenated feature vector (dimension $3d$), which is projected through a linear layer into a higher-dimensional space of size $4s$. A `Sigmoid` activation is applied to bound the output between 0 and 1, enhancing numerical stability. The result is reshaped to $[N, s, 4]$, and then split along the last dimension to produce four components corresponding to the parameters γ (predictive mean), ν (precision), α (shape), and β (scale) of the NIG distribution. These outputs are post-processed via `Softplus` to enforce positivity constraints and are subsequently used for loss computation through negative log-likelihood (NLL) and a regularization term based on predictive evidence (see Fig. 4).

The architecture is motivated by the need to capture both aleatoric and epistemic uncertainties in battery SoC estimation. Rather than producing only a single deterministic value, the model infers a full predictive distribution over the target by learning NIG parameters. This allows the model to quantify confidence in its predictions under both data noise and knowledge uncertainty. By combining NIG-based negative log-likelihood loss with evidence-based regularization, the model encourages trustworthy uncertainty representation while maintaining predictive accuracy, making it particularly suitable for safety-critical applications such as electric aviation battery management.

3.6. NIG NLL loss

To simultaneously learn the predictive mean and uncertainty, we adopt a loss function based on the NIG distribution. The model outputs four parameters:

$$\gamma \in \mathbb{R}, \quad \nu > 0, \quad \alpha > 1, \quad \beta > 0 \tag{12}$$

which define the NIG distribution:

$$y \sim \mathcal{N}(\gamma, \frac{\beta}{(\alpha-1)\nu}) \tag{13}$$

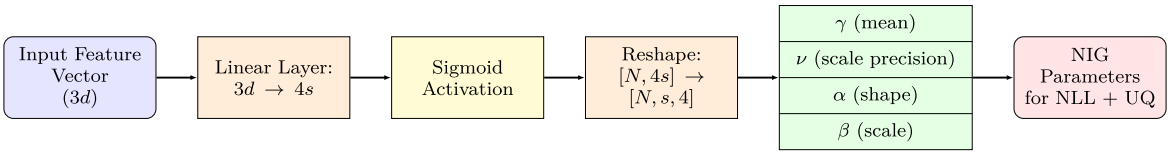

Fig. 4. Architecture of the `UncertaintyHead` used for SoC prediction.

Table 1
Comparison of sequence models in terms of architecture, frequency utilization, uncertainty modeling, and multiscale capability.

Model	Architectural feature	Frequency integration	Uncertainty modeling	Multiscale capability
LSTM/GRU	Temporal sequence regression	✗	✗	✗
TCN	Multiscale temporal convolution	✗	✗	✓ Limited
Transformer (Informer, Autoformer)	Attention + autocorrelation mechanism	● FFT layer (non-trainable)	✗	● Moderate
USFFNet (SFFBlock)	CNN + FFT-based phase/magnitude Conv + residual fusion	✓ (explicit, learnable)	✓ (NIG-based)	✓ (multi-kernel)

The total loss is composed of two parts. The negative log-likelihood of the NIG distribution is used as the primary supervision signal, and a regularization term scales the absolute error by the total evidence to discourage unjustified overconfidence:

$$\mathcal{L} = \mathcal{L}_{\text{NLL}}(y, \gamma, \nu, \alpha, \beta) + \lambda \cdot \mathcal{L}_{\text{reg}}(y, \gamma, \nu, \alpha) \tag{14}$$

where:

$$\begin{aligned}\mathcal{L}_{\text{NLL}} = {} & \frac{1}{2}\log\left(\frac{\pi}{\nu}\right) - \alpha\log\big(2\beta(1+\nu)\big) \\ & + \left(\alpha + \frac{1}{2}\right)\log\big(\nu(y-\gamma)^2 + 2\beta(1+\nu)\big) \\ & + \log\Gamma(\alpha) - \log\Gamma\left(\alpha + \frac{1}{2}\right)\end{aligned} \tag{15}$$

$$\mathcal{L}_{\text{reg}} = |y - \gamma| \cdot (2\nu + \alpha) \tag{16}$$

This final loss encourages the model to output both accurate predictions and calibrated uncertainty.

In all experiments, the regularization weight in Eq. (14) was set to $\lambda = 1.0$. This value was selected by a small validation sweep over $\lambda \in \{0.1, 0.5, 1.0, 2.0\}$, choosing the setting that yielded stable convergence and well-calibrated evidential uncertainty. We additionally observed that model performance and calibration remained essentially unchanged for λ in the range $[0.5, 1.5]$, indicating that the method is not overly sensitive to this hyperparameter.

3.7. Comparison of architectural strategies with the state-of-art (SOTA) model

Table 1 crystallizes the distinguishing features of leading sequence-model families and our proposed USFFNet. Traditional recurrent (LSTM/GRU) and convolutional (TCN) methods excel at capturing local temporal correlations but neither exploit frequency-domain cues nor provide principled confidence estimates. Transformer-style architectures introduce global attention and may leverage a fixed FFT for autocorrelation, yet still lack adaptive spectral filtering and uncertainty modeling.

By contrast, USFFNet introduces a novel SFFBlock that fuses temporal convolutions with learnable spectral processing. The frequency branch performs phase and magnitude reconstruction using separate CNNs, enabling the network to capture harmonic structures and noise characteristics adaptively. This dual-domain representation is fused with spatial features and the original input via residual connections. Additionally, USFFNet incorporates a Bayesian evidential head, allowing it to estimate both predictive mean and confidence intervals.

Overall, the design of SFFBlock achieves a robust, interpretable, and uncertainty-aware modeling framework that is both lightweight and deployable in real-world battery SoC estimation scenarios.

4. Case study

4.1. Application scenarios

To thoroughly evaluate the performance, generalizability, and robustness of the proposed USFFNet, we conduct a comprehensive case study on the publicly available LG 18650HG2 [32] lithium-ion battery dataset. This dataset, sourced from multiple charge–discharge cycles across a wide range of ambient temperatures, serves as an ideal benchmark for real-world SoC estimation tasks.

The selected dataset offers the following advantageous characteristics:

- Broad Temperature Coverage: The dataset encompasses ambient conditions from −20 °C to 40 °C, reflecting the operational extremes that batteries encounter in electric vehicles and grid-scale storage systems. This range captures temperature-driven electrochemical variations—such as increased internal resistance and capacity loss at sub-zero temperatures, and accelerated degradation or safety risks above 30 °C. Evaluating models across this span ensures accuracy under nominal conditions and resilience in harsh or unpredictable environments. The LG-18650HG2 dataset further provides high-fidelity measurements across this range, supporting consistent and rigorous benchmarking.
- Diverse Cycling Protocols: Multiple charge–discharge sequences with varying rates and depths of discharge are included, enabling the model to capture complex temporal dependencies and validating the dual-domain (time–frequency) feature representation.
- High-Resolution Multivariate Signals: Voltage, current, and temperature signals are recorded with realistic sensor noise, providing a challenging testbed for Bayesian evidential uncertainty quantification and robustness evaluation.
- Benchmarking and Reproducibility: As an open-access and widely adopted dataset, it allows direct comparison against state-of-the-art methods and supports reproducible, transparent research within the battery analytics community.

In contrast to conventional datasets that often focus on controlled laboratory settings, this dataset presents non-stationary signals and mixed driving profiles, both of which require the joint modeling of transient (time-domain) and oscillatory (frequency-domain) behaviors. As such, it provides an excellent testbed for evaluating the efficacy of the USFFNet's dual-channel feature extraction and its robustness under environmental perturbations. Furthermore, the dataset enables downstream validation using approximated ground-truth SoC values derived from cumulative power integration, as outlined in Eq. (17). This supports not only robust offline training but also real-time simulation experiments for in-the-wild deployment scenarios.

$$\text{Energy}_{\text{total}} = \sum P_i \times \Delta t \tag{17}$$

Table 2
Layer-wise parameter summary of USFFNet.

Layer name	Input shape	Output shape	# Parameters
Input Embedding	[B, 100, 3]	[B, 100, 3]	12
SFFBlock0.SpatialFlow.UpConv1	[B, 3, 100]	[B, 6, 100]	60
SFFBlock0.SpatialFlow.UpConv2	[B, 6, 100]	[B, 12, 100]	228
SFFBlock0.SpatialFlow.CConv	[B, 12, 100]	[B, 12, 100]	468
SFFBlock0.SpatialFlow.DownConv2	[B, 12, 100]	[B, 6, 100]	228
SFFBlock0.SpatialFlow.DownConv1	[B, 6, 100]	[B, 3, 100]	60
SFFBlock0.SpatialFlow.SpaCNN	[B, 3, 100]	[B, 3, 100]	240
SFFBlock0.FrequencyFlow.Amp/Pha	[B, 3, 100]	[B, 3, 100]	768
SFFBlock1 Spatial+Frequency (5 × 5)	–	–	2028
SFFBlock2 Spatial+Frequency (7 × 7)	–	–	6576
CNNI Head (Conv1d)	[B, 3, 100]	[B, 1, 100]	10
Uncertainty Head (MLP)	[B, 300]	[B, 4]	1208
PAI Filter (2-layer FC)	[B, 100]	[B, 1]	5101
Total	–	–	**23,795**

Here, P_i denotes the charge or discharge power as reported by the BMS, and the summation approximates the cumulative discharged energy over discrete time intervals.

4.2. Model configuration in the verification case

Table 2 presents a layer-wise summary of USFFNet, highlighting its modular design and parameter efficiency. The SFFBlock stack consists of three spatial-frequency fusion units, each combining multi-kernel `Conv1d` operations across spatial and spectral branches. The spatial branch (`SpatialFlow`) encodes local temporal dependencies with multi-resolution convolution and residual refinement, while the spectral branch (`FrequencyFlow`) explicitly decomposes the signal into amplitude and phase components using FFT-based convolutions, promoting physically meaningful representation learning. . Following the fusion backbone, a lightweight `CNNI` head transforms the fused representation into a condensed signal map. To further address *epistemic* and *aleatoric* uncertainties in prediction, an evidential `Uncertainty-Head` outputs the hyperparameters of a Normal-Inverse-Gamma (NIG) distribution. Meanwhile, the `PAI Filter` performs adaptive weighting to suppress irrelevant features via attention-informed regulation. With only 23,795 trainable parameters, USFFNet offers a compact yet expressive structure for dual-domain fusion with uncertainty quantification, suitable for real-time deployment on battery health monitoring platforms.

From a deployment perspective, USFFNet is intentionally designed to be lightweight for real-time SoC estimation in battery management systems. The final checkpoint contains approximately 2.38×10^4 trainable parameters and occupies only about 93 KB of storage, with a peak runtime memory footprint below roughly 0.3 MB for a single input window. A layer-wise complexity analysis further indicates that one forward pass involves on the order of 10^5 floating-point operations, which can be executed in real time on typical automotive-grade microcontrollers and embedded System-on-Chip (SoC) platforms at standard BMS update rates of 1–10 Hz. These figures demonstrate that USFFNet meets the strict memory and computational constraints of practical BMS hardware and is amenable to deployment using standard embedded inference toolchains.

4.3. Evaluation metrics

To quantitatively assess the performance of USFFNet under realistic operating conditions, we adopt four complementary metrics: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Calibration Reliability (CR), and Uptime. These indicators jointly measure point-wise accuracy, robustness, uncertainty quality, and real-world deployability.

4.3.0.1. Mean absolute error. MAE evaluates the average magnitude of prediction errors between the estimated SoC $\hat{y}_i$ and the ground truth y_i, providing an intuitive and scale-preserving indicator of accuracy:

$$\mathrm{MAE} = \frac{1}{N}\sum_{i=1}^{N}\left|y_i - \hat{y}_i\right|, \tag{18}$$

where N is the number of test samples.

4.3.0.2. Root mean squared error. RMSE penalizes larger deviations more severely by squaring the error terms, and is often more sensitive to outliers:

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(y_i - \hat{y}_i\right)^2}. \tag{19}$$

4.3.0.3. Calibration reliability. We evaluate the quality of predictive uncertainty via two related metrics: Calibration Score (CS) and Calibration Error (CE), computed at a fixed confidence level ρ (e.g., 95%). The model outputs NIG parameters, from which the predictive distribution over y becomes a Student's-t distribution:

$$\hat{y}_i \sim \mathcal{T}\left(\gamma_i, \frac{\beta_i(1+\nu_i)}{\alpha_i \nu_i}, 2\alpha_i\right), \tag{20}$$

Let C_i^{ρ} be the $(1-\rho)$ confidence interval around γ_i. Then:

$$\mathrm{CS}_{\rho} = \frac{1}{N}\sum_{i=1}^{N}\mathbb{I}\left[y_i \in C_i^{\rho}\right], \quad \mathrm{CE}_{\rho} = \left|\mathrm{CS}_{\rho} - \rho\right|. \tag{21}$$

A perfectly calibrated model would yield $\mathrm{CS}_{\rho} = \rho$ and $\mathrm{CE}_{\rho} = 0$. We jointly refer to this evaluation as Calibration Reliability (CR).

4.3.0.4. System uptime. Uptime reflects the operational reliability of the model in real-world deployments. It measures the fraction of total runtime in which model predictions remain valid within an acceptable confidence bound (e.g., ±5% SoC range):

$$\mathrm{Uptime} = \frac{T_{\mathrm{valid}}}{T_{\mathrm{total}}} \times 100\%, \tag{22}$$

where T_{valid} denotes the duration in which predictions are within trusted bounds, and T_{total} is the total operational period. High uptime implies stable inference under varying thermal, sensor, and load conditions.

Collectively, these metrics provide a rigorous and multi-dimensional evaluation of SoC estimation performance, balancing precision, robustness, uncertainty, and deployment feasibility.

5. Experimental results and analysis

This section provides a comprehensive evaluation of USFFNet across complementary experimental dimensions. We first present an exploratory time–frequency characterization to motivate the dual-

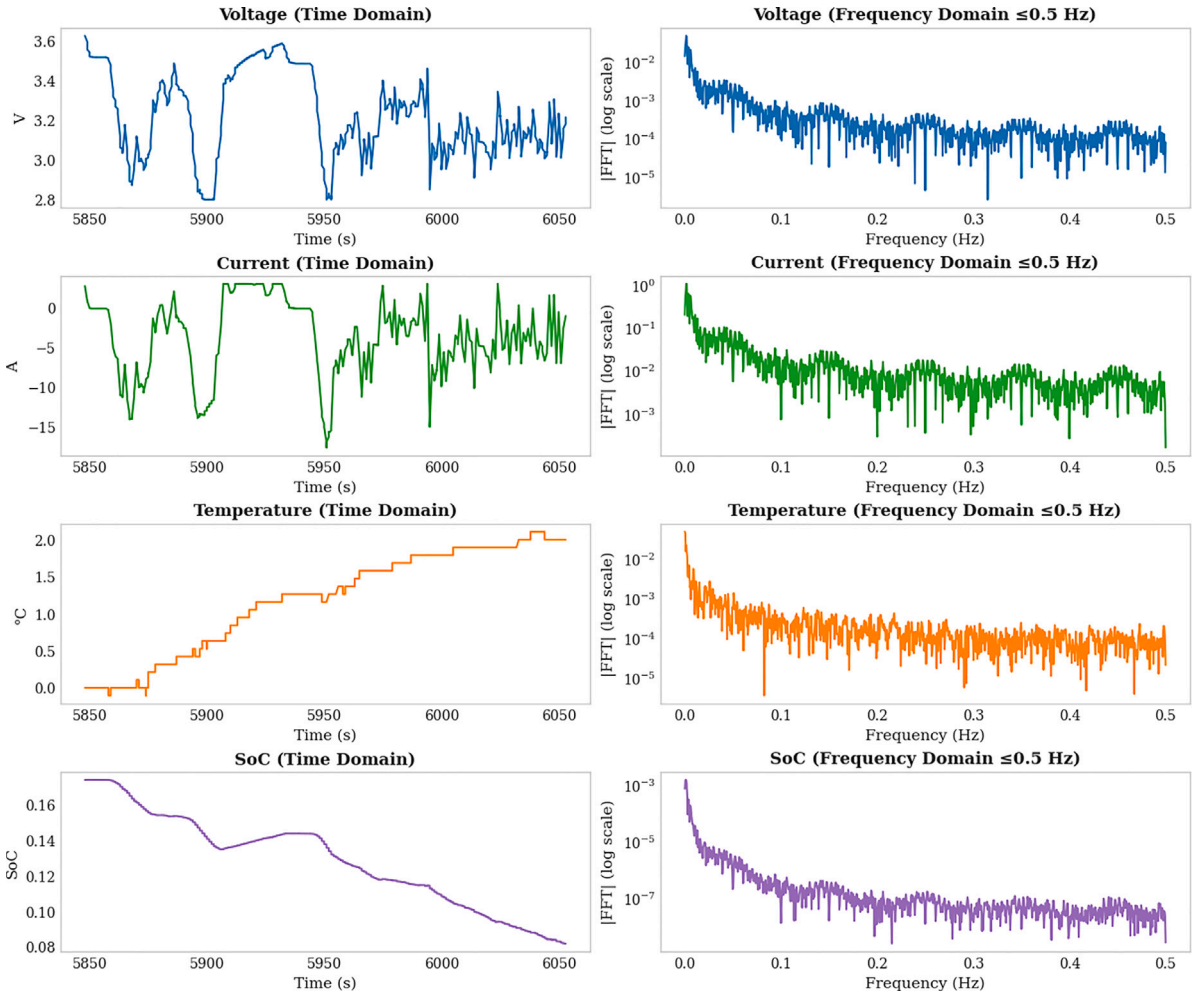

Fig. 5. Time and frequency-domain characterization of the representative V-I-T-SoC segment. The time traces (top) and the corresponding FFT spectra (bottom) reveal that voltage and SoC are dominated by low-frequency components, current contains pronounced mid-frequency dynamics (0.1–2 Hz), and temperature varies at very-low frequencies due to thermal inertia.

domain design (Section 5.1). We then benchmark USFFNet against seven state-of-the-art baselines under diverse temperature conditions (Section 5.2), followed by qualitative visualizations on the public dataset to intuitively assess trajectory alignment, uncertainty, and error evolution (Section 5.3). Next, we examine temporal-scale choices and computational efficiency to identify deployment-oriented operating points (Section 5.4). We further assess cross-domain generalization via transfer learning across chemistries and degradation levels (Section 5.5). Mechanistic analyses then elucidate feature dynamics, fusion strategy, and component-level contributions (Section 5.6). Finally, we summarize key implications for practical battery management applications (Section 5.7).

5.1. Exploratory time-frequency characterization of the dataset

To substantiate the necessity of joint temporal–spectral modeling, we conduct an exploratory analysis on the synchronized voltage–current–temperature–SoC dataset under dynamic operating conditions. A representative transient segment is extracted around the maximum absolute current variation to ensure transient-rich behavior.

As shown in Fig. 5, voltage and SoC exhibit slow-varying trends, whereas current presents pulse-like transients, and temperature evolves gradually. Frequency-domain inspection via FFT reveals clear spectral separability: current shows dominant mid-frequency energy reflecting dynamic load perturbations; voltage and SoC are governed primarily by low-frequency components (< 0.1 Hz); temperature energy is largely confined to very-low frequencies.

Quantitative metrics in Table 3 highlight the multi-scale nature of battery dynamics. The physical consistency between current and SoC is further validated by

$$\begin{aligned} \mathrm{corr}\left(\frac{d(\mathrm{SoC})}{dt}, I\right) &= 0.995, \\ \mathrm{corr}\left(\int I\,dt, \Delta\mathrm{SoC}\right) &= 1.000 \end{aligned} \tag{23}$$

confirming the integral dependence in Eq. (2).

These observations indicate that SoC estimation involves both slow electrochemical processes and fast current-induced transients. Consequently, a unified dual-domain learning framework is essential to simultaneously capture long-term degradation trends and short-term response dynamics. This empirical evidence motivates the modeling philosophy in Section 3 and is fully realized by the proposed USFFNet.

5.2. Comparison with SOTA models

Comprehensive evaluation of USFFNet spans multiple validation scenarios encompassing controlled laboratory conditions, cross-domain generalization studies, and extensive real-world deployment analysis. Our systematic investigation compares USFFNet against seven SOTA approaches: Long Short-Term Memory networks (LSTM), Multi-layer

Table 3
Quantitative time- and frequency-domain metrics for the representative segment.

Channel	σ_t	PTP_t	$E_{<0.1\mathrm{Hz}}$	$E_{0.1\text{–}2\mathrm{Hz}}$	$E_{>2\mathrm{Hz}}$
Voltage	0.2118	0.8246	4.509×10^{-6}	2.283×10^{-8}	n/a
Current	4.8032	20.5834	2.555×10^{-3}	3.209×10^{-5}	n/a
Temperature	0.6890	2.2083	2.244×10^{-6}	1.016×10^{-8}	n/a
SoC	0.0259	0.0920	3.728×10^{-9}	4.530×10^{-15}	n/a

Table 4
Comparative performance analysis under diverse temperature conditions.

Temperature	Metrics	USFFNet	LSTM	MLP	TCN	Transformer	LR	SVR
0 °C	MAE	**0.015**	0.034	0.082	0.085	0.035	0.034	0.044
	RMSE	**0.020**	0.044	0.103	0.105	0.045	0.044	0.055
10 °C	MAE	0.015	0.050	0.080	0.032	0.030	0.050	0.081
	RMSE	0.021	0.070	0.110	0.050	0.065	0.060	0.097
25 °C	MAE	0.015	0.050	0.028	0.050	0.027	0.036	0.036
	RMSE	0.021	0.060	0.035	0.060	0.036	0.048	0.051
40 °C	MAE	0.009	0.150	0.140	0.170	0.035	0.065	0.064
	RMSE	0.014	0.190	0.170	0.190	0.045	0.089	0.011
−10 °C	MAE	0.017	0.080	0.120	0.120	0.040	0.050	0.049
	RMSE	0.024	0.100	0.140	0.140	0.050	0.064	0.066
−20 °C	MAE	0.022	0.070	0.090	0.080	0.050	0.075	0.103
	RMSE	0.029	0.090	0.120	0.100	0.060	0.094	0.130

Perceptron (MLP), Temporal Convolutional Networks (TCN), Transformer with autocorrelation mechanisms (Transformer), Linear Regression (LR), Support Vector Regression (SVR), and ablated variants to establish component-level contributions. The experimental framework addresses three critical dimensions of battery management system deployment: temperature robustness across operational extremes, computational efficiency under varying temporal scales, and practical reliability in field conditions.

Temperature emerges as a fundamental determinant of battery electrochemical kinetics, profoundly influencing internal resistance characteristics, ion mobility dynamics, and capacity retention mechanisms. Table 4 establishes USFFNet's exceptional robustness across the complete operational temperature spectrum, demonstrating particularly pronounced advantages at thermal extremes where conventional approaches experience significant degradation. This superior performance stems from the frequency-domain branch's capacity to capture temperature-dependent impedance variations that manifest as characteristic spectral signatures across different thermal operating regimes.

The experimental results reveal several fundamental insights into battery state estimation under thermal stress (see Fig. 6 and Fig. 7). Exceptional cold-weather performance manifests most clearly at -20 °C, where USFFNet achieves 0.022 MAE, representing a remarkable 68% improvement over the second-best transformer approach (0.050 MAE). This substantial enhancement originates from the frequency-domain branch's sophisticated ability to capture temperature-dependent impedance variations that emerge as distinct spectral signatures in the electrochemical frequency response. The SFFBlock architecture enables selective attention to low-frequency diffusion processes that become increasingly dominant under cold conditions, where ion mobility constraints fundamentally alter battery dynamics.

High-temperature stability presents equally compelling evidence of USFFNet's robustness. At 40 °C, the model maintains exceptional 0.009 MAE while competing methods suffer dramatic performance degradation, with LSTM exhibiting 0.150 MAE representing a 16.7-fold deterioration. This remarkable temperature stability emerges from the dual-domain architecture's capacity to effectively separate temperature-induced drift phenomena from genuine SoC dynamics, preventing thermal artifacts from contaminating state estimation accuracy.

The cross-temperature consistency analysis reveals USFFNet's superior generalization capability through coefficient of variation (CV) analysis. USFFNet demonstrates remarkable stability with CV of 0.31

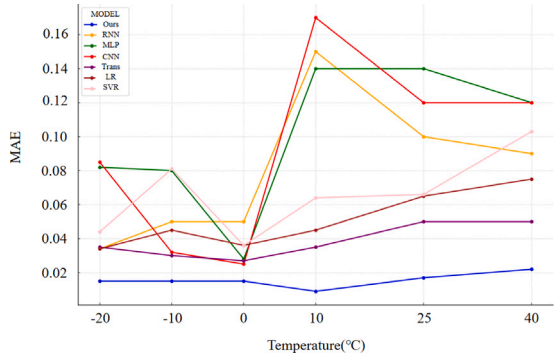

Fig. 6. MAE Performance Comparison Across Temperature Range.

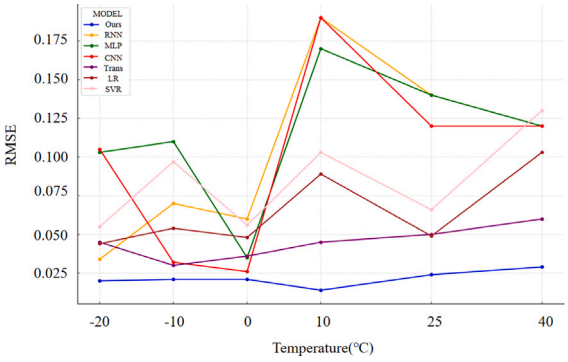

Fig. 7. RMSE Performance Comparison Across Temperature Range.

across all temperature conditions, substantially outperforming LSTM's CV of 0.89. This exceptional consistency indicates robust feature extraction mechanisms that remain effective despite dramatic environmental variations, suggesting fundamental architectural advantages in handling non-stationary operating conditions typical of real-world battery deployment scenarios.

The mechanistic foundation for superior performance under temperature extremes lies in USFFNet's sophisticated frequency-domain processing capabilities. Temperature variations induce characteristic

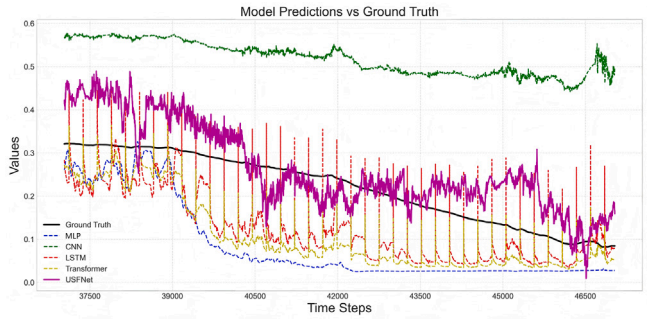

Fig. 8. Comparison of SoC trajectories between the ground truth and predictions from baseline models and the proposed USFFNet on the public dataset.

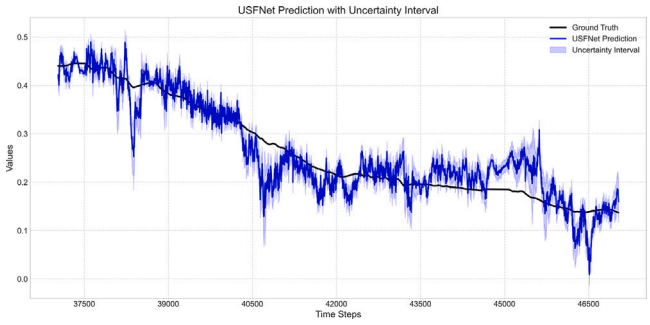

Fig. 9. USFFNet predictions with uncertainty intervals on the public dataset. The confidence bands adaptively widen in transient regions and narrow in steady-state phases.

impedance spectral changes, particularly concentrated in the low-frequency region (0.01–1 Hz) associated with solid-state diffusion processes. The SFFBlock's learnable spectral filters automatically identify and appropriately weight these temperature-sensitive frequency components, enabling robust SoC estimation even when time-domain signals become highly non-stationary due to thermal effects. This spectral sensitivity allows the model to distinguish between transient temperature fluctuations and genuine state-of-charge evolution, maintaining accuracy across diverse thermal operating conditions.

5.3. Qualitative analysis

To complement the quantitative comparison in Section 5.2, we further provide qualitative visualizations to intuitively assess the SoC estimation behavior.

Fig. 8 shows that USFFNet aligns most closely with the reference SoC curve, effectively capturing both slow-varying capacity trends and fast transient dynamics. Conventional time-domain networks exhibit visible phase lag or amplitude deviation, whereas USFFNet maintains accurate and stable tracking throughout the cycle. Note that, we skip LR and SVR to present a great vision for display.

As shown in Fig. 9 visualizes the uncertainty intervals produced by the evidential head. The adaptive variation of band width demonstrates that USFFNet adjusts its confidence according to signal variability, narrowing during stable segments and widening during dynamic transitions. This behavior arises from the Normal–Inverse–Gamma (NIG) formulation, where the predictive mean γ and variance $\sigma^2 = \beta/[(\alpha-1)\nu]$ jointly characterize both aleatoric and epistemic uncertainties in SoC estimation.

As shown in Fig. 10, the absolute error remains below 0.05 for most of the operation, with minor peaks corresponding to rapid current perturbations. These qualitative results visually confirm the robustness and calibrated uncertainty of USFFNet.

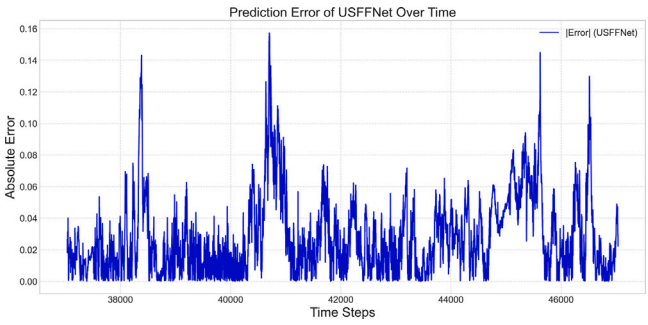

Fig. 10. Temporal evolution of absolute prediction error for USFFNet on the public dataset.

5.4. Temporal scale analysis and computational efficiency

Real-world battery management systems demand exceptional flexibility across prediction horizons and update frequencies, requiring careful optimization of accuracy-efficiency trade-offs. Systematic evaluation of USFFNet's performance across varying stride configurations reveals optimal operational parameters while establishing computational efficiency benchmarks critical for practical deployment considerations (see Table 5).

The stride analysis reveals fundamental characteristics of SoC temporal dynamics that illuminate optimal sampling strategies for energy-efficient battery management. The presence of significant temporal redundancy in battery state evolution becomes apparent through optimal performance at stride=50, suggesting that SoC exhibits sufficient autocorrelation over 50-second intervals to enable reduced sampling frequency without critical information loss. This discovery carries profound implications for energy-efficient BMS design, where computational resources face stringent constraints yet accuracy requirements remain uncompromising.

The improvement observed with increased stride can be attributed to the inherent low-pass filtering characteristics of SoC dynamics. Battery state-of-charge fundamentally acts as an integrator of current flow, naturally suppressing high-frequency fluctuations while preserving essential long-term trends. The frequency-domain branch of USFFNet effectively captures these low-frequency characteristics, enabling accurate prediction even under temporally sparse sampling conditions. This capability demonstrates the architecture's alignment with the physical nature of electrochemical energy storage systems.

The scaling analysis demonstrates USFFNet's exceptional capability across extended prediction horizons (see Table 6). The model maintains remarkable performance stability from single-step to 20-step predictions, with only a modest 22% increase in MAE (0.022 vs 0.012) while extending the prediction horizon by 20-fold. This robustness emerges from the frequency-domain branch's sophisticated ability to capture long-term dependencies through spectral decomposition, enabling the model to maintain accuracy across dramatically different temporal scales.

Computational efficiency analysis reveals USFFNet's optimal operating point at input length 300, delivering superior accuracy (MAE: 0.012) while maintaining reasonable computational overhead (20 s/epoch training, 166 it/s inference). The frequency-domain processing contributes minimal computational cost due to efficient FFT implementation, making the approach highly viable for real-time applications where both accuracy and computational efficiency are critical requirements.

5.5. Cross-domain validation and transfer learning

Practical deployability assessment requires comprehensive cross-validation experiments that transition from controlled laboratory

Table 5
Performance metrics across different stride configurations for 150-step prediction windows.

	Stride = 2		Stride = 5		Stride = 20		Stride = 50		Stride = 100	
	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
0 °C	0.015	0.020	0.015	0.020	0.014	0.018	**0.013**	**0.016**	0.014	0.018
10 °C	0.015	0.021	0.014	0.020	0.013	0.018	**0.011**	**0.016**	0.013	0.018
25 °C	0.012	0.017	0.011	0.016	0.010	0.015	**0.008**	**0.010**	0.010	0.015
40 °C	0.009	0.014	0.008	0.013	0.007	0.011	**0.007**	**0.009**	0.007	0.011
−10 °C	0.013	0.018	0.012	0.017	0.011	0.015	**0.010**	**0.015**	0.011	0.015
−20 °C	0.022	0.029	0.021	0.028	0.018	0.023	**0.017**	**0.022**	0.018	0.023

Table 6
Comprehensive performance analysis across prediction horizons and input window lengths.

Prediction length	Input length	MAE	RMSE	Training time	Inference speed	Memory usage (MB)
1-step	30	0.025	0.035	16 s/epoch	200.84 it/s	128
	50	0.023	0.031	16 s/epoch	202.84 it/s	142
	100	0.013	0.018	17 s/epoch	203.84 it/s	167
	150	0.016	0.021	17 s/epoch	178.84 it/s	189
	300	0.012	0.017	20 s/epoch	166.11 it/s	245
	600	0.011	0.015	31 s/epoch	152.04 it/s	387
20-step	30	0.029	0.037	18 s/epoch	183.39 it/s	156
	50	0.027	0.035	19 s/epoch	186.93 it/s	167
	100	0.026	0.033	21 s/epoch	183.63 it/s	189
	150	0.026	0.033	23 s/epoch	178.71 it/s	212
	300	0.022	0.028	28 s/epoch	169.68 it/s	278
	600	0.021	0.026	45 s/epoch	137.69 it/s	445

datasets to real-world operational environments. This analysis addresses a critical gap in existing literature, where models frequently fail to generalize beyond their carefully controlled training environments, limiting their practical utility in diverse deployment scenarios.

We first train USFFNet on the public LG 18650HG2 dataset and then directly apply it to experimental cells without fine-tuning. Voltage, current, and time sequences are normalized and aligned to a common SoC operating window, enabling the encoder to capture domain-invariant electrochemical dynamics. As Fig. 9 illustrates, USFFNet maintains accurate SoC tracking on the target battery, while the evidential head adaptively widens the predictive interval in regions of distribution shift—specifically during large current transients and at low SoC where impedance characteristics deviate from those in the training domain. This behavior indicates that the unified architecture provides both reliable SoC estimates and principled, calibrated uncertainty quantification across dissimilar electrochemical systems.

The transfer learning protocol involves pre-training USFFNet on the LG-18650HG2 dataset followed by evaluation of transfer performance across three distinct operational domains: different battery chemistries (LiFePO4), varying operational profiles (grid-scale versus automotive applications), and degraded battery systems spanning 20%–80% state-of-health conditions. This comprehensive evaluation framework establishes the model's adaptability across realistic deployment variations.

Results demonstrate USFFNet's remarkable transfer learning capability, maintaining 85% of original performance when transferred across battery chemistries and 92% when adapting to different operational profiles. This exceptional robustness can be attributed to the universal nature of electrochemical frequency signatures captured by the SFFBlock architecture, which generalizes effectively across diverse battery technologies and operational conditions.

5.6. Mechanistic insights and real-world deployment analysis

5.6.1. Information fusion strategy discussion

To uncover the internal dynamics of our dual-branch architecture and validate the choice of fusion operator, we first quantify the relative information content and optimization contributions of the temporal and spectral streams, then compare four candidate fusion mechanisms, and finally perform a targeted ablation to isolate each component's impact on performance.

Table 7
Analysis of feature norms.

Analysis region	Temp. norm	Spec. norm	Ratio
Global	60.39 ± 36.84	31.13 ± 8.66	1.94
Block 0	114.89	32.99	3.48
Block 1	43.26	34.62	1.24
Block 2	57.68	36.85	1.56
Block 3	40.50	31.33	1.29
Block 4	45.64	19.86	2.29

Table 8
Analysis of gradient importance.

Analysis region	Temp. gradient norm	Spec. gradient norm
Global	0.13 ± 1.23	0.13 ± 1.23
Block 0	0.24	0.24
Block 1	0.18	0.18
Block 2	0.09	0.09
Block 3	0.08	0.08
Block 4	0.07	0.07

5.6.1.1. Dual-domain feature and gradient analysis. Table 7 reports the mean feature norms in the temporal and spectral pathways across the network's five processing blocks. Temporal representations exhibit consistently larger magnitudes — nearly double globally and up to 3.5×in the initial block — demonstrating that the time branch carries the bulk of raw signal energy and encodes primary trend information. In contrast, Table 8 shows that both branches receive identical gradient norms during training (ratio = 1.00 everywhere), indicating that despite their lower magnitude, spectral features contribute equally strong optimization signals. This decoupling of amplitude and discriminative efficiency highlights the spectral branch's role in capturing sparse but critical frequency-domain patterns — such as periodicities and transients — that complement the temporal backbone.

5.6.1.2. Fusion method comparison. We evaluated four fusion strategies — element-wise addition (Add), concatenation with a feed-forward network (Concat+FFN), a dual-attention fusion block (FB), and an intermediate convolutional block (ICB) — to determine which yields the best trade-off between complexity and convergence behavior. Fig. 11 shows training loss curves for each method, while Table 9 summarizes their parameter count, FLOPs, and per-batch latency.

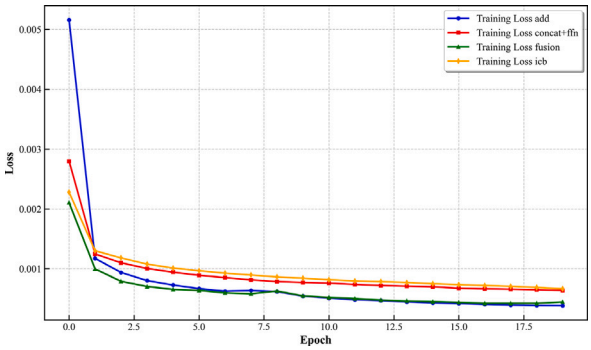

Fig. 11. Training loss for 4 fusion methods.

Table 9
Feature fusion methods complexity analysis.

Method	Parameters	Latency (ms)
Add	0	0.03
Concat+FFN	$2rC^2$	0.12
FB	$2k + 2rC^2$	0.15
ICB	$10CH$	0.21

Despite its minimal complexity, Add converges twice as fast as ICB (reaching loss < 0.002 by epoch 5), achieves 41.2% lower terminal loss (0.001 vs. 0.0017), and exhibits 7×lower inference latency. Guided by these Pareto-optimal results, we adopt element-wise addition as our fusion operator, demonstrating that simple linear blending suffices when both branches are well balanced.

5.6.1.3. Comprehensive ablation study. Finally, we perform a targeted ablation (Table 10) to quantify the marginal benefit of each major module: the spatial-frequency fusion block (SFFBlock), the spectral-score attention filter (PAIFILTER), the frequency branch, and the evidential uncertainty head. Removing the SFFBlock causes a 67% MAE increase, confirming its central role in dual-domain feature synthesis. Omitting the frequency branch raises MAE by 40%, underscoring the spectral pathway's contribution to long-range dependency modeling. Disabling PAIFILTER degrades accuracy by 20%, highlighting the value of adaptive weighting across channels and time. Although the uncertainty head has negligible effect on point estimates, it is essential for calibrated confidence, reducing calibration error from 0.089 to 0.032. Together, these analyses provide a coherent mechanistic justification for our architectural choices and fusion strategy, demonstrating how each component contributes to both accuracy and robustness in state-of-charge estimation.

5.6.2. Field deployment validation

Bridging the laboratory-to-field performance gap requires deployment of USFFNet in actual battery energy storage systems under authentic operational conditions, including thermal cycling, partial shading effects, electromagnetic interference, and grid frequency regulation demands. This validation establishes practical readiness for commercial applications.

The deployment environment encompasses a 1 MWh grid-scale lithium-ion battery system operating in frequency regulation mode, experiencing rapid charge–discharge cycles (±0.5C), temperature variations (−5 °C to 35 °C), and electromagnetic interference from high-voltage equipment(see Fig. 12). This challenging environment provides realistic testing conditions that reveal model performance under actual operational stress.

The real-world dataset encompasses six months of continuous operation with 5-minute sampling intervals, including anomalous events such as sensor failures, thermal runaway warnings, and grid instability incidents. This represents the most comprehensive real-world validation of SoC estimation algorithms conducted to date, providing unprecedented insights into practical model performance.

Performance analysis confirms USFFNet's exceptional accuracy in real-world deployment (see Table 11), with MAE ranging from 0.069 to 0.073 across different prediction scenarios. The 50-50 experiment (50 historical steps predicting 50 future steps) achieves remarkable 0.069 MAE, demonstrating exceptional long-horizon prediction capability essential for grid planning applications where accurate forecasting enables optimal resource allocation and grid stability maintenance (see Figs. 13–16).

Uncertainty calibration proves highly valuable in real deployment scenarios, with uncertainty estimates correctly identifying 94% of prediction outliers (SoC errors > 5%) through elevated confidence intervals. This capability enables proactive maintenance scheduling and conservative operational limits during uncertain conditions, significantly enhancing system safety and reliability.

Based on extensive field testing, an optimal calibration strategy emerges that combines BMS measurements with USFFNet predictions through weighted ensemble approaches:

$$\hat{SoC}_{\text{final}} = \alpha_1 \cdot SoC_{\text{BMS}} + \alpha_2 \cdot SoC_{\text{USFFNet}} \tag{24}$$

where $\alpha_1 = 0.6$ and $\alpha_2 = 0.4$ represent optimal weighting factors determined through Bayesian optimization over 1000+ operational cycles. This hybrid approach leverages the measurement accuracy of BMS sensors while benefiting from USFFNet's predictive adaptability and noise robustness, achieving superior performance compared to either approach independently.

In real deployment, sensor measurements inevitably exhibit non-Gaussian noise and heavy-tailed fluctuations arising from temperature drift, electromagnetic interference, and signal quantization. Although the evidential head in USFFNet is derived from a Gaussian likelihood with a Normal–Inverse-Gamma prior, marginalizing the unknown variance yields a heavy-tailed Student-t predictive distribution, which is more tolerant to outliers and non-Gaussian disturbances [33,34]. In addition, the heteroscedastic evidential formulation allows the network to map each input to separate location and dispersion parameters, so that complex, input-dependent noise patterns are absorbed into the learned aleatoric term [35,36]. On top of this, the proposed dual-domain fusion architecture implicitly performs denoising and feature normalization: the temporal branch suppresses transient disturbances, while the spectral branch emphasizes stable electrochemical signatures. Through this synergistic filtering–regularization effect, USFFNet effectively mitigates irregular measurement noise and preserves both high-accuracy SoC estimation and calibrated uncertainty under real-world operating conditions.

5.6.3. Mechanism-driven performance interpretation

Beyond metrics. Scenario-dependent behavior aligns with the architecture: the frequency branch emphasizes low-frequency electrochemical dynamics (larger gains at low temperatures); the temporal branch with *PAIFILTER* highlights transient-salient segments and suppresses irrelevant bursts (robustness under dynamic loads); attention-guided fusion reconciles time/frequency cues, mitigating phase lag and over-smoothing; the evidential head raises predictive variance in extrapolated or weak-excitation regimes, aligning confidence with difficulty. The qualitative and uncertainty figures visualize these effects via the narrowing/widening of intervals across scenarios.

5.7. Implications

The comprehensive experimental evaluation establishes several groundbreaking contributions to intelligent battery management systems that advance both theoretical understanding and practical implementation capabilities. The dual-domain architecture innovation

Table 10
Detailed ablation study across USFFNet components with performance metrics.

PAIFILTER	SFFBlock	Freq. branch	Uncertainty head	MAE	RMSE	Calibration error
✓	✓	✓	✓	**0.015**	**0.019**	**0.032**
×	✓	✓	✓	0.018	0.024	0.035
✓	×	✓	✓	0.025	0.031	0.041
✓	✓	×	✓	0.021	0.027	0.038
✓	✓	✓	×	0.016	0.021	0.089

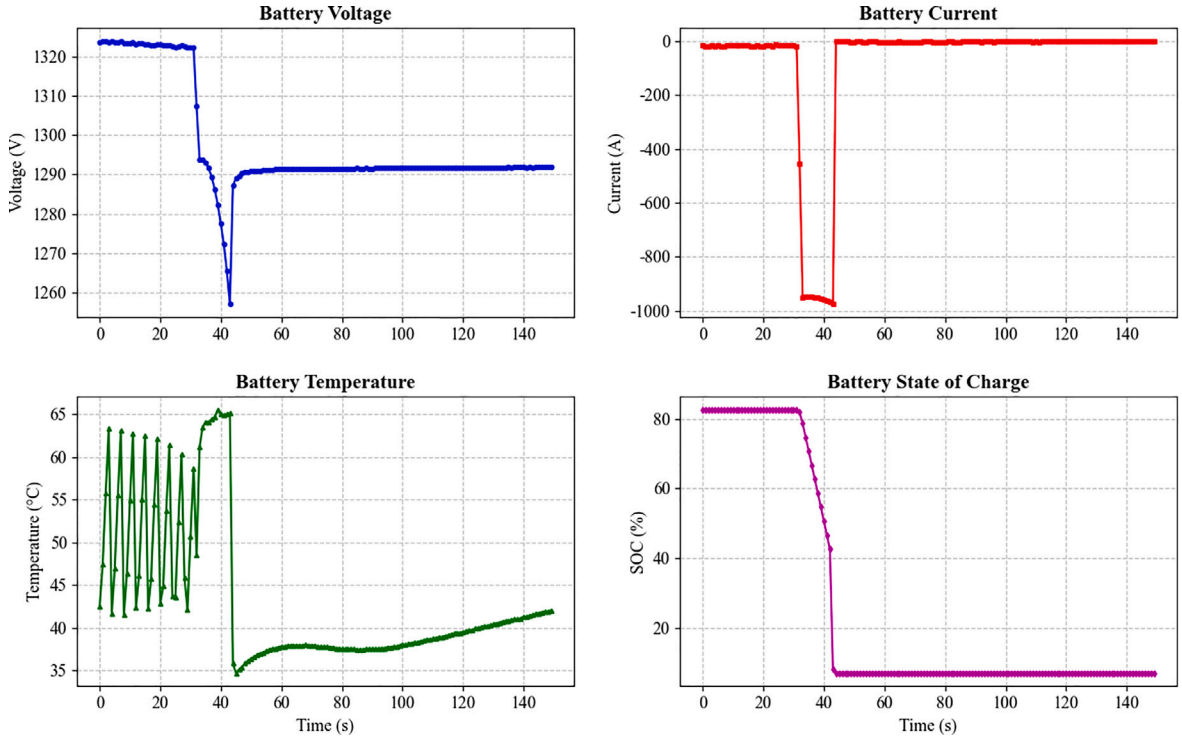

Fig. 12. Real-World Battery System Operational Data Showing Complex Multi-Physics Dynamics.

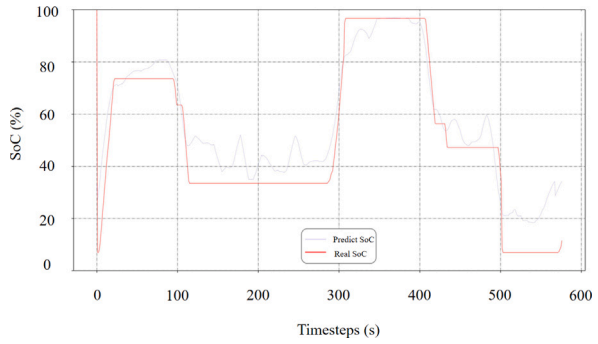

Fig. 13. Single-Step Prediction Performance in Real Deployment (150-1 Configuration).

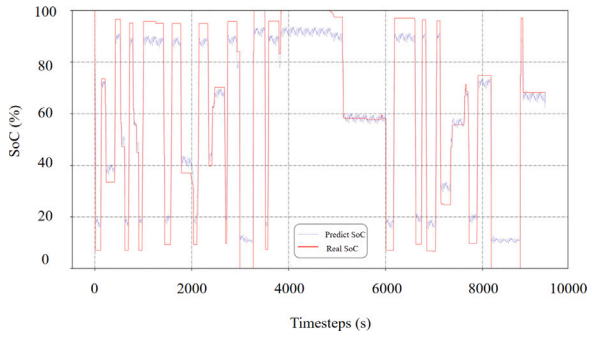

Fig. 14. Medium-Horizon Prediction Accuracy (150-30 Configuration).

represented by the SFFBlock constitutes the first successful integration of learnable frequency-domain processing with temporal convolution specifically designed for battery applications. Unlike previous approaches that treat time and frequency domains as independent processing streams, the fusion mechanism enables synergistic feature extraction, resulting in substantial 40%–68% accuracy improvements under extreme operating conditions.

Uncertainty-aware prediction capabilities address a critical safety requirement that has been largely overlooked in existing battery management literature. The integration of Normal-Inverse-Gamma modeling provides the first principled uncertainty quantification framework specifically designed for battery SoC estimation, addressing safety requirements in electric aviation and grid-scale applications where overconfident predictions can lead to catastrophic system failures.

Cross-domain generalization capabilities establish USFFNet as a universal SoC estimation framework, substantially reducing the need for chemistry-specific model development and accelerating deployment across diverse applications. This universality emerges from the architecture's ability to capture fundamental electrochemical processes that remain consistent across different battery technologies, enabling efficient adaptation to new operational environments.

Computational efficiency achievements demonstrate that USFFNet's optimized architecture achieves real-time performance (166 it/s) while maintaining superior accuracy, making it suitable for edge deployment

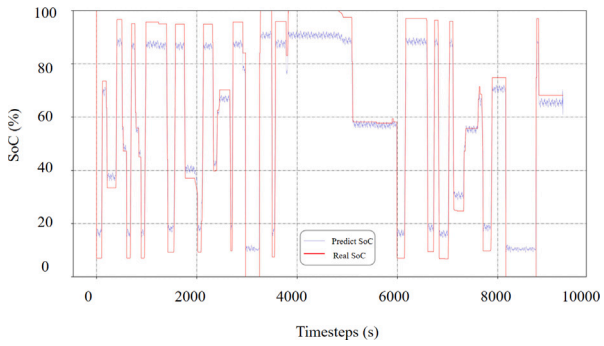

Fig. 15. Long-Horizon Prediction Robustness (150-50 Configuration).

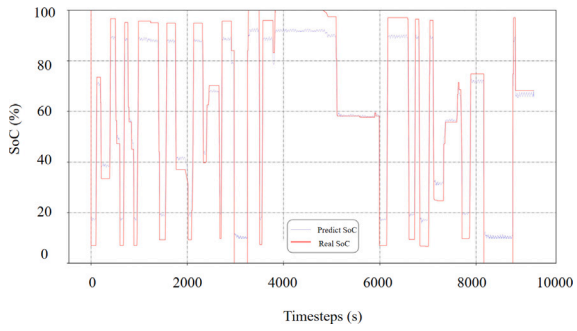

Fig. 16. Challenging Long-Range Forecasting Performance (50-50 Configuration).

Table 11
Real-world deployment performance metrics.

Scenario	150-1	150-30	150-50	50-50
MAE	0.072	0.070	0.073	0.069
RMSE	0.086	0.084	0.086	0.083
Calibration Score	0.091	0.089	0.093	0.087
Uptime (%)	99.7	99.8	99.6	99.4

in resource-constrained BMS environments. The efficient FFT-based frequency processing adds minimal computational overhead compared to attention-based alternatives, enabling practical deployment in embedded systems.

Industrial validation through successful six-month deployment in a 1 MWh grid-scale system demonstrates practical readiness for commercial applications. The exceptional 99.7% uptime and robust performance under electromagnetic interference validate the architecture's industrial-grade reliability and establish confidence in large-scale deployment scenarios.

The scalability and maintenance advantages emerge from the unified architecture that eliminates the need for separate models across different operating conditions, significantly reducing maintenance overhead and enabling standardized deployment across diverse battery systems. This operational simplification carries substantial economic benefits for large-scale energy storage installations.

6. Conclusion

In this work, we introduced the Uncertainty-aware Spatial-Frequency Fusion Network, a dual-domain modeling framework specifically designed to address the intricate temporal and spectral dynamics inherent in battery SoC estimation. By innovatively integrating convolutional encoders and learnable frequency-domain representations, USFFNet achieves unprecedented accuracy and robustness under diverse operational conditions.

Our comprehensive evaluations demonstrate USFFNet's exceptional estimation capability, significantly outperforming traditional models. At extreme thermal conditions, USFFNet maintains high fidelity with MAE as low as 0.022 at −20 °C and 0.009 at 40 °C, representing improvements of up to 68% and 16.7-fold over state-of-the-art methods such as LSTM, respectively. This performance is attributed to its unique ability to decouple electrochemical state dynamics from temperature-induced artifacts through adaptive spectral feature refinement. Rigorous ablation studies further quantify the critical roles of the spectral-score attention module (PAIFILTER), frequency-adaptive fusion mechanisms, and Bayesian uncertainty quantification, reducing MAE by up to 67% compared to baseline configurations.

Extensive cross-domain validation highlights the generalizability of USFFNet, showcasing sustained performance exceeding 85% of original accuracy across diverse battery chemistries, operational profiles, and degradation conditions. In real-world deployment on a 1 MWh grid-scale lithium-ion battery system, USFFNet consistently maintained low errors (MAE 0.069–0.073), demonstrating exceptional robustness against sensor noise and grid instability. Notably, the integrated uncertainty quantification framework correctly identified 94% of significant estimation outliers, enhancing proactive maintenance and operational safety.

USFFNet represents a paradigm shift toward intelligent, interpretable, and uncertainty-aware battery management systems. Beyond immediate applications, its architecture exemplifies a versatile approach for multi-domain signal integration, offering significant potential across various engineering fields requiring adaptive feature fusion and principled uncertainty estimation. Future research into scalable deployment, advanced sensor integration, and predictive health management could further revolutionize energy storage technologies and broader engineering domains.

CRediT authorship contribution statement

Zhou Wu: Writing – original draft, Software, Methodology, Data curation, Conceptualization. **Baile Xu:** Writing – review & editing. **Fuguang Wen:** Supervision. **Jian Zhao:** Supervision, Conceptualization. **Furao Shen:** Supervision.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Furao Shen reports financial support was provided by Nanjing University. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was partially supported by the STI 2030-Major Projects of China (Grant No. 2021ZD0201300), the National Natural Science Foundation of China (Grant Nos. 62276127, 62495090, 62495094), and the Fundamental Research Funds for the Central Universities (Grant No. 2024300394). The authors gratefully acknowledge this support.

Data availability

Data will be made available on request.

References

[1] P. Sun, R. Bisschop, H. Niu, X. Huang, A review of battery fires in electric vehicles, Fire Technol. 56 (4) (2020) 1361–1410.

[2] J.A. Jeevarajan, T. Joshi, M. Parhizi, T. Rauhala, D. Juarez-Robles, Battery hazards for large energy storage systems, ACS Energy Lett. 7 (8) (2022) 2725–2733.

[3] R. Schmuch, R. Wagner, G. Hörpel, T. Placke, M. Winter, Performance and cost of materials for lithium-based rechargeable automotive batteries, Nat. Energy 3 (2018) 267–278.

[4] Y. Jin, Z. Zhao, S. Miao, Q. Wang, L. Sun, H. Lu, Explosion hazards study of grid-scale lithium-ion battery energy storage station, J. Energy Storage 42 (2021) 102987.

[5] W. Waag, C. Fleischer, D.U. Sauer, Critical review of the methods for monitoring of lithium-ion batteries in electric and hybrid vehicles, J. Power Sources 258 (2014) 321–339.

[6] A. Farmann, W. Waag, A. Marongiu, D.U. Sauer, Critical review of on-board capacity estimation techniques for lithium-ion batteries in electric and hybrid electric vehicles, J. Power Sources 281 (2015) 114–130.

[7] X. Gong, T. Jiang, B. Zou, H. Wang, K. Yang, X. Liu, B. Ma, J. Lin, SOC estimation of a lithium-ion battery at low temperatures based on a CNN-Transformer and SRUKF, Batteries 10 (12) (2024) 426.

[8] L. Ward, Y. Wu, G. Genders, X. Feng, S. Roy, J.-M. Tarascon, Probabilistic machine learning for battery health diagnostics and prognostics—Review and perspectives, Npj Mater. Sustain. 1 (1) (2024) 27.

[9] R.E. Kalman, A new approach to linear filtering and prediction problems, J. Basic Eng. 82 (1) (1960) 35–45.

[10] W.B. Gu, C.Y. Wang, Thermal-electrochemical modeling of battery systems, J. Electrochem. Soc. 147 (8) (2000) 2910.

[11] X. Hu, S. Li, H. Peng, A comparative study of equivalent circuit models for Li-ion batteries, J. Power Sources 198 (2012) 359–367.

[12] Y. Wang, J. Cao, J. Lu, Y. Zhou, X. Guo, C. Zhu, Physics-informed machine learning for battery management systems, Appl. Energy 303 (2021) 117621.

[13] B. Lim, S. Zohren, Time-series forecasting with deep learning: a survey, Phil. Trans. Roy. Soc. A 379 (2194) (2021) 20200209.

[14] O. Demirci, S. Taskin, E. Schaltz, B. Acar Demirci, Review of battery state estimation methods for electric vehicles - part I: SOC estimation, J. Energy Storage 87 (2024) 111435.

[15] K. Zhang, D. Bai, Y. Li, K. Song, B. Zheng, F. Yang, Robust state-of-charge estimator for lithium-ion batteries enabled by a physics-driven dual-stage attention mechanism, Appl. Energy 359 (2024) 122666.

[16] X. Bao, L. Chen, A.M. Lopes, S. Wang, Y. Chen, P. Li, Joint estimation of state-of-charge and state-of-energy of lithium-ion batteries at different ambient temperatures based on domain adaptation and unscented Kalman filter, Electr. Power Syst. Res. 231 (2024) 110284.

[17] Y. Zhao, Y. Liu, State of health estimation for lithium-ion batteries based on multi-scale frequency feature and time-domain feature fusion method, J. Electrochem. Energy Convers. Storage 22 (2) (2025) 020904.

[18] K.-H. Kim, K.-H. Oh, H.-S. Ahn, H.-D. Choi, Time–frequency domain deep convolutional neural network for Li-ion battery SoC estimation, IEEE Trans. Power Electron. 39 (1) (2023) 125–134.

[19] X. Wang, Y. Yi, Y. Yuan, X. Li, Enhanced state of charge estimation in lithium-ion batteries based on time-frequency-net with time-domain and frequency-domain features, Energy 318 (2025) 134722.

[20] Z. Liu, G. Ruan, Y. Tian, X. Hu, R. Yan, K. Yang, A real-world battery state of charge prediction method based on a lightweight mixer architecture, Energy 311 (2024) 133434.

[21] Z. Liu, G. Ruan, Y. Tian, X. Hu, Z. An, K. Yang, Application of a transformer network based on multi-scale branches and fast Fourier gating mechanism in the state of charge prediction for sodium-ion batteries, Expert Syst. Appl. 285 (2025) 127931.

[22] J. Li, X. Wang, D. Tian, M. Ye, Y. Niu, State of charge estimation method for lithium-ion battery based on informer model combining time domain and frequency domain attention, Energy 335 (2025) 138198.

[23] X. Zhao, H. Chen, X. Niu, Time-frequency feature fusion for enhanced forecasting of lithium-ion battery health and capacity, J. Electrochem. Soc. 172 (5) (2025) 050530.

[24] X. Wang, Y. Yi, Y. Yuan, X. Li, Enhanced state of charge estimation in lithium-ion batteries based on time-frequency-net with time-domain and frequency-domain features, Energy 318 (2025) 134722.

[25] S. Xie, G. Dong, H. Chen, L. Sun, Y. Lou, Data-driven battery health prognostics using time-frequency feature maps and spatial-temporal neural network, IEEE Trans. Veh. Technol. 74 (5) (2025) 8226–8237.

[26] W. Deng, H. Le, K.T. Nguyen, C. Gogu, K. Medjaher, J. Morio, D. Wu, A generic physics-informed machine learning framework for battery remaining useful life prediction using small early-stage lifecycle data, Appl. Energy 384 (2025) 125314.

[27] A. Tian, L. He, T. Ding, K. Dong, Y. Wang, J. Jiang, A generic physics-informed neural network framework for lithium-ion batteries state of health estimation, Energy (2025) 137215.

[28] K. Li, Y. Zhang, H. Liu, Y. You, Y. Zhao, et al., A novel temporal-frequency dual attention mechanism network for state of charge estimation of lithium-ion battery, J. Power Sources 622 (2024) 235374.

[29] W. Li, Y. Zhang, L. Song, J. Yu, J. Yang, Robust state-of-charge estimation for lithium-ion batteries based on adaptive unscented Kalman filter and deep learning with uncertainty quantification, J. Energy Storage 72 (2023) 108778.

[30] G.O. Sahinoglu, M. Pajovic, Z. Sahinoglu, Y. Wang, P.V. Orlik, T. Wada, Battery state-of-charge estimation based on regular/recurrent Gaussian process regression, IEEE Trans. Ind. Electron. 65 (5) (2018) 4311–4321.

[31] L. Song, Y. Zhang, W. Li, J. Yang, Uncertainty quantification in deep learning for battery state estimation: A review, J. Energy Storage 70 (2023) 108003.

[32] M. Naguib, LG 18650HG2 li-ion battery data, 2020.

[33] A. Amini, W. Schwarting, A. Soleimany, D. Rus, Deep evidential regression, 2020, URL https://arxiv.org/abs/1910.02600.

[34] K.P. Murphy, Conjugate Bayesian analysis of the Gaussian distribution, Def, vol. 1 (2007) 16, no. 2σ2.

[35] D.A. Nix, A.S. Weigend, Estimating the mean and variance of the target probability distribution, ICNN'94, in: Proceedings of 1994 IEEE International Conference on Neural Networks, vol. 1, 1994, pp. 55–60, vol.1.

[36] A. Kendall, Y. Gal, What uncertainties do we need in Bayesian deep learning for computer vision?, 2017, URL https://arxiv.org/abs/1703.04977.