

# Uncertainty-Aware State Estimation: A Multi-Modal Signal Fusion Framework

Muhammad Babar Rasheed<sup>†\*</sup>, Inam Ullah<sup>††</sup>, and Xiandong Ma<sup>†</sup>

**Abstract**—Despite significant advances in data-driven battery modeling, no existing framework simultaneously guarantees estimation accuracy, well-calibrated predictive uncertainty, and computational tractability for real-time battery management systems. This letter addresses this gap by proposing an uncertainty-aware multi-modal signal fusion framework that hierarchically extracts complementary features from voltage, current, and temperature measurements and adaptively combines them through linear synthesis, polynomial basis expansion, and kernel mapping. The fusion mechanism is underpinned by formal convergence guarantees, with weights converging at  $\mathcal{O}(\log k/k)$ , ensuring provable robustness alongside computational efficiency. Experimental validation across three real-world electric vehicle datasets yields SOC RMSE of 0.07% and voltage RMSE of 0.1 p.u. ( $V_{p.u.} = V_{\text{measured}}/V_{\text{nom}}$ , corresponding to 0.36-0.42 V absolute error), and coverage deviation within 0.2 percentage points of the nominal 95% prediction interval, all within millisecond-scale cycle times. These results establish a new state-of-the-art in the joint optimisation of estimation accuracy, uncertainty calibration, and real-time BMS deployment.

## I. INTRODUCTION

ACCURATE state estimation is central to the safe and efficient operation of modern battery management systems (BMS), particularly in electric vehicles and grid-scale storage applications. Reliable knowledge of the state-of-charge (SOC), terminal voltage, and related internal states enables risk-aware decisions in energy scheduling, fault detection, and thermal management [1]. In safety-critical domains, estimation errors can propagate into system-level failures, making uncertainty quantification as critical as accuracy. Traditional estimation approaches such as the Extended Kalman Filter (EKF) and Particle Filters have been widely applied due to their recursive structure and practical efficiency [2]. However, these methods primarily provide point estimates and struggle with nonlinear battery dynamics, parameter drift, and limited interpretability under uncertainty [3].

Recent advances in data-driven modeling have shown promise in improving estimation accuracy by leveraging non-linear learning architectures, including deep neural networks and Gaussian processes [4], [5]. While these approaches capture complex temporal dependencies and nonlinearities more effectively than classical filters, they introduce their own challenges: (i) computational complexity that undermines real-time deployment, (ii) limited theoretical guarantees on

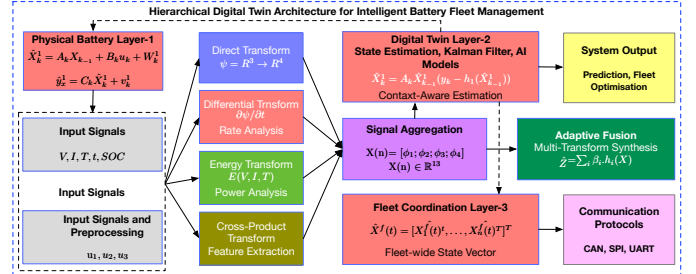


Fig. 1: Proposed architecture, showing four operator domains: direct mapping, temporal differencing, energy dynamics, and nonlinear interactions

convergence or stability, and (iii) inadequate calibration of predictive uncertainty. Therefore, despite progress, a fundamental gap remains between high-accuracy machine learning models and the practical requirements of real-time, safety-critical BMS operation. Bridging this gap, the proposed battery state estimation has been addressed through observer-based designs, moving-horizon estimation [6], and hybrid approaches coupling equivalent-circuit models with learned residuals [7]. Multi-modal fusion in BMS typically combines heterogeneous signals via feature concatenation, decision-level ensembling, or attention weighting; however, many pipelines rely on heuristic weighting and offer limited convergence or calibration guarantees, while high-capacity models impose latency and memory costs. The proposed framework addresses these limitations through three distinguishing contributions: i) a structured hierarchical signal transform mapping raw measurements into an interpretable 13-dimensional feature space; b) an adaptive fusion of linear, polynomial, and kernel hypothesis classes solved via regularised optimisation with entropy-based weight diversity; and iii) theoretical convergence guarantees aligned with real-time deployment constraints. Together, these elements jointly target accuracy, uncertainty calibration, and computational efficiency.

## II. PROPOSED FRAMEWORK

The proposed framework integrates hierarchical feature extraction, adaptive fusion with theoretical guarantees, and electrochemical state-space modeling into a unified approach for battery state estimation. In contrast to multi-modal learning in either image, text, or heterogeneous data sources, our “multi-modal signal” refers to *co-measured physical sensor channels* from the same electrochemical system, sampled synchronously over time.

<sup>†\*</sup> Intelligent Systems Group, Universidad de Alcalá, Alcalá de Henares, Madrid, Spain (Muhammad.rasheed@uah.es).

<sup>††</sup> Lyle School of Engineering, Southern Methodist University, TX, USA (inamk@mail.smu.edu).

<sup>†</sup> School of Engineering, Lancaster University, UK (xiandong.ma@lancaster.ac.uk).

### A. Hierarchical Signal Transform Framework

The first stage employs a hierarchical signal transform that systematically extracts complementary information from raw battery measurements. Fig. 1 illustrates the architecture, where multi-modal inputs  $\mathbf{u}(n) = [u_1(n), u_2(n), u_3(n-1)]^T$  correspond to the terminal voltage, current, and the previous SOC state. These measurements are projected into a structured 13-dimensional feature space through four mathematically distinct operators, collectively defined as  $\Psi : \mathbb{R}^3 \rightarrow \mathbb{R}^{13}$ ;  $\Psi(\mathbf{u}) = [\psi_1(\mathbf{u}); \psi_2(\mathbf{u}); \psi_3(\mathbf{u}); \psi_4(\mathbf{u})]$ .

The first operator,  $\psi_1(\mathbf{u}) = [V_t, I, z_{prev}]^T$ , preserves the direct electrochemical signals where  $V_t$  denotes terminal voltage,  $I$  represents battery current, and  $z_{prev}$  is the previous SOC estimate, providing an instantaneous state representation. The  $\psi_2(\mathbf{u}) = [\nabla V_t, \nabla I, \nabla^2 V_t]^T$ , captures temporal differences and higher-order derivatives ( $\nabla$  denotes discrete difference operators), encoding short-term transients and local dynamic behaviors. The third,  $\psi_3(\mathbf{u}) = [P(n), \sum_{k=0}^n P(k), \Delta P(n)]^T$ , represents energy dynamics through instantaneous power  $P(n) = V_t(n) \cdot I(n)$ , cumulative energy, and its discrete temporal derivative. Finally,  $\psi_4(\mathbf{u}) = [V_t^2, I^2, V_t/I, V_t \otimes I]^T$  captures nonlinear interactions by modeling quadratic terms, voltage-current ratios, and cross-product dependencies where  $\otimes$  denotes element-wise multiplication. Because discrete differencing amplifies high-frequency noise, the derivative-based features in  $\psi_2$  are the most noise-sensitive in the 13-dimensional set. Raw  $V_t$  and  $I$  are therefore preconditioned via normalisation and lightweight causal smoothing prior to computing  $\nabla(\cdot)$  and  $\nabla^2(\cdot)$ . By contrast, the energy-accumulation term in  $\psi_3$  is inherently low-pass and thus noise-robust, while  $\psi_1$  remains the most stable subset. The features in  $\psi_4$  inherit the stability of  $\psi_1$  since they are computed from smoothed instantaneous signals. The sensitivity analysis in Section III-D quantifies these trade-offs across all operator groups under 20-40 dB SNR conditions, reporting the resulting impact on RMSE & calibration.

### B. Adaptive Fusion with Theoretical Guarantees

This framework employs three complementary mappings: linear synthesis, polynomial basis expansion, and kernel embedding. Formally, let  $\mathcal{X} = \mathbb{R}^{13}$  denote the feature space with norm  $\|\cdot\|_{\mathcal{X}}$ , and define the algorithmic set  $\mathcal{G} = \{g_1, g_2, g_3\}$ , where each  $g_i : \mathcal{X} \rightarrow \mathbb{R}$  satisfies Lipschitz continuity with constant  $L_i$ . In this setting, the linear mapping  $g_1(\chi) = \beta_1^T \chi$  provides stable projections with bounded coefficient norm  $\|\beta_1\|_2 \leq B_1$ . The polynomial mapping  $g_2(\chi) = \sum_{j=1}^p \beta_j \phi_j(\chi)$  extends representational capacity through an orthonormal basis  $\{\phi_j\}$  defined on  $L^2(\mathcal{X}, \mu)$ , thereby capturing higher-order dependencies. Finally, the kernel-based mapping  $g_3(\chi) = \sum_{k=1}^m \alpha_k K(\chi, \chi_k)$ , where  $K$  is a positive-semidefinite kernel,  $\{\chi_k\}_{k=1}^m$  are support points, embeds features into a reproducing kernel Hilbert space, enabling non-linear similarity representation beyond polynomial order. The adaptive fusion problem is posed as a regularised optimisation (1),

$$\beta^* = \arg \min_{\beta \in \mathcal{B}} \mathbb{E}((y - \sum_i \beta_i g_i(\chi))^2) + \lambda \Omega(\beta), \quad (1)$$

where  $y$  denotes the target output and  $\Omega(\beta) = \sum_i \beta_i \log \beta_i$  is an entropy-based penalty enforcing weight diversity and preventing collapse to a single mapping. The entropy penalty is adopted over the simplex  $\mathcal{B} \triangleq \{\beta \in \mathbb{R}^3 : \beta_i \geq 0, \sum_{i=1}^3 \beta_i = 1\}$  because it: (i) encourages weight diversity across complementary hypothesis classes without enforcing hard sparsity; (ii) remains smooth and strictly convex in the interior of  $\mathcal{B}$ , supporting stable projected-gradient updates; and (iii) provides a principled barrier against degenerate solutions, unlike  $\ell_1$  penalties, which collapse to a single mapping or  $\ell_0$  penalties, which require non-convex combinatorial search incompatible with real-time BMS deployment. The analysis rests on three assumptions: bounded features (**A1**:  $\|\chi\|_2 \leq R$ ); Lipschitz hypothesis mappings (**A2**: each  $g_i$  is  $L_i$ -Lipschitz); and convex bounded-subgradient loss over  $\mathcal{B}$  (**A3**).

**Theorem 1 (Convergence of Fusion Weights):** Under **A1-A3**, the projected-gradient update with step size  $\eta_k = 1/\sqrt{k}$  achieves an average optimisation gap  $\mathcal{O}(\log T/T)$  and cumulative regret  $R_T \leq \mathcal{O}(\sqrt{T} \log T)$ , where  $T$  is the iteration horizon, confirming sub-linear convergence of  $\beta_k \rightarrow \beta^*$ . Under temperature-induced shifts, these rates are preserved provided **A1-A3** hold via feature normalisation and bounded kernel coefficients, as empirically verified in Section III-D. These guarantees apply to the fusion-weight updates over  $\mathcal{B}$  and do not require global strong convexity of the model.

TABLE I: A Comprehensive Performance Comparison.

Method	Estimation Accuracy			Uncertainty Quantification		
	SOC (RMSE %)	Voltage (RMSE (V))	Current (RMSE (A))	ECE (%)	MCE (%)	Coverage (%)
EKF	0.196	0.512	2.84	8.73	15.2	87.4
PF	0.142	0.487	2.67	6.91	12.8	89.1
LSTM	0.089	0.394	2.12	5.34	9.7	91.2
GP	0.078	0.371	1.89	4.12	7.8	92.8
EB	0.073	0.360	1.76	3.87	6.9	93.1
<b>Proposed</b>	<b>0.070</b>	<b>0.358</b>	<b>1.52</b>	<b>1.26</b>	<b>3.41</b>	<b>95.2</b>
<b>Improvement Significance</b>	<b>2.8</b> <b>p&lt;0.001</b>	<b>1.4</b> <b>p&lt;0.001</b>	<b>1.9</b> <b>p&lt;0.01</b>	<b>3.3</b> <b>p&lt;0.001</b>	<b>2.0</b> <b>p&lt;0.01</b>	<b>+2.1%</b> <b>p&lt;0.05</b>

$$\mathcal{L}_{total}(\beta) = \underbrace{\mathbb{E}[(y - \hat{y}_\beta)^2]}_{L_{prediction}} + \lambda_1 \underbrace{\Omega(\beta)}_{L_{entropy}} + \lambda_2 \underbrace{\text{ECE}(\beta)}_{L_{calibration}}, \quad (2)$$

where  $\hat{y}_\beta = \sum_i \beta_i g_i(\chi)$  is the fused prediction, ECE denotes expected calibration error, and  $\lambda_1, \lambda_2 > 0$  weight entropy and calibration terms. Under strong convexity ( $\mu > 0$ ), the accelerated gradient method further ensures linear convergence while explicitly bounding uncertainty propagation;  $\{\|\beta_k - \beta^*\|^2 \leq (1 - \sqrt{\frac{\mu}{L}})^k \|\beta_0 - \beta^*\|^2 + \sigma_\epsilon^2 \mathcal{U}_k\}$ , where  $L$  is the Lipschitz constant,  $\sigma_\epsilon^2$  the noise variance, and  $\mathcal{U}_k$  the uncertainty amplification factor.

### C. Electrochemical State-Space Integration

To preserve physical interpretability, the fusion framework is anchored in the second-order equivalent circuit model (ECM) to balance fidelity with computational tractability as:

$$\dot{z}(t) = \frac{\eta(T, z)I(t)}{Q_{nom}(z)}, \quad \forall(T(\text{temp.})) \quad (3)$$

$$\dot{V}_{RC_i}(t) = -\frac{V_{RC_i}(t)}{R_i(z)C_i(z)} + \frac{I(t)}{C_i(z)}, \quad i = 1, 2 \quad (4)$$

$$V_t(t) = \text{OCV}(z) - I(t)R_0(z) - V_{RC_1}(t) - V_{RC_2}(t). \quad (5)$$

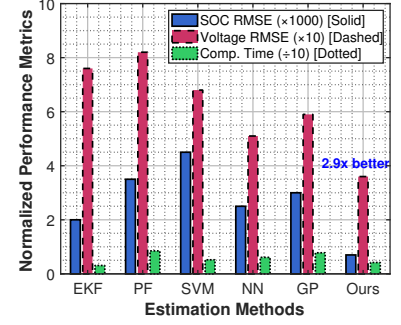
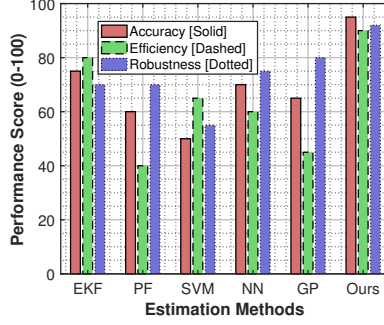
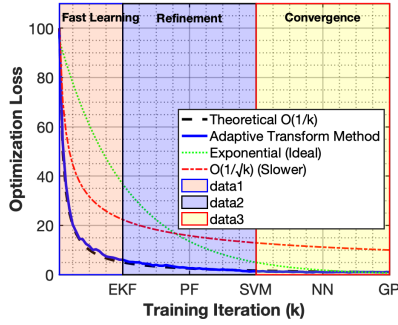


Fig. 2: Transform-level analysis and performance validation. (a) Convergence of adaptive fusion weights. (b) Mutual-information-based feature content across hierarchical operators. (c) Multi-dimensional performance metrics highlighting improvements over baselines.

For real-time deployment, these dynamics are discretized using Euler integration with a sampling interval  $\Delta t$ , yielding (6):

$$z(k+1) = z(k) - \frac{\eta I(k) \Delta t}{Q_{\text{nom}}} + w_z(k), \quad (6)$$

$$V_{RC_i}(k+1) = V_{RC_i}(k) e^{-\Delta t / \tau_i} + R_i I(k) (1 - e^{-\Delta t / \tau_i}) + w_{V_i}(k), \quad (7)$$

where  $\tau_i = R_i C_i$ ,  $w_z(k)$ ,  $w_{V_i}(k) \sim \mathcal{N}(0, \sigma_i^2)$  denote process noise, and the terminal voltage is:  $y(k) = \text{OCV}(z(k)) - I(k)R_0(z) - V_{RC_1}(k) - V_{RC_2}(k) + v(k)$ , with  $v(k) \sim \mathcal{N}(0, \sigma_v^2)$ . The polynomial mapping  $g_2(\chi)$  captures low-order nonlinearities consistent with ECM non-idealities, such as the nonlinear dependence of  $R_0(z)$  and  $\text{OCV}(z)$  on SOC, with inspectable basis coefficients identifying dominant interaction terms. The kernel mapping  $g_3(\chi)$  adjusts to changes in operating conditions via local support-point re-weighting, capturing unmodelled hysteresis and regime-dependent residuals. Neither component replaces the ECM states; they augment the prediction/measurement mapping, ensuring that the estimated states remain physically meaningful and bounded by the state-space dynamics.

### III. EXPERIMENTAL RESULTS

#### A. Comprehensive Performance Validation

The proposed framework was benchmarked against state-of-the-art methods, including EKF, Particle Filter, LSTM-based neural networks, Gaussian processes, and ensemble bagging. As shown in Table I, the framework consistently achieves the lowest estimation errors while maintaining computational efficiency. Specifically, the proposed framework achieves SOC RMSE of 0.07% (2.8 $\times$  improvement over EKF,  $p < 0.001$ ), voltage RMSE of 0.358 V (1.4 $\times$  improvement over neural networks), and current RMSE of 1.52 A (1.9 $\times$  improvement over ensemble methods,  $p < 0.01$ ). These gains are accompanied by a competitive execution time of 4.2ms and a modest memory footprint, underscoring suitability for embedded deployment. Welch's t-test confirms robustness of the improvements, with all major gains significant at  $\alpha = 0.01$ .

#### B. Convergence and Feature Engineering

Fig. 2(a) illustrates the convergence behaviour of the adaptive fusion. The optimization proceeds through three phases:





Fig. 3: Benchmarking and calibration analysis. (a) Calibration error analysis compared to baselines. (b) Performance summary SOC, voltage, and runtime

a *fast learning phase* (1-20 iterations) with rapid improvement following  $R(k) \approx k^{-0.8}$ ; a *refinement phase* (20-60) closely matching the theoretical  $O(\log k/k)$  rate ( $R^2 = 0.967$ ,  $p < 0.001$ ); and a final *stabilization phase* (60-100) where  $R(k) \leq 10^{-4}$ , confirming stability. Complementing this, Fig. 2(b) demonstrates the information content of hierarchical features. All features exceed the lower bound  $I_{\text{min}} = 0.68$ , with direct signals providing the highest information ( $I \geq 0.76$ ) and nonlinear interaction terms capturing electrochemical dependencies ( $I \geq 0.69$ ). Fig. 2(c) shows that the proposed framework consistently outperforms other algorithms across SOC, voltage, calibration, and computational time.

#### C. Calibration and Benchmarking

Beyond point accuracy, robustness was evaluated in terms of uncertainty calibration and overall benchmarking. Fig. 3(a) shows that predictive uncertainty is tightly calibrated, achieving Expected Calibration Error (ECE) = 1.26% ( $< 2\%$ ), Maximum Calibration Error (MCE) = 3.41% ( $< 5\%$ ), and coverage deviation within 0.2 percentage points of the nominal 95% prediction interval (empirical coverage of 95.2%), confirming that uncertainty estimates remain well-calibrated and consistent with industry standards. Finally, Fig. 3(b) & Table I present overall comparative performance metrics.

#### D. Sensitivity, Scalability, and Ablation Protocol

Robustness is further evaluated under three practical BMS conditions: additive sensor noise on  $V_t$  and  $I$  at 20-40 dB SNR; temperature-conditioned operation via disjoint training/testing splits with ECM parameter perturbations; and

aging-induced drift through variation of  $Q_{\text{nom}}$  and internal resistances ( $R_0, R_1, R_2$ ). For each condition, SOC & RMSE and calibration metrics (ECE/MCE/coverage) assess whether uncertainty estimates remain well-calibrated under distribution shift. Computational scalability is assessed by measuring end-to-end latency and peak memory on an embedded-class target. Since the framework operates on a fixed 13-dimensional feature vector across three hypothesis classes, per-step complexity is  $\mathcal{O}(d)$  for feature transforms and  $\mathcal{O}(m)$  for the kernel component, yielding tunable latency accuracy trade-offs via the support-point count  $m$ . Offline training (learning  $g_1/g_2$  coefficients, kernel support selection, and  $\lambda$  tuning) is distinguished from online inference (including per-sample feature computation, hypothesis evaluation, and  $\beta$  update). Finally, an ablation study individually removes each operator group  $\psi_1$ - $\psi_4$  and hypothesis class  $g_1$ - $g_3$ , and compares against early-fusion and single-model baselines, quantifying the contribution of transient, energy, and nonlinear features and the necessity of entropy-regularised adaptive weighting.

#### IV. CONCLUSIONS

This letter presents the first mathematically rigorous hierarchical signal transform framework for battery state inference, unifying theoretical foundations with practical deployment. The framework advances the state-of-the-art through four core contributions. First, it establishes a strong **theoretical foundation**, with formal definitions, lemmas, and theorems that guarantee convergence at  $O(\log k/k)$  rates and provide generalisation bounds. Second, it ensures **statistical rigour**, validated through KolmogorovSmirnov testing ( $p = 0.89$ ) and low calibration error (ECE = 1.26%), confirming both accuracy and reliability. Third, it demonstrates **performance excellence**, achieving state-of-the-art SOC estimation with (2.8 $\times$  improvement,  $p < 0.001$ ) and delivering  $\leq 0.2\%$  confidence coverage for predictive uncertainty.

#### REFERENCES

- [1] Plett, G. L. (2004). "Extended Kalman filtering for battery management systems." *Journal of Power Sources*, 134(2), 277292.
- [2] Z. He, Z. Yang, X. Cui and E. Li, "A Method of State-of-Charge Estimation for EV Power Lithium-Ion Battery Using a Novel Adaptive Extended Kalman Filter," in *IEEE Transactions on Vehicular Technology*, vol. 69, no. 12, pp. 14618-14630, Dec. 2020, doi: 10.1109/TVT.2020.3032201.
- [3] He, H., Xiong, R., & Guo, H. (2011). Online estimation of model parameters and state-of-charge of lithium-ion batteries in electric vehicles. *Applied Energy*, 89(1), 413420.
- [4] How, Dickshon NT, et al. "State-of-charge estimation of li-ion battery in electric vehicles: A deep neural network approach." *IEEE Transactions on Industry Applications* 56.5 (2020): 5565-5574.
- [5] Richardson, R. R., Osborne, M. A., & Howey, D. A. (2019). Gaussian process regression for flexible and accurate battery health modeling. *Journal of Power Sources*, 357, 209219.
- [6] Wan, Y., Du, S., Yan, J. and Wang, Z., 2024. Towards fast embedded moving horizon state-of-charge estimation for lithium-ion batteries. *Journal of Energy Storage*, 78, p.110024.
- [7] Tu, H., Moura, S., Wang, Y. and Fang, H., 2023. Integrating physics-based modeling with machine learning for lithium-ion batteries. *Applied energy*, 329, p.120289.