Les Cahiers du GERAD ISSN: 0711–2440

Distributionally robust extended Kalman filtering for lithium-ion battery state-of-charge estimation

D. Fisher, K. Sheshyekani, A. Lesage-Landry

G–2026–34

July 2026

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2026
– Bibliothèque et Archives Canada, 2026

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2026
– Library and Archives Canada, 2026

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Distributionally robust extended Kalman filtering for lithium-ion battery state-of-charge estimation

Donald Fisher [a, b, c]

Keyhan Sheshyekani [b]

Antoine Lesage-Landry [a, b, c]

[a] *GERAD, Montréal (Qc), Canada, H3T 1J4*

[b] *Polytechnique Montréal, Montréal (Qc), Canada, H3T 0A3*

[c] *MILA, Montréal (Qc), Canada, H2S 3H1*

donald.fisher@etud.polymtl.ca
keyhan.sheshyekani@polymtl.ca
antoine.lesage=landry@polymtl.ca

July 2026
Les Cahiers du GERAD
G–2026–34

Abstract : Lithium-ion batteries (LIBs) are currently the core energy storage component of electric vehicles (EVs). Real-time, accurate, and robust state-of-charge (SOC) estimation is critical for reliable range prediction and for preventing cell degradation due to overcharging or deep discharge. Current model-based estimation methods, based on equivalent circuit models and filters, such as the extended Kalman filter (EKF) and unscented Kalman filter (UKF), provide effective real-time SOC estimation. However, they typically assume a Gaussian process and measurement noise with known mean and covariance, which may not hold in practice. Distributionally robust Kalman filtering (DRKF) utilizes the Wasserstein distance to construct an ambiguity set around a nominal process and measurement noise, thereby accounting for errors in the nominal assumptions. In this paper, we propose the distributionally robust extended Kalman filter (DREKF), which broadens this methodology to nonlinear problems and applies it to the battery state estimation problem. We benchmark the DREKF against the EKF and UKF using simulations based on an open-source urban dynamometer driving schedule dataset. Results across multiple temperatures demonstrate that the DREKF achieves accuracy and computation time comparable to conventional EKF and UKF approaches under nominal conditions, while offering improved resilience to measurement outliers.

Keywords: Electric vehicles; state-of-charge estimation; lithium-ion battery; equivalent circuit model; extended Kalman filter; distributionally robust optimization

Résumé : Les batteries lithium-ion sont présentement le choix principal pour le stockage d'énergie dans les véhicules électriques (*electric vehicle*, EV). Il est donc essentiel d'avoir une estimation précise et fiable de l'état de charge (*state of charge*, SOC) en temps réel afin de prévoir l'autonomie de la batterie d'une manière fiable et d'éviter la dégradation due à une surcharge ou à une décharge profonde. Les méthodes d'estimation actuelles basées sur des modèles de circuits équivalents et des filtres, telles que le filtre de Kalman étendu (*extended Kalman filter*, EKF) et le filtre de Kalman sans parfum (*unscented Kalman filter*, UKF), fournissent une estimation efficace du SOC en temps réel. Cependant, elles supposent généralement un bruit de processus et un bruit de mesure gaussien, avec une moyenne et une covariance connues, ce qui peut ne pas être le cas dans la pratique. Le filtre de Kalman robuste en distribution utilise la distance de Wasserstein pour construire un ensemble d'ambiguïté autour d'un bruit de processus et de mesure nominal, se prémunissant ainsi contre toutes les déviations dans les valeurs nominales. Dans ce travail, nous proposons un nouveau filtre robuste en distribution de Kalman étendu (*distributionally robust extended Kalman filter*, DREKF), qui étend cette méthodologie au scénario nonlinéaire et on l'applique au problème de l'estimation de l'état de la batterie. Nous comparons le DREKF contre l'EKF et l'UKF utilisant des simulations basées sur un ensemble de données ouvertes provenant d'un programme de conduite sur dynamomètre urbain. Les résultats obtenus pour un éventail de températures démontrent que le DREKF atteint une précision comparable à celle des approches EKF et UKF conventionnelles dans des conditions nominales, tout en offrant une meilleure résilience aux valeurs aberrantes des mesures.

Mots clés : Véhicules électriques; estimation de l'état de charge; batterie de lithium-ion; modélisation du circuit électrique équivalent; filtre de Kalman étendu; optimisation distributionnellement robuste

Acknowledgements: This work was supported in part by Serafin Solutions and Mitacs through the Mitacs Accelerate Program under grant IT33680.

1 Introduction

Lithium-ion batteries (LIBs) are widely used in currently manufactured electric vehicles (EVs), and it is projected that the global EV stock will expand to 140 million by 2030 [6]. LIBs are the most popular choice for energy storage in EVs due to their high energy density, long lifespan, light weight, and low self-discharge rate relative to other battery chemistries [13]. To ensure the battery's lifespan [35] and alleviate range anxiety [2], battery management systems must provide accurate and robust real-time state-of-charge (SOC) estimation.

Accurate SOC estimation remains challenging due to the inherent nonlinearity of battery dynamics and the highly variable operating environments of EVs [21]. Rapid acceleration, regenerative braking, and switching operations of power electronic devices in EVs result in substantial transient current variations [4]. Acquisition systems are also susceptible to electromagnetic interference, leading to spikes in the voltage and current noise, which typically exhibit non-Gaussian characteristics [4, 30]. When faced with non-Gaussian noise or significant initial estimation bias, the performance of estimators relying on Gaussian assumptions degrades, leading to reduced accuracy and convergence issues [33,39]. This motivates the use of a distributionally robust estimation framework.

Related Works: Battery SOC estimation can be broadly categorized into four methods: direct measurement methods, electrochemical models (EMs), data-driven methods, and equivalent circuit models (ECMs).

Direct measurement methods include Coulomb counting and the open-circuit voltage (OCV) characterization method. Coulomb counting is capable of real-time use and is computationally light. However, it suffers from cumulative errors due to sensor inaccuracies, integration error and initial SOC estimation uncertainty, leading to significant deviations over time [3, 25]. OCV characterization alone is unsuitable for real-time estimation due to the required rest periods, which are not feasible for EVs [37].

EMs utilize internal electrochemical kinetics and Li-ion transport behaviour to predict the current SOC. The main EMs for LIBs are the pseudo-2-d model [15] and the single particle model [24,41]. These methods are based on partial differential equations, leading to computational bottlenecks in online applications. As such, EMs are typically only preferred in battery chemistries where the relationship between the SOC and OCV cannot be utilized, e.g., lithium-sulphur [12].

Data-driven approaches like deep learning have also been utilized for LIB SOC estimation [16]. Given sufficient training data and an appropriate architecture, e.g., recurrent neural networks with long short-term memory [38], the SOC value can be predicted accurately without the need for a dedicated model. However, the blackbox nature of data-driven methods means that these models lack the interpretability and reliability guarantees needed for such a critical tool.

ECMs utilize circuit models with resistors and capacitors to mimic the polarization present in LIBs that cause the load voltage to differ from the OCV [43]. Kalman filters (KFs) can then be employed using the state transition and measurement functions derived from the ECM [3], and are widely used due to their relatively low computational complexity and satisfactory tracking performance over time [5, 20, 42]. Although the standard KF is effective in mitigating cumulative errors and correcting state estimation deviations in real-time, its recursive computations are inherently derived from linear system assumptions. Given the nonlinear behaviour of LIBs, the conventional KF is unable to accurately track SOC dynamics. This has led to the use of the extended Kalman filter (EKF) [5, 30] and unscented Kalman filter (UKF) [29]. To mitigate the filter's high sensitivity to ECM parameters, adaptive KF techniques have also been developed. For instance, [19] introduces an adaptive UKF that updates the system and measurement noise covariances in real-time, yielding improved performance over standard EKF and UKF implementations. Additionally, [40] highlights the advantages of incorporating variable parameters within an EKF to achieve a reduction in maximum estimation error, particularly at high C-rates. The aforementioned filters assume that the measurement and process noise follow a Gaussian distribution and that their characteristics are known during the estimation process. This, in turn, limits the applicability and effectiveness of these filters in practical deployments where measurement and process noise characteristics have uncertainty and may be non-Gaussian.

In this work, rather than relying on exact knowledge of the process and measurement noise distributions, we instead seek state estimators that perform well under the worst-case distributions within a pre-specified ambiguity set centred around a nominal model. Within the literature, [31] establishes a distributionally robust Kalman filter (DRKF) in the linear setting. Subsequently, [8] has developed formulations for both finite and infinite-horizon problems, while [14] has proposed a computationally efficient DRKF by deriving theoretical conditions under which the Kalman gain may be computed offline. This paper introduces a distributionally robust SOC estimator for lithium-ion batteries with nonlinear dynamics under uncertain and potentially non-Gaussian noise. The main contributions are as follows:

1. We design a distributionally robust extended Kalman filter (DREKF) that integrates order-2 Wasserstein ambiguity sets into an EKF-based estimation framework, enabling improved robustness to noise distribution mismatch and measurement outliers in nonlinear settings.
2. We propose a computationally efficient DREKF for real-time LIB SOC estimation to hedge against non-Gaussian noise, in which worst-case covariance updates are selectively recomputed based on voltage residuals, preserving real-time feasibility despite the inclusion of semidefinite programs (SDPs).
3. We illustrate the performance of the DREKF on an open-source LIB dataset across multiple temperatures, initialization conditions, and corrupted measurement scenarios, demonstrating improved robustness relative to the EKF and UKF.

The remainder of the paper is organized as follows. In Section 2, we introduce the preliminary background for our method. Section 3 presents the DREKF algorithm. Section 4 presents the empirical evaluation based on the open-source LG 18650HG2 Li-ion Battery dataset [17]. Finally, we lay out our conclusions in Section 5.

2 Preliminaries

We now present the relevant background on battery modelling and KFs.

2.1 Equivalent circuit modelling

ECMs enable the development of a fully defined state-space and measurement function through the use of Kirchhoff's voltage law (KVL) and Kirchhoff's current law (KCL). ECMs consist of the following components to model typical LIB behaviour:

- A voltage source representing the OCV of the battery. We denote the OCV as a polynomial function of the battery SOC [3, 29], $u_{\mathrm{oc}}(s) \in \mathbb{R}$, where s represents the SOC.
- A resistor $r_0 > 0$, which captures the instantaneous voltage drop that occurs when a current is applied.
- A set of N resistor-capacitor (RC) pairs, parameterized by r_j and c_j, $j = 1, 2, \ldots, N$, which models the voltage relaxation and polarization effects [3, 29].

The number of RC pairs N is chosen to balance between improved accuracy and additional computational burden. We set $N = 2$ (2-RC) because it yields a good level of accuracy in terms of voltage response, while only subject to limited computational complexity [3, 29]. The resulting circuit is shown in Figure 1. We let $i(t) \in \mathbb{R}$ denote the applied current at time t, adopting the convention that discharge current is positive, and let $u_{\mathrm{L}}(t) \in \mathbb{R}$ denote the load voltage at time t.

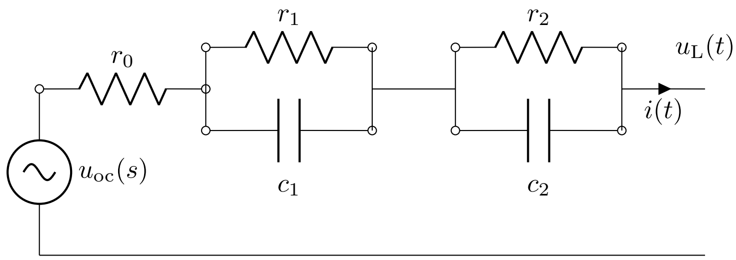

Figure 1: 2-RC ECM.

Applying KVL and KCL, $u_{\mathrm{L}}(t)$ is given by:

$$u_{\mathrm{L}}(t) = u_{\mathrm{oc}}(s) - r_0 i(t) - u_1(t) - u_2(t), \tag{1}$$

where $u_1(t)$ and $u_2(t)$ represent the transient voltage drops across the first and second RC pairs, respectively. The polynomial function for $u_{\mathrm{oc}}(s)$, as well as the values for the resistors and capacitors, are typically fitted with data from hybrid pulse power characterization (HPPC) tests, conducted across a range of temperatures [3]. This is briefly discussed later in Section 2.2. The dynamics of u_1 and u_2 are governed by the following ordinary differential equations (ODEs):

$$\dot{u}_1 = \frac{1}{c_1}\left(i(t) - \frac{u_1}{r_1}\right) \tag{2}$$

$$\dot{u}_2 = \frac{1}{c_2}\left(i(t) - \frac{u_2}{r_2}\right), \tag{3}$$

where $\dot{u}_1$ and $\dot{u}_2$ denote the time derivatives of the respective RC pair voltages. To track the evolution of the SOC, we employ Coulomb counting [25]. Letting s_t denote the SOC at time t in seconds and assuming the prior SOC s_{t-1} is known, the updated SOC is given by:

$$s_t = s_{t-1} - \frac{1}{3600Q}\int_{t-1}^{t} i(\eta)\mathrm{d}\eta, \tag{4}$$

where $Q > 0$ is the battery capacity in Ah, and the negative sign reflects the convention that positive current corresponds to discharge. To obtain a discretized state-space model that captures both the SOC evolution and transient voltage dynamics, we solve ODEs (2)–(3) analytically using the integrating factor

$$\alpha_j = \mathrm{e}^{-\frac{\Delta t}{\tau_j}}, \quad j = 1, 2, \dots, N, \tag{5}$$

where Δt is the sampling interval and $\tau_j = r_j c_j$ is the time constant. We index the discretized time horizon by $k \in \mathbb{N}$, corresponding to a sampling period, Δt. A process noise vector $\mathbf{w}_k \in \mathbb{R}^{3\times 1}$ is included to account for current sensor noise, modelling inaccuracies, and unmodelled hysteresis effects. Combining (4) and the discretized solutions to (2)–(3) yields the following discrete-time state dynamics:

$$\begin{bmatrix} s_k \\ u_{1,k} \\ u_{2,k} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \alpha_1 & 0 \\ 0 & 0 & \alpha_2 \end{bmatrix} \begin{bmatrix} s_{k-1} \\ u_{1,k-1} \\ u_{2,k-1} \end{bmatrix} + \begin{bmatrix} -\frac{\Delta t}{3600Q} \\ r_1(1-\alpha_1) \\ r_2(1-\alpha_2) \end{bmatrix} i_k + \mathbf{w}_k, \tag{6}$$

where α_1 and α_2 are defined in (5). To obtain a discretized measurement equation, a scalar measurement noise term $v_k \in \mathbb{R}$ is added to (1), which captures voltage sensor noise and residual modelling error. Hereinafter, $u_{\mathrm{L},k}$ denotes the measured terminal voltage, while $\hat{u}_{\mathrm{L},k}$ denotes the model-based estimate of the load voltage at time step k. The corresponding measurement equation is:

$$\hat{u}_{\mathrm{L},k} = u_{\mathrm{oc}}(s_k) - r_0 i_k - u_{1,k} - u_{2,k} + v_k. \tag{7}$$

We re-express (6) and (7) compactly in the standard discrete-time stochastic state-space form and obtain

$$\begin{aligned} \mathbf{x}_k &= \mathbf{A}\mathbf{x}_{k-1} + \mathbf{B}\mathbf{u}_k + \mathbf{w}_k \\ \mathbf{y}_k &= \mathbf{C}\mathbf{x}_k + \mathbf{D}\mathbf{u}_k + \mathbf{v}_k, \end{aligned} \tag{8}$$

where $\mathbf{x}_k \in \mathbb{R}^{n_x}$ denotes the system state vector, $u_k \in \mathbb{R}^{n_u}$ denotes our control input $i(t)$, $\mathbf{y}_k \in \mathbb{R}^{n_y}$ denotes the measured output, and $\mathbf{v}_k$ represents a generalized measurement noise vector at time step k. The matrices $\mathbf{A} \in \mathbb{R}^{n_x \times n_x}$ and $\mathbf{B} \in \mathbb{R}^{n_x \times n_u}$ are directly obtained from (6), while $\mathbf{C} \in \mathbb{R}^{n_y \times n_x}$ and $\mathbf{D} \in \mathbb{R}^{n_y \times n_u}$ follow from the measurement Equation (7). For the 2-RC ECM, we have dimensions $n_x = 3$, $n_u = 1$, and $n_y = 1$. Note that the measurement equation is nonlinear due to the polynomial nature of $u_{\text{oc}}(s_k)$. The state vector is written component-wise as

$$\mathbf{x}_k = \begin{bmatrix} x_{1,k} & x_{2,k} & \dots & x_{n_x,k} \end{bmatrix}^\top,$$

where $x_{\ell,k}$, $\ell = 1, 2, \dots N$, denotes the ℓ-th component at time step k.

2.2 Estimating battery model parameters

The reliability and accuracy of the ECM in estimating s_k and $u_{\text{L},k}$ depend on the precision with which its parameters are identified, because each parameter directly shapes the voltage response used by the estimator. The numerical values of these parameters determine both the magnitude and dynamics of voltage transients during current changes. For example, a higher r_0 increases the instantaneous voltage drop under load, while a larger τ_j slows the relaxation behaviour, increasing the time taken for the load voltage to settle to the OCV value following a current pulse. If these values are inaccurate, the model produces an incorrect predicted terminal voltage $\hat{u}_{\text{L},k}$, and the resulting error propagates into the SOC estimation. The ohmic resistance r_0, polarization resistances r_j, and polarization capacitances c_j all vary with SOC s and temperature T. Accordingly, each parameter is identified as a function of these variables, e.g., $r_0(s,T)$, $r_j(s,T)$, and $c_j(s,T)$. We collect all model parameters into the vector $\boldsymbol{\psi} = \begin{bmatrix} r_0, \ r_1, \ c_1, \ r_2, \ c_2, \ u_{\text{oc}}(s) \end{bmatrix}^\top \in \mathbb{R}^6$, where the dependence on s and T is left implicit for notational brevity. Parameter identification is performed by minimizing the sum of the squared error between $\hat{u}_{\text{L},k}$ and $u_{\text{L},k}$ over a test dataset of K samples. Our parameter estimation problem is given as

$$\hat{\boldsymbol{\psi}} \in \arg\min_{\boldsymbol{\psi}} \sum_{k=1}^{K} \left(u_{\text{L},k} - \hat{u}_{\text{L},k}(\boldsymbol{\psi})\right)^2, \tag{9}$$

which can be solved offline for experimental purposes [23] or in real-time as a recursive least squares problem [22, 28, 32, 34]. The HPPC test is commonly used to generate the data required for this identification procedure, as it excites the battery with controlled current pulses at fixed SOC setpoints [3], enabling clear observation of both the OCV and the transient voltage dynamics associated with each RC pair. The test is carried out over several temperatures, and the resulting identified parameter functions $r_0(s,T)$, $r_j(s,T)$, and $c_j(s,T)$ are subsequently used to populate and dynamically update the system matrices $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$, and $\mathbf{D}$ in (8).

2.3 Extended Kalman filtering

For battery SOC estimation, an EKF is required over the standard KF, as the latter is not well-suited to the nonlinear dynamics of LIBs [36]. The EKF operates similarly to the KF with a state prediction and state correction step. However, the nonlinear dynamics are linearized at operating points of the state estimates. In the prediction step, the EKF forecasts the state and its uncertainty based on the nonlinear state transition model:

$$\hat{\mathbf{x}}_{k|k-1} = f(\hat{\mathbf{x}}_{k-1|k-1}, \mathbf{u}_k) + \mathbf{w}_k \tag{10}$$

$$\mathbf{P}_{k|k-1} = \mathbf{F}_k \mathbf{P}_{k-1|k-1} \mathbf{F}_k^\top + \mathbf{Q}_k, \tag{11}$$

where $\hat{\mathbf{x}}_{k|k-1}$ is the prior state estimate at time k given measurements until $k-1$, $\mathbf{P}_{k|k-1} \in \mathbb{R}^{n_x \times n_x}$ is the corresponding prior error covariance matrix, and $\mathbf{Q}_k \in \mathbb{R}^{n_x \times n_x}$ represents the process noise covariance at time k. The nonlinear state transition is described by the mapping $f(\hat{\mathbf{x}}_{k-1|k-1}, \mathbf{u}_k)$ with $\mathbf{F}_k \in \mathbb{R}^{n_x \times n_x}$ denoting its Jacobian matrix evaluated at the posterior state estimate $\hat{\mathbf{x}}_{k-1|k-1}$. In the correction step, the EKF updates the Kalman gain $\mathbf{K}_k \in \mathbb{R}^{n_x \times n_y}$ by using the measurement function $\mathbf{H}_k \in \mathbb{R}^{n_y \times n_x}$, which is the linearized matrix of the nonlinear measurement function $h(\hat{\mathbf{x}}_{k|k-1})$. The state estimate is then updated using the residual error, defined as the difference between the observed measurement $\mathbf{y}_k$ and the predicted measurement, $\hat{\mathbf{y}}_k = h(\hat{\mathbf{x}}_{k|k-1}) + \mathbf{v}_k$. This leads to:

$$\begin{aligned}
\mathbf{K}_k =& \mathbf{P}_{k|k-1}\mathbf{H}_k^\top \left(\mathbf{H}_k \mathbf{P}_{k|k-1}\mathbf{H}_k^\top + \mathbf{R}_k\right)^{-1} \\
\hat{\mathbf{x}}_{k|k} =& \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k \left(\mathbf{y}_k - \hat{\mathbf{y}}_k\right) \\
\mathbf{P}_{k|k} =& \left(\mathbf{I} - \mathbf{K}_k \mathbf{H}_k\right) \mathbf{P}_{k|k-1},
\end{aligned}$$

where $\hat{\mathbf{x}}_{k|k}$ and $\mathbf{P}_{k|k}$ are the posterior state estimate and error covariance matrix, respectively, and $\mathbf{R}_k \in \mathbb{R}^{n_y \times n_y}$ is the measurement noise covariance matrix. To obtain $\mathbf{H}_k$ in the context of LIB SOC estimation, we differentiate (7) with respect to the prior state estimate, which yields:

$$\mathbf{H}_k = \begin{bmatrix} u'_{\mathrm{oc}}(s_k) & -1 & -1 \end{bmatrix}, \tag{12}$$

where $u'_{\mathrm{oc}}(s_k)$ is the derivative of $u_{\mathrm{oc}}(s_k)$ with respect to s_k evaluated at time k. Additionally, we define the measurement $\mathbf{y}_k \coloneqq u_{\mathrm{L},k}$ and $\hat{\mathbf{y}}_k \coloneqq \hat{u}_{\mathrm{L},k}$. Note that this linearization is an approximation and causes the EKF to no longer be an optimal estimator. Although the unscented transform in the UKF often gives better results for nonlinear estimation problems [19], the use of the EKF allows us to extend the DRKF to nonlinear problems, as presented next.

3 Proposed methodology

We now formulate our DREKF.

3.1 Notation

Let $\mathcal{S}^n$ denote the set of all symmetric matrices in $\mathbb{R}^{n\times n}$. We use $\mathcal{S}^n_+ \subset \mathcal{S}^n$ and $\mathcal{S}^n_{++} \subset \mathcal{S}^n$ to represent the cone of symmetric positive semidefinite and positive definite matrices in $\mathcal{S}^n$, respectively. For any $\mathbf{A}, \mathbf{B} \in \mathcal{S}^n$, $\mathbf{A} \succeq \mathbf{B}$ and $\mathbf{A} \succ \mathbf{B}$ mean that $\mathbf{A} - \mathbf{B} \in \mathcal{S}^n_+$ and $\mathbf{A} - \mathbf{B} \in \mathcal{S}^n_{++}$, respectively. The set of Borel probability measures with support $\Omega \subseteq \mathbb{R}^n$ is denoted by $\mathcal{P}(\Omega)$, while the subset of measures with finite second-order moments is denoted by $\mathcal{P}_2(\Omega)$.

3.2 Problem setup

Consider the state-space problem defined similarly to (8). We also assume that the nominal distributions for the initial state $\hat{\mathbf{x}}_0$, process noise $\hat{\mathbf{w}}_k$, and measurement noise $\hat{\mathbf{v}}_k$ for all k are Gaussian with distributions $\hat{\mathbb{Q}}_{\mathbf{x}_0} = \mathrm{N}(\hat{\mathbf{x}}_0, \hat{\boldsymbol{\Sigma}}^-_{\mathbf{x}_0})$, $\hat{\mathbb{Q}}_{\mathbf{w}_k} = \mathrm{N}(\hat{\mathbf{w}}_k, \hat{\boldsymbol{\Sigma}}_{\mathbf{w}_k})$, and $\hat{\mathbb{Q}}_{\mathbf{v}_k} = \mathrm{N}(\hat{\mathbf{v}}_k, \hat{\boldsymbol{\Sigma}}_{\mathbf{v}_k})$, respectively, where $\hat{\mathbf{x}}_0, \hat{\mathbf{w}}_k \in \mathbb{R}^{n_x}$ and $\hat{\mathbf{v}}_k \in \mathbb{R}^{n_y}$ are the mean vectors, while $\hat{\boldsymbol{\Sigma}}^-_{\mathbf{x}_0}$, $\hat{\boldsymbol{\Sigma}}_{\mathbf{w}_k} \in \mathcal{S}^{n_x}_+$ and $\hat{\boldsymbol{\Sigma}}_{\mathbf{v}_k} \in \mathcal{S}^{n_y}_+$ are the nominal covariance matrices. In practice, these nominal distributions are simply estimations as the true distributions governing $\mathbf{x}_0$, $\mathbf{w}_k$, and $\mathbf{v}_k$ are never directly observable. This motivates the use of an estimator that performs well if the true distributions deviate from the nominal assumptions.

Because the state estimate and measurement are coupled through the observation model, a worst-case distribution is sought over their joint distribution rather than over each marginal independently. At each time step k, the prediction step of the filter propagates both the state estimation and process noise uncertainty through the state transition dynamics, such that the process noise $\mathbf{w}_k$ is absorbed into a prior state estimate distribution, $\mathbb{P}_{\hat{\mathbf{x}}_{k|k-1}} \in \mathcal{P}_2(\mathbb{R}^{n_x})$ [14, 31]. The measurement noise distribution, defined as $\mathbb{P}_{\mathbf{v}_k} \in \mathcal{P}_2(\mathbb{R}^{n_y})$, remains independent [14, 31]. We write the joint distribution as

$\mathbb{P}_k = \mathbb{P}_{\hat{\mathbf{x}}_{k|k-1}} \otimes \mathbb{P}_{\mathbf{v}_k}$, where $\otimes$ denotes the product measure. Similarly, we define the nominal joint distribution as $\hat{\mathbb{Q}}_k = \hat{\mathbb{Q}}_{\hat{\mathbf{x}}_{k|k-1}} \otimes \hat{\mathbb{Q}}_{\mathbf{v}_k}$. To account for the uncertainty between $\mathbb{P}_k$ and $\hat{\mathbb{Q}}_k$, we consider an ambiguity set $\mathcal{B}_{\boldsymbol{\theta}}(\hat{\mathbb{Q}}_k) \subset \mathcal{P}_2(\mathbb{R}^{n_x}) \otimes \mathcal{P}_2(\mathbb{R}^{n_y})$ consisting of all joint probability distributions that are sufficiently close to the nominal distribution, $\hat{\mathbb{Q}}_k$. We control this distance through the use of the order-2 Wasserstein distance with respect to the Euclidean norm $||\cdot||_2$ [7], denoted as $\mathrm{W}_2(\mathbb{P}_1, \mathbb{P}_2)$ for two probability measures $\mathbb{P}_1$ and $\mathbb{P}_2$. Formally, $\mathcal{B}_{\boldsymbol{\theta}}(\hat{\mathbb{Q}}_k)$ is defined as:

$$
\begin{aligned}
\mathcal{B}_{\boldsymbol{\theta}}(\hat{\mathbb{Q}}_k) \coloneqq \big\{ & \mathbb{P}_k = \mathbb{P}_{\hat{\mathbf{x}}_{k|k-1}} \otimes \mathbb{P}_{\mathbf{v}_k} \\
& \mathbb{P}_{\hat{\mathbf{x}}_{k|k-1}} \in \mathcal{P}_2(\mathbb{R}^{n_x}),\ \mathbb{P}_{\mathbf{v}_k} \in \mathcal{P}_2(\mathbb{R}^{n_y}) \\
& \mathrm{W}_2(\mathbb{P}_{\hat{\mathbf{x}}_{k|k-1}}, \hat{\mathbb{Q}}_{\mathbf{x}_{k|k-1}}) \leq \theta_{\mathbf{x}},\ \mathrm{W}_2(\mathbb{P}_{\mathbf{v}_k}, \hat{\mathbb{Q}}_{\mathbf{v}_k}) \leq \theta_{\mathbf{v}} \big\},
\end{aligned}
$$

where $\boldsymbol{\theta} = (\theta_{\mathbf{x}}, \theta_{\mathbf{v}})^\top \geq \mathbf{0}$ are the radii of the Wasserstein ball for the process and measurement noise, respectively. With the ambiguity set defined, we seek a state estimator that minimizes the worst-case mean-square error across all distributions with respect to the ambiguity set. The distributionally robust minimum mean-square (DR-MMSE) takes the form:

$$
\min_{\boldsymbol{\phi}_k \in \mathcal{F}_k} \max_{\mathbb{P}_k \in \mathcal{B}_{\boldsymbol{\theta}}(\hat{\mathbb{Q}}_k)} \mathbb{E}^{\mathbb{P}_k} \left[\|\mathbf{x}_k - \boldsymbol{\phi}_k(\mathbf{y}_k)\|^2 \right], \tag{13}
$$

where $\boldsymbol{\phi}_k : \mathbb{R}^{n_y} \to \mathbb{R}^{n_x}$ denotes the state estimator, and $\mathcal{F}_k$ denotes the set of all admissible estimators. The inner maximization selects the least-favourable distribution within the ambiguity set $\mathcal{B}_{\boldsymbol{\theta}}(\hat{\mathbb{Q}}_k)$, and the radii $\theta_{\mathbf{x}}$ and $\theta_{\mathbf{v}}$ act as hyperparameters that govern the size of the ambiguity set. As $\theta_{\mathbf{x}}, \theta_{\mathbf{v}} \to 0$, the ambiguity set collapses to the nominal distribution and the problem recovers the standard MMSE solution, while larger radii yield a more robust but potentially overly conservative estimator. [31].

While (13) is generally nonconvex, [7, 14] have shown that if the nominal distributions $\hat{\mathbb{Q}}_{\mathbf{x}_0}, \hat{\mathbb{Q}}_{\mathbf{w}_k}$, and $\hat{\mathbb{Q}}_{\mathbf{v}_k}$ are assumed to be Gaussian, then (13) has a tractable reformulation. The resulting estimator preserves the recursive structure of the KF, but replaces the nominal covariances for the process and measurement noise with worst–case covariances, in line with the robustness consideration.

We now present the DRKF [14, 26] and then devise the DREKF in Section 3.3. The DRKF begins with a prediction step: At time k and given the previous posterior estimate $(\hat{\mathbf{x}}_{k-1|k-1}, \hat{\boldsymbol{\Sigma}}_{\mathbf{x}_{k-1|k-1}})$, the prior state estimate, $\hat{\mathbf{x}}_{k|k-1}$, and the pseudo-nominal prior state covariance, $\hat{\boldsymbol{\Sigma}}_{\mathbf{x}_{k|k-1}}$ are calculated as:

$$
\hat{\mathbf{x}}_{k|k-1} = \mathbf{A}\hat{\mathbf{x}}_{k-1|k-1} + \mathbf{B}\mathbf{u}_k + \mathbf{w}_k \tag{14}
$$

$$
\hat{\boldsymbol{\Sigma}}_{\mathbf{x}_{k|k-1}} = \mathbf{A}\hat{\boldsymbol{\Sigma}}_{\mathbf{x}_{k-1|k-1}}\mathbf{A}^\top + \hat{\boldsymbol{\Sigma}}_{\mathbf{w}_k}. \tag{15}
$$

Then, instead of directly using the pseudo-nominal prior state covariance $\hat{\boldsymbol{\Sigma}}_{\mathbf{x}_{k|k-1}}$ and measurement covariance $\hat{\boldsymbol{\Sigma}}_{\mathbf{v}_k}$, the DR formulation computes the least–favourable covariances by solving the following SDP [14]:

$$
\max_{\boldsymbol{\Sigma}_{\mathbf{x}_{k|k-1}}, \boldsymbol{\Sigma}_{\mathbf{x}_{k|k}}, \boldsymbol{\Sigma}_{\mathbf{v}_k}, \mathbf{Y}, \mathbf{Z}} \mathrm{Tr}[\boldsymbol{\Sigma}_{\mathbf{x}_{k|k}}] \tag{16a}
$$

$$
\text{s.t.} \begin{bmatrix} \boldsymbol{\Sigma}_{\mathbf{x}_{k|k-1}} - \boldsymbol{\Sigma}_{\mathbf{x}_{k|k}} & \boldsymbol{\Sigma}_{\mathbf{x}_{k|k-1}}\mathbf{C}_k^\top \\ \mathbf{C}_k\boldsymbol{\Sigma}_{\mathbf{x}_{k|k-1}} & \mathbf{C}_k\boldsymbol{\Sigma}_{\mathbf{x}_{k|k-1}}\mathbf{C}_k^\top + \boldsymbol{\Sigma}_{\mathbf{v}_k} \end{bmatrix} \succeq 0 \tag{16b}
$$

$$
\begin{bmatrix} \hat{\boldsymbol{\Sigma}}_{\mathbf{x}_{k|k-1}} & \mathbf{Y} \\ \mathbf{Y}^\top & \boldsymbol{\Sigma}_{\mathbf{x}_{k|k-1}} \end{bmatrix} \succeq 0, \quad \begin{bmatrix} \hat{\boldsymbol{\Sigma}}_{\mathbf{v}_k} & \mathbf{Z} \\ \mathbf{Z}^\top & \boldsymbol{\Sigma}_{\mathbf{v}_k} \end{bmatrix} \succeq 0 \tag{16c}
$$

$$
\mathrm{Tr}[\boldsymbol{\Sigma}_{\mathbf{x}_{k|k-1}} + \hat{\boldsymbol{\Sigma}}_{\mathbf{x}_{k|k-1}} - 2\mathbf{Y}] \leq \theta_{\mathbf{x}}^2 \tag{16d}
$$

$$
\mathrm{Tr}[\boldsymbol{\Sigma}_{\mathbf{v}_k} + \hat{\boldsymbol{\Sigma}}_{\mathbf{v}_k} - 2\mathbf{Z}] \leq \theta_{\mathbf{v}}^2, \tag{16e}
$$

where $\boldsymbol{\Sigma}_{\mathbf{x}_{k|k}}, \boldsymbol{\Sigma}_{\mathbf{x}_{k|k-1}}, \hat{\boldsymbol{\Sigma}}_{\mathbf{x}_{k|k-1}} \in \mathcal{S}_+^{n_x}$, $\boldsymbol{\Sigma}_{\mathbf{v}_k} \in \mathcal{S}_{++}^{n_y}$, $\mathbf{Y} \in \mathbb{R}^{n_x \times n_x}$, $\mathbf{Z} \in \mathbb{R}^{n_y \times n_y}$. The decision variables $\boldsymbol{\Sigma}_{\mathbf{x}_{k|k-1}}$ and $\boldsymbol{\Sigma}_{\mathbf{x}_{k|k}}$ represent the prior and posterior state covariances, respectively, while $\boldsymbol{\Sigma}_{\mathbf{v}_k}$ represents the measurement covariance. The objective (16a) finds the pair of $\mathbf{\Sigma}_{\mathbf{x}_{k|k-1}}$ and $\mathbf{\Sigma}_{\mathbf{v}_k}$ which maximizes $\mathbf{\Sigma}_{\mathbf{x}_{k|k}}$ within the ambiguity set. Constraint (16b) ensures that $\mathbf{\Sigma}_{\mathbf{x}_{k|k}}$ respects the Kalman filter mechanism. The auxiliary variables $\mathbf{Y}$ and $\mathbf{Z}$ and the resulting constraints (16c)−(16e) follow from the use of Schur complement arguments and Gelbrich bounds to convert the original infinite-dimensional optimization problem into a convex finite-dimensional problem [9], and ensure that $\mathbf{\Sigma}_{\mathbf{x}_{k|k-1}}$ and $\mathbf{\Sigma}_{\mathbf{v}_k}$ lie within the specified Wasserstein radii of the ambiguity set $\mathcal{B}_{\boldsymbol{\theta}}(\hat{\mathbb{Q}}_k)$. Interested readers are referred to [9] for the full derivation of (16). Solving (16) yields the least-favourable covariances ($\mathbf{\Sigma}^*_{\mathbf{x}_{k|k-1}}$ and $\mathbf{\Sigma}^*_{\mathbf{v}_k}$), which are used to determine the robust Kalman gain as

$$\mathbf{K}^*_k = \mathbf{\Sigma}^*_{\mathbf{x}_{k|k-1}} \mathbf{C}_k^\top \left(\mathbf{C}_k \mathbf{\Sigma}^*_{\mathbf{x}_{k|k-1}} \mathbf{C}_k^\top + \mathbf{\Sigma}^*_{\mathbf{v}_k}\right)^{-1}. \tag{17}$$

The posterior state estimate and covariance are determined via

$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}^*_k\big(\mathbf{y}_k - \hat{\mathbf{y}}_k\big), \tag{18}$$

$$\hat{\mathbf{\Sigma}}_{\mathbf{x}_{k|k}} = (\mathbf{I} - \mathbf{K}^*_k \mathbf{C}_k)\mathbf{\Sigma}^*_{\mathbf{x}_{k|k-1}}. \tag{19}$$

where $\mathbf{I} \in \mathbb{R}^{n_x \times n_x}$ is the identity matrix. The Wasserstein radii $\theta_{\mathbf{x}}$ and $\theta_{\mathbf{v}}$ determine the estimator's conservativeness and serve as tunable hyperparameters selected from the data. Notably, for quadratic loss functions under a W_2-ambiguity set centred at a Gaussian distribution, the worst-case distribution is known to remain Gaussian. This follows from the fact that the optimal adversarial distribution is an affine mapping of the nominal one, as demonstrated in [18].

3.3 Distributionally robust extended Kalman filter

The DRKF framework of Section 3.2 assumes a linear output mapping $\mathbf{C}_k$. To extend DRKF to the nonlinear setting, we adopt a methodology akin to EKF, which we refer to as DREKF. We begin by assuming that $\mathbf{x}_k$ and $\mathbf{y}_k$ are governed by the nonlinear state and measurement functions $f(\hat{\mathbf{x}}_{k-1|k-1}, \mathbf{u}_k)$ and $h(\hat{\mathbf{x}}_{k|k-1})$, respectively. The prior state estimate, $\hat{\mathbf{x}}_{k|k-1}$, is computed exactly as is in (10). The nonlinear state transition matrix is obtained via the Jacobian $\mathbf{F}_k = \frac{\partial f}{\partial \hat{\mathbf{x}}}\big|_{\hat{\mathbf{x}}_{k-1|k-1}}$. The pseudo-nominal prior state covariance $\hat{\mathbf{\Sigma}}_{\mathbf{x}_{k|k-1}}$ is obtained by using $\mathbf{F}_k$ to approximate the linear dynamics. The prediction step becomes:

$$\hat{\mathbf{x}}_{k|k-1} = f(\hat{\mathbf{x}}_{k-1|k-1}, \mathbf{u}_k) + \mathbf{w}_k \tag{20}$$

$$\hat{\mathbf{\Sigma}}_{\mathbf{x}_{k|k-1}} = \mathbf{F}_k \mathbf{\Sigma}_{\mathbf{x}_{k-1|k-1}} \mathbf{F}_k^\top + \hat{\mathbf{\Sigma}}_{\mathbf{w}_k}. \tag{21}$$

Similarly, the linearized time-varying measurement matrix is obtained via the Jacobian $\mathbf{H}_k = \frac{\partial h}{\partial \hat{\mathbf{x}}}\big|_{\hat{\mathbf{x}}_{k|k-1}}$. Matrix $\mathbf{H}_k$ replaces matrix $\mathbf{C}_k$ within the SDP constraints (16b), the robust gain calculation (17), and the posterior state covariance update (19). Because $\mathbf{H}_k$ is linear, it preserves the Kalman filter's update mechanism for propagating the uncertainty in constraint (16b), as well as when updating the Kalman gain and the posterior state covariance [26]. The updated robust Kalman gain is then computed as:

$$\mathbf{K}^*_k = \mathbf{\Sigma}^*_{\mathbf{x}_{k|k-1}} \mathbf{H}_k^\top \left(\mathbf{H}_k \mathbf{\Sigma}^*_{\mathbf{x}_{k|k-1}} \mathbf{H}_k^\top + \mathbf{\Sigma}^*_{\mathbf{v}_k}\right)^{-1}, \tag{22}$$

and the state and covariance updates are now computed using the linearized approximation:

$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}^*_k\big(\mathbf{y}_k - \hat{\mathbf{y}}_k\big) \tag{23}$$

$$\hat{\mathbf{\Sigma}}_{\mathbf{x}_{k|k}} = (\mathbf{I} - \mathbf{K}^*_k \mathbf{H}_k)\mathbf{\Sigma}^*_{\mathbf{x}_{k|k-1}}. \tag{24}$$

In the context of LIB SOC estimation, nonlinearity is driven solely by the OCV polynomial $u_{\mathrm{oc}}(s)$. Because the state-transition function is linear, $\mathbf{F}_k = \mathbf{A}\ \forall\ k$ while $\mathbf{H}_k$ is defined by (12).

The DREKF is presented in Algorithm 1. We begin by predicting the state prior estimate, $\hat{\mathbf{x}}_{k|k-1}$, and the pseudo-nominal prior state covariance matrix, $\hat{\mathbf{\Sigma}}_{\mathbf{x}_{k|k-1}}$ via (20) and (21) (Line 7). Matrix $\mathbf{C}_k$

Algorithm 1 Distributionally Robust Extended Kalman Filter (DREKF)

Inputs:
1: Wasserstein radii $\theta_{\mathbf{x}}, \theta_{\mathbf{v}}$
2: Covariances $\hat{\mathbf{\Sigma}}_{\mathbf{w}}, \hat{\mathbf{\Sigma}}_{\mathbf{v}}, \hat{\mathbf{\Sigma}}_{\mathbf{x}_0}$,
3: Initial state estimate and noise means $\hat{\mathbf{x}}_0, \hat{\mathbf{w}}_k, \hat{\mathbf{v}}_k$,
4: Update Tolerance ϵ
Output: State estimates $\{\hat{\mathbf{x}}_k\}_{k=1}^{T}$
5: **for** $k = 1$ to T **do**
6: Compute Jacobian $\mathbf{F}_k = \frac{\partial f}{\partial \hat{\mathbf{x}}}\Big|_{\hat{\mathbf{x}}_{k-1|k-1}}$
7: Predict $\hat{\mathbf{x}}_{k|k-1}$ and $\hat{\mathbf{\Sigma}}_{\mathbf{x}_{k|k-1}}$ via (20) and (21)
8: Obtain measurement $\mathbf{y}_k$ and predict $\hat{\mathbf{y}}_k$ via (7)
9: Compute $\mathbf{H}_k = \frac{\partial h}{\partial \hat{\mathbf{x}}}\Big|_{\hat{\mathbf{x}}_{k|k-1}}$
10: **if** $\|\mathbf{y}_k - \hat{\mathbf{y}}_k\| \geq \epsilon$ or $k = 1$ **then**
11: Substitute $\mathbf{H}_k$ into (16b)
12: Solve SDP (16) to obtain $\mathbf{\Sigma}^*_{\mathbf{x}_{k|k-1}}, \mathbf{\Sigma}^*_{\mathbf{v}_k}$
13: **else**
14: $\mathbf{\Sigma}^*_{\mathbf{x}_{k|k-1}}, \mathbf{\Sigma}^*_{\mathbf{v}_k} = \mathbf{\Sigma}^*_{\mathbf{x}_{k-1|k-2}}, \mathbf{\Sigma}^*_{\mathbf{v}_{k-1}}$
15: **end if**
16: Compute $\mathbf{K}^*_k$ via (22)
17: Compute $\hat{\mathbf{x}}_{k|k}, \hat{\mathbf{\Sigma}}_{\mathbf{x}_{k|k}}$ via (23) and (24)
18: **end for**

is approximated by $\mathbf{H}_k$ (Line 9) and we substitute it into (16b) to solve the SDP (16) (Line 11– 12). The robust Kalman gain is then computed via (22) and the state estimate and covariance are updated via (23) and (24) (Line 17).

A drawback of the DREKF is that it requires $\mathbf{H}_k$ to be updated at each time step k. This in turn necessitates recomputing the worst-case covariance at each time step, and prevents the use of a more efficient version of (16), where the Kalman gain is computed offline [14]. To improve computational efficiency, we draw inspiration from works based on event-triggered state estimation [10, 11] and assume that if the predicted measurement is very close to the actual measurement, the worst-case distribution from the previous timestep can be considered a reasonable approximation, and we do not require a recomputation of (16). Therefore, we recompute the worst-case covariances only when the difference between the measured terminal voltage and the estimated value exceeds a certain tolerance, $\epsilon \geq 0$ (Lines 10 and 13). We highlight this modification in grey in Algorithm 1. This additional hyperparameter can be tuned, depending on the desired accuracy and computational efficiency gain of the user and its effect is studied in Section 4.3. Setting $\epsilon = 0$, effectively ensures that (16) is recomputed at every timestep, leading to the nominal DREKF.

4 Numerical study

We now evaluate the performance of the DREKF algorithm for LIB SOC estimation. Numerical simulations are carried out in MATLAB and Simulink. We use the open-source LG 18650HG2 Li-ion Battery dataset [17], which contains tests at $-20°$C, $-10°$C, $0°$C, $10°$C, $25°$C, and $40°$C. Specifically, the HPPC tests are used to characterize the battery behaviour, and the urban dynamometer driving schedule (UDDS) tests are used to evaluate the model's performance against baselines, viz., EKF and UKF. The UDDS dataset gives us a sample time $\Delta t = 0.1$s.

We use the MATLAB `Simscape` battery library and the `tfest` function method to obtain the ECM parameter vector $\hat{\boldsymbol{\psi}}$ in (9) from the HPPC test. We assume that the battery SOH is constant at 100%. The inclusion of a varying SOH is a topic for future work. The reference SOC is obtained via Coulomb counting using high-precision current measurements and a known initial SOC, $x_{1,0} = 1.0$. To evaluate the correction ability of the various Kalman filters, we also observe the performance when $x_{1,0}$ does not match the true value and later set the initial condition to $x_{1,0} = 0.75$.

For each filter, the process and measurement covariances are tuned iteratively until a stable value is obtained. The radii $\theta_{\mathbf{x}}$ and $\theta_{\mathbf{v}}$ are determined through the use of `Optuna` [1], where 100 trials are carried out to find the best performing values with respect to the root mean square error (RMSE). Through this, $\boldsymbol{\theta} = (0.01, 0.001)^\top$ is selected for our DREKF. Our implementation is readily available on Github.

4.1 SOC estimation accuracy

Table 1 reports averages for the SOC estimation error for the nominal UDDS data across temperature and initial SOC conditions. We consider the mean absolute error (MAE) and the RMSE as primary metrics for assessing the accuracy, favouring the RMSE as it gives higher weights to larger deviations from the reference value. We also look at the average 75$^{\text{th}}$ percentile and the number of outlier errors within our dataset to evaluate the spread of the estimation errors. Outliers are defined per test as samples exceeding 1.5× the inter-quartile range of the absolute error. These results are presented in Table 1.

Table 1: Average performance metrics across $-20°$C to $40°$C for nominal UDDS dataset at $\epsilon = 0.01$.

$x_{1,0}$	Method	MAE (%)	RMSE (%)	75$^{\text{th}}$ pct.(%)	No. Outliers
	EKF	3.14	4.46	3.85	3,347.67
1.00	UKF	2.85	3.23	3.55	3,933.00
	DREKF	**2.77**	**3.10**	**3.46**	**2,937.33**
	EKF	3.53	4.42	4.33	4,270.33
0.75	UKF	**2.89**	**3.30**	**3.58**	4134.67
	DREKF	3.03	3.69	**3.58**	**4015.33**

When averaging the metrics across all temperatures, the proposed DREKF consistently outperforms EKF across all metrics. Under the correct initial SOC ($x_{1,0} = 1.0$), it also marginally outperforms UKF. With $x_{1,0} = 0.75$, UKF performs slightly better, partly due to its faster settling time, which is illustrated in Figure 2. Crucially, these results demonstrate that DREKF incurs no significant performance penalty from its robustness under nominal conditions, while also preserving a settling time comparable to the baseline filters when SOC correction is required.

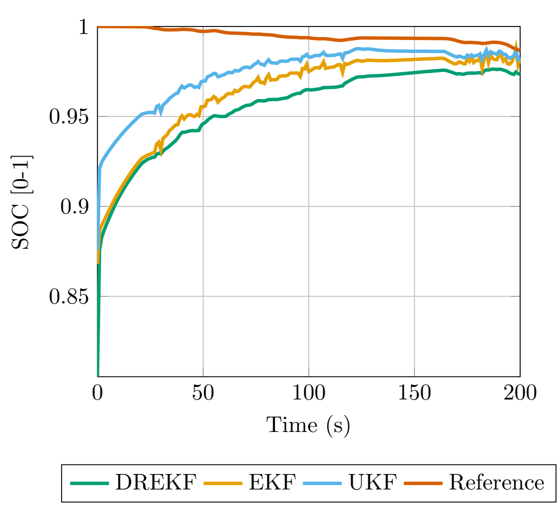

Figure 2: Convergence speed for the tested estimators at $25°$C.

4.2 Robustness against corrupted data

To evaluate the robustness of our approach, we employ a Sliced-Wasserstein (SW) outlier generation method [27] on the UDDS dataset. This method allows us to generate out-of-distribution outliers that can mimic non-Gaussian noise by thresholding the SW-distance between the sample and the outlier distributions. In our corrupted UDDS dataset, we specify 1% of the current measurements to be corrupted by a SW-distance threshold of 0.01, while 1% of the voltage points are corrupted by a SW-distance 1×10^{-3}. The results are presented in Table 2.

Table 2: Performance comparison for the 1% corrupted UDDS dataset at $x_{1,0} = 1.0$ and $\epsilon = 0.01$.

Temp	Method	MAE (%)	RMSE (%)	75th pct. (%)	No. Outliers
$-20°$C	EKF	8.64	9.93	10.40	3397
	UKF	7.42	7.89	8.73	4781
	DREKF	**7.14**	**7.51**	**8.36**	**3322**
$-10°$C	EKF	5.47	6.87	5.18	14435
	UKF	**4.04**	4.46	**4.27**	11041
	DREKF	**4.04**	**4.38**	4.40	**8279**
0°C	EKF	4.80	8.24	3.49	23676
	UKF	1.89	2.23	2.70	256
	DREKF	**1.76**	**2.09**	**2.56**	**83**
10°C	EKF	2.04	3.31	2.48	4268
	UKF	1.91	2.22	2.68	1734
	DREKF	**1.87**	**2.17**	**2.63**	**1602**
25°C	EKF	1.12	1.46	1.59	4131
	UKF	0.95	1.28	1.43	3622
	DREKF	**0.88**	**1.16**	**1.29**	**2673**
40°C	EKF	1.14	1.50	1.75	1534
	UKF	0.90	1.28	1.49	1969
	DREKF	**0.86**	**1.20**	**1.43**	**1462**

We observe in Table 2 that the DREKF achieves lower RMSE, MAE, and 75^{th} percentile errors than both the EKF and UKF for all temperatures, except $-10°$C, where the UKF yields a lower MAE and 75^{th} percentile error. Even in this case, the DREKF outperforms the EKF and produces fewer outliers than the UKF. Across all temperatures, the primary benefit of the proposed DREKF is a reduction in the upper tail of the estimation error distribution, as evidenced by consistently lower 75^{th} percentile errors and fewer outliers. This reflects the fact that distributional robustness limits the sensitivity of the filter's gain to outliers, due to the consideration of the worst-case distribution. We further illustrate this in Figure 3b, where the DREKF exhibits narrower whiskers and fewer outliers relative to EKF and UKF.

We note a clear performance degradation at $-20°$C across all estimators. Exceptionally, some outliers at this temperature correspond to anomalously low estimation errors, as shown in Figure 3a. At this extreme cold, the battery's internal and polarization resistances increase sharply and nonlinearly in the battery, which the parameter lookup tables fail to capture as accurately compared to the other temperatures. As such, while distributional robustness effectively mitigates the effect of non-Gaussian noise, it cannot compensate for larger, systematic model inaccuracies. Nevertheless, we observe that the DREKF is effective in reducing the maximum estimation error by approximately 5% relative to the UKF and 25% relative to the EKF.

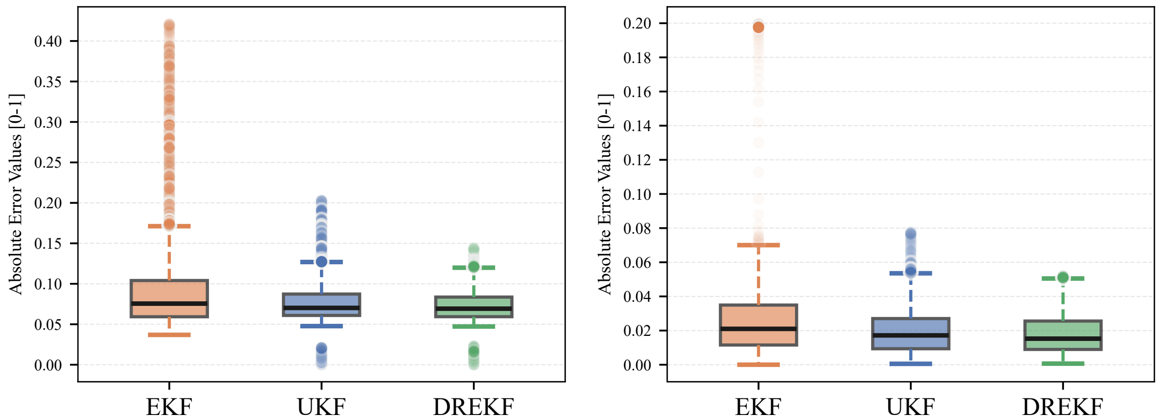

(a) SOC error distribution at -20°C **(b) SOC error distribution at 0°C**

Figure 3: SOC error distribution of estimators for $x_{1,0} = 1.0$.

4.3 Computational complexity

To satisfy real-time operation requirements, each state estimate must be computed before the next measurement is acquired. In this experiment, this corresponds to a maximum allowable computation time of 0.1 s per sample. We evaluate the impact of the ϵ tolerance value by carrying out simulations with varying values. We choose values within the range of $0.001 - 0.02$ V, as this roughly translates to a SOC difference of $0.1 - 2\%$ based on the battery's corresponding OCV-SOC profile. We compute the average computation times across all temperature conditions, using the corrupted UDDS dataset. All trials are conducted on a 64-bit Lenovo Thinkpad with 16GB RAM and a 13th Gen Intel Core i7 processor. `Simulink's Profiler` is used to time each filter during a simulation. For the DREKF, we present the runtimes alongside the average accuracy metrics in Table 3.

Table 3: Evaluation time of the DREKF for different ϵ.

ϵ	Average Computation Time (ms/sample)	MAE (%)	RMSE (%)	75^{th} pct. error (%)
0.001	23.6	2.93	3.38	3.70
0.005	5.15	2.82	3.16	3.56
0.01	2.51	2.76	3.44	3.44
0.02	1.15	2.75	3.43	3.43

Compared to the other filters, the EKF exhibits the lowest mean execution time at 0.634 ms per sample. The UKF incurs a moderate increase in overhead, averaging 2.39 ms per sample. For the DREKF, the cost is significantly higher at the tightest tolerance of $\epsilon = 0.001$, where the SDP is recomputed at effectively every timestep. However, Table 3 shows that the computation cost remains comparable to the UKF at tolerances of $\epsilon = 0.005$ and 0.01, while being more efficient at $\epsilon = 0.02$. These results demonstrate that the implementation of ϵ is effective in reducing the computational demands without sacrificing the robustness of the DREKF, and the algorithm can be implemented with a marginal overhead relative to the benchmark filters. Critically, all variants operated well within the 0.1 s sampling interval, validating the real-time feasibility of the DREKF.

5 Conclusion

This paper proposes a distributionally robust extended Kalman filter (DREKF) for Lithium-ion battery state-of-charge estimation. By incorporating a Wasserstein-based ambiguity set into the estimation framework, the proposed method improves robustness to deviations in the nominal measurement and process noise distributions, while preserving computational efficiency for real-time operations. Simulation results across multiple temperatures and operating conditions demonstrate that the DREKF achieves accuracy comparable to the conventional extended Kalman filter and unscented Kalman filter (UKF) approaches under nominal conditions, while outperforming them when faced with measurement outliers. Future work includes developing a distributionally robust extension to the UKF to further improve estimation performance.

References

[1] Takuya Akiba, View Profile, Shotaro Sano, View Profile, Toshihiko Yanase, View Profile, Takeru Ohta, View Profile, Masanori Koyama, and View Profile. Optuna: A Next-Generation Hyperparameter Optimization Framework. In The 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, ACM Conferences, pages 2623–2631. July 2019.

[2] Fayez Alanazi. Electric Vehicles: Benefits, Challenges, and Potential Solutions for Widespread Adaptation. Applied Sciences, 13(10), May 2023.

[3] Balakumar Balasingam. Robust Battery Management System Design with MATLAB. Artech House Publishers, Norwood, MA, USA, June 2023.

[4] Qi Chen, Feng Zhang, Liang Su, Baoxing Lin, Sien Chen, and Yong Zhang. State Parameter Fusion Estimation for Intelligent Vehicles Based on IMM-MCCKF. Applied Sciences, 14(11), May 2024.

[5] Zhenjie Cui, Weihao Hu, Guozhou Zhang, Zhenyuan Zhang, and Zhe Chen. An Extended Kalman Filter Based SOC estimation Method for Li-ion Battery. Energy Reports, 8:81–87, August 2022.

[6] Ana Luiza Carvalho Ferrer and Antonio Márcio Tavares Thomé. Carbon Emissions in Transportation: A Synthesis Framework. Sustainability, 15(11):8475, May 2023.

[7] Matthias Gelbrich. On a Formula for the L2 Wasserstein Metric between Measures on Euclidean and Hilbert Spaces. Mathematische Nachrichten, 147(1):185–203, 1990.

[8] Joudi Hajar, Taylan Kargin, Vikrant Malik, and Babak Hassibi. Distributionally Robust Kalman Filtering over an Infinite-Horizon. In ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 1–5, April 2025.

[9] Astghik Hakobyan and Insoon Yang. Wasserstein Distributionally Robust Control of Partially Observable Linear Stochastic Systems. IEEE Transactions on Automatic Control, 69(9):6121–6136, September 2024.

[10] Jiarao Huang, Dawei Shi, and Tongwen Chen. Robust event-triggered state estimation: A risk-sensitive approach. AUTOMATICA, 99:253–265, January 2019.

[11] Jiarao Huang, Dawei Shi, and Tongwen Chen. Event-triggered robust state estimation for systems with unknown exogenous inputs. Automatica, 122:109248, December 2020.

[12] Zhijia Huang, Luis D. Couto, Chitra Dangwal, Shujie Xiao, Wei Lv, Dong Zhang, and Scott J. Moura. On Electrochemical Model-Based State Estimation for Lithium–Sulfur Batteries. IEEE Transactions on Control Systems Technology, 32(3):849–861, May 2024.

[13] Md. Monirul Islam, Xiao Zhong, Zeyi Sun, Haoyi Xiong, and Wenqing Hu. Real-time Frequency Regulation Using Aggregated Electric Vehicles in Smart Grid. Computers & Industrial Engineering, 134:11–26, August 2019.

[14] Minhyuk Jang, Astghik Hakobyan, and Insoon Yang. On the Steady-State Distributionally Robust Kalman Filter. In 2025 IEEE 64th Conference on Decision and Control (CDC), pages 2385–2392, December 2025.

[15] Ali Jokar, Barzin Rajabloo, Martin Désilets, and Marcel Lacroix. Review of Simplified Pseudo-two-Dimensional Models of Lithium-ion batteries. Journal of Power Sources, 327:44–55, September 2016.

[16] Sang-Bum Kim and Sang-Hyun Lee. Application of Deep Learning Techniques for the State of Charge Prediction of Lithium-Ion Batteries. Applied Sciences, 14(17):8077, January 2024.

[17] Philip Kollmeyer, Carlos Vidal, Mina Naguib, and Michael Skells. LG 18650HG2 Li-ion Battery Data and Example Deep Neural Network xEV SOC Estimator Script, March 2020.

[18] Daniel Kuhn, Peyman Mohajerin Esfahani, Viet Anh Nguyen, and Soroosh Shafieezadeh-Abadeh. Wasserstein Distributionally Robust Optimization: Theory and Applications in Machine Learning. In Operations Research & Management Science in the Age of Analytics, INFORMS TutORials in Operations Research, chapter 6, pages 130–166. INFORMS, October 2019.

[19] Xinyou Lin, Yunliang Tang, Jing Ren, and Yimin Wei. State of Charge Estimation with the Adaptive Unscented Kalman Filter Based on an Accurate Equivalent Circuit Model. Journal of Energy Storage, 41:102840, September 2021.

[20] Zheng Liu, Zhenhua Zhao, Yuan Qiu, Benqin Jing, Chunshan Yang, and Huifeng Wu. Enhanced State of Charge Estimation for Li-ion Batteries Through Adaptive Maximum Correntropy Kalman Filter with Open Circuit Voltage Correction. Energy, 283:128738, November 2023.

[21] Valentina Lucaferri, Michele Quercio, Antonino Laudani, and Francesco Riganti Fulginei. A Review on Battery Model-Based and Data-Driven Methods for Battery Management Systems. Energies, 16(23):7807, January 2023.

[22] Changlong Ma, Chao Wu, Luoya Wang, Xueyang Chen, Lili Liu, Yuping Wu, and Jilei Ye. A Review of Parameter Identification and State of Power Estimation Methods for Lithium-ion Batteries. Processes, 12(10):2166, October 2024.

[23] Mohamed A. A. Mohamed, Tung Fai Yu, Grace Ramsden, James Marco, and Thomas Grandjean. Advancements in Parameter Estimation Techniques for 1RC and 2RC Equivalent Circuit Models of Lithium-ion Batteries: A Comprehensive Review. Journal of Energy Storage, 113:115581, March 2025.

[24] Scott Moura, Federico Bribiesca Argomedo, Reinhardt Klein, Anahita Mirtabatabaei, and Miroslav Krstic. Battery State Estimation for a Single Particle Model With Electrolyte Dynamics. IEEE Transactions on Control Systems Technology, 25:1–16, June 2016.

[25] Kiarash Movassagh, Arif Raihan, Balakumar Balasingam, and Krishna Pattipati. A Critical Look at Coulomb Counting Approach for State of Charge Estimation in Batteries. Energies, 14(14):4074, January 2021.

[26] Viet Anh Nguyen, Soroosh Shafieezadeh-Abadeh, Daniel Kuhn, and Peyman Mohajerin Esfahani. Bridging Bayesian and Minimax Mean Square Error Estimation via Wasserstein Distributionally Robust Optimization. Math. Oper. Res., 48(1):1–37, February 2023.

[27] Julien Pallage and Antoine Lesage-Landry. Wasserstein Distributionally Robust Shallow Convex Neural Networks. INFORMS Journal on Optimization, 8(1):61–93, January 2026.

[28] Xianghua Peng, Jingyuan Yin, Longfei Sun, Zeyu Ye, and Tongzhen Wei. Parameter Identification of Lithium-ion Battery Equivalent Circuit Model Based on Limited Memory Recursive Least Squares Algorithm with Variable Forgetting Factor. J. Phys.: Conf. Ser., 2246(1):012090, April 2022.

[29] Gregory L. Plett. Battery Management Systems, Volume II: Equivalent-Circuit Methods. Artech House, December 2015.

[30] Dengliang Qi, Jingan Feng, Xiangdong Ni, and Lei Wang. Maximum Correntropy Extended Kalman Filter for Vehicle State Observation. Int.J Automot. Technol., 24(2):377–388, April 2023.

[31] Soroosh Shafieezadeh Abadeh, Viet Anh Nguyen, Daniel Kuhn, and Peyman Mohajerin Mohajerin Esfahani. Wasserstein Distributionally Robust Kalman Filtering. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, Advances in Neural Information Processing Systems, volume 31. Curran Associates, Inc., 2018.

[32] Qiang Song, Yuxuan Mi, and Wuxuan Lai. A Novel Variable Forgetting Factor Recursive Least Square Algorithm to Improve the Anti-Interference Ability of Battery Model Parameters Identification. IEEE Access, 7:61548–61557, 2019.

[33] Weihua Song, Ranran Liu, Xiaona Jin, and Wei Guo. SOC Estimation for Lithium-ion Batteries Based on Weighted Multi-Innovation Sage–Husa Adaptive EKF. Energies, 18(16), August 2025.

[34] Xiangdong Sun, Jingrun Ji, Biying Ren, Chenxue Xie, and Dan Yan. Adaptive Forgetting Factor Recursive Least Square Algorithm for Online Identification of Equivalent Circuit Model Parameters of a Lithium-ion Battery. Energies, 12(12):2242, January 2019.

[35] Junjie Tao, Shunli Wang, Wen Cao, Paul Takyi-Aninakwa, Carlos Fernandez, and Josep M. Guerrero. A Comprehensive Review of State of Charge and State of Health Estimation for Lithium-ion Battery Energy Storage Systems. Ionics, 30(10):5903–5927, October 2024.

[36] Zhuo Wang, Daniel T. Gladwin, Matthew J. Smith, and Stefan Haass. Practical State Estimation Using Kalman Filter Methods for Large-Scale Battery Systems. Applied Energy, 294:117022, July 2021.

[37] Zuolu Wang, Guojin Feng, Dong Zhen, Fengshou Gu, and Andrew Ball. A Review on Online State of Charge and State of Health Estimation for Lithium-ion Batteries in Electric Vehicles. Energy Reports, 7:5141–5161, November 2021.

[38] Kei Long Wong, Michael Bosello, Rita Tse, Carlo Falcomer, Claudio Rossi, and Giovanni Pau. Li-Ion Batteries State-of-Charge Estimation Using Deep LSTM at Various Battery Specifications and Discharge Cycles. In Proceedings of the Conference on Information Technology for Social Good, GoodIT '21, pages 85–90, New York, NY, USA, September 2021. Association for Computing Machinery.

[39] Muyao Wu, Linlin Qin, and Gang Wu. State of Charge Estimation of Power Lithium-ion Battery Based on an Adaptive Time Scale Dual Extended Kalman filter. Journal of Energy Storage, 39:102535, July 2021.

[40] Jaejung Yun, Yeonho Choi, Jaehyung Lee, Seonggon Choi, and Changseop Shin. State-of-Charge Estimation Method for Lithium-ion Batteries Using Extended Kalman Filter With Adaptive Battery Parameters. IEEE Access, 11:90901–90915, 2023.

[41] Dong Zhang, Satadru Dey, Luis D. Couto, and Scott J. Moura. Battery Adaptive Observer for a Single-Particle Model With Intercalation-Induced Stress. IEEE Transactions on Control Systems Technology, 28(4):1363–1377, July 2020.

[42] Xuan Zheng and Zhuqian Zhang. State of Charge Estimation at Different Temperatures Based on Dynamic Thermal Model for Lithium-ion Batteries. Journal of Energy Storage, 48:104011, April 2022.

[43] Wenlu Zhou, Yanping Zheng, Zhengjun Pan, and Qiang Lu. Review on the Battery Model and SOC Estimation Method. Processes, 9(9):1685, September 2021.