Academic Journal of Applied Sciences
ISSN: 3105-0549 (Print), ISSN: 3105-0557 (Online) | Vol. 4, No. 1, 2026

A Four-Timescale Continuous-Time Model with Sliding-Window Calibration for State-of-Charge Prediction on Battery-Powered Devices

Haiyang Yu *

Nanjing University of Science and Technology, Nanjing, China

* **Corresponding author**: (Email: 3334177044@qq.com)

Abstract: Predicting how long a battery-powered device will run before its state of charge (SOC) reaches a usable limit is a forward problem in continuous time: the quantity of interest is a first-passage time of a state driven by inputs whose timescales differ by about seven orders of magnitude, from processor load within seconds to capacity loss over the service life. A model resolved on one timescale misplaces the others. This paper develops a continuous-time model in which one state equation is closed by four terms, each owning a band: a discharge-power decomposition into baseline, user and environmental parts, an hourly integration whose power coefficients are recalibrated on a sliding window, an aging coefficient on the day-to-month band, and an Arrhenius–solid-electrolyte-interphase capacity trajectory beyond a month. Sixteen public datasets supply the coefficients and the usage statistics, and nine of them carry the parameters the model needs and serve as validation runs. Across the nine runs the mean absolute error is 6.0786 and the root mean square error is 7.0300 in SOC percentage points. The predicted mean falls below the observed mean in every run, and across the five tabulated runs the gap is 4.20 to 5.68 points, with that offset acc ounting for 97% to 99% of the mean absolute error in four of them, which places the dominant error in the calibration rather than in the structure. At full charge the predicted endurance runs from 3.0 h under extreme usage to 11.9 h in low-power mode, a factor of 4.0 that a user can move within by changing settings alone.

Keywords: State of Charge, Multi-Timescale Continuous-Time Model, Sliding-Window Calibration.

1. Introduction

The discharge rate of a battery-powered device is governed by different terms at different horizons. Within an hour the workload in execution dominates; over a month the capacity lost to aging does, since fluctuations of the load average out while the loss of cyclable lithium does not. A working day lies between the two, and the three timescales differ by about seven orders of magnitude. Predicting the time left before the charge reaches a usable limit therefore requires all three scales, and an estimate that omits one of them drifts as the horizon lengthens.

Two properties of the underlying dynamics complicate the estimate. The discharge rate is nonlinear in ambient temperature, because internal resistance, ionic diffusion and thermal-management load all vary with temperature and none of them vary linearly. The capacity that normalises the discharge rate is itself a slow state: coefficients identified on a fresh cell drift as the cell ages, and a parameter set fixed at manufacture ceases to describe the device once a tenth of the usable capacity has been lost.

Three research lines supply partial answers. Observer-based circuit models estimate the present state of charge from terminal voltage and current, and they presuppose current sensing and a validated circuit representation. Data-driven regressors map usage features onto state of charge or state of health, with accuracy that holds inside the training support. Mechanistic aging studies describe capacity loss over months. Each line reports on a single clock, whereas the forward estimate spans all three.

The formulation used below keeps one state equation and lets the four bands supply its terms. A prediction at any horizon therefore rests on the same coefficient set, from the sub-hour scale on which the setting vector fixes the instantaneous drain to the year scale on which interphase growth sets the capacity. Capacity, nominal panel power and the thermal coefficient enter as coefficients, so a change of device changes the coefficient set and leaves the form of the equation intact. Sixteen public datasets supply those coefficients and the usage statistics, and the coefficients used below are calibrated and validated on the lithium-ion cell of a smartphone. The model returns the endurance at a given setting, and the validation fixes the error attached to that number.

The contributions of this paper are as follows:

(1) A four-band closure of one state equation in which the rated capacity is promoted from a constant to an Arrhenius–solid-electrolyte-interphase trajectory, so that usage intensification and capacity fade enter a single expression with separate exponents.

(2) A sliding-window calibration that holds the power coefficients at their daily value and corrects them from the residual of the previous step, with the window length as the only resolution parameter of the fit.

(3) A validation on nine public runs whose error decomposition separates a level offset from point-to-point scatter and shows which of the two the calibration controls.

2. Related Work

2.1. State-of-Charge Estimation

Reviews of SOC estimation separate model-based from data-driven routes and compare them on accuracy, sensing burden and computational cost [1][2]. The model-based side is dominated by equivalent-circuit representations identified from pulse or drive-cycle data, and recent work resolves fast

polarization and slow diffusion on separate horizons [3]. These methods deliver present SOC under a measured load, with error budgets written for sensor noise and model order. The data-driven side maps voltage, current and temperature histories onto SOC without a physical circuit, at the cost of a training campaign that has to cover the operating envelope.

2.2. Capacity Fade over Months

Aging models work on the other clock. Mechanistic formulations tie capacity loss to interphase growth, with the film thickness following a square-root-of-time law whose rate constant is thermally activated [4, 5]. Reviews of remaining-life prediction catalogue the regression and sequence models applied to the same quantity and note their dependence on the coverage of the training campaign [6]. The two families meet in the capacity term: a mechanistic fade law supplies the trajectory that a usage model needs, and a usage model supplies the duty cycle over which a fade law averages.

2.3. Device Energy Models and Parameter Identification

Device-level energy models fall into behaviour models, which infer usage structure from measured traces [7], and component models, which fit a display power law from instrumentation [8] or cost uplink and downlink traffic [9]. An endurance estimate needs the component split, because the terms it produces respond to different settings. Temperature enters every one of those terms: internal resistance, diffusion and thermal-management load all move with it, and a review of the temperature effect collects the mechanisms and their relative weight [10]. Identification of such models is a sensitivity problem in its own right, since a parameter that the output barely responds to cannot be fitted from data with any confidence [11].

3. Four-Timescale Continuous-Time Model

Four bands organise the model, split where the dominant term changes. Under an hour the setting vector and the hardware response fix the drain; between an hour and a day the coefficients drift and the integration carries the state; between a day and a month the usable capacity starts to shrink; beyond a month the capacity trajectory sets the trend on its own. One state equation runs through all four, and each band contributes the term that dominates inside it.

3.1. Discharge Power Decomposition

Charge bookkeeping on the cell gives the state equation, with $Q(t)$ the charge held by the cell. The regime analysed holds the charging power at zero:

$$\frac{dQ(t)}{dt} = -P_{\text{discharge}}(t) + P_{\text{charge}}(t) \tag{1}$$

With $P_{\text{charge}} \equiv 0$ the state equation measures the drain alone. A charging term would tie the result to infrastructure availability and dwell time, which mixes the determinants of drain with the schedule of the charger; the discharge-only regime keeps the two separable. Discharge power separates into a baseline, a user term and an environment term:

$$P_{\text{discharge}}(t) = P_{\text{base}} + \Delta P_{\text{user}}(t) + \Delta P_{\text{env}}(t) \tag{2}$$

The user term collects the application load and the three hardware domains that respond to it, with N_{app} the number of concurrently active applications, $s(t)$ the normalised screen brightness, $\mathbf{1}_{\text{on}}(t)$ the screen-on indicator and u_{CPU}, u_{GPU} the processor utilisations:

$$\begin{cases} \Delta P_{\text{app}}(t) = \beta N_{\text{app}}(t) \\ \Delta P_{\text{screen}}(t) = P_{\text{screen,base}}\, s(t) \mathbf{1}_{\text{on}}(t) \\ \Delta P_{\text{CPU}}(t) = P_{\text{CPU,base}} + P_{\text{CPU,act}}\, u_{\text{CPU}}(t) \\ \Delta P_{\text{GPU}}(t) = P_{\text{GPU,base}} + P_{\text{GPU,act}}\, u_{\text{GPU}}(t) \end{cases} \tag{3}$$

Both processor terms carry a floor and a load-proportional part, so their drain stays positive at zero utilisation, and β is the incremental power of one active application. Brightness scales a nominal panel power, which removes the panel contribution at a dark screen and leaves the rest of the display path in place.

Two effects sit in the environment term. Temperature acts twice, as a global efficiency factor η_T and as an explicit thermal-management term, with κ and ζ the linear and quadratic temperature coefficients and ΔT the deviation from the reference temperature T_{ref}. The second-order polynomial trades the nonlinearity of internal resistance against the cost of evaluating the model at every step:

$$\begin{cases} \Delta T(t) = T(t) - T_{\text{ref}} \\ \eta_T(t) = 1 + \kappa \Delta T(t) + \zeta \Delta T^2(t) \\ \Delta P_{\text{temp}}(t) = \lambda \Delta T(t) \eta_T(t) \end{cases} \tag{4}$$

Weak signal raises transmit power, so the radio term sums over location classes, with L the indoor, suburban and subway classes, $P_{l,\text{avg}}$ the power increment of class l, u_{comm} the communication utilisation and α_{comm} its coefficient:

$$\Delta P_{\text{loc}}(t) = \sum_{l \in \mathrm{L}} \mathbf{1}_l(t) P_{l,\text{avg}} \left[1 + \alpha_{\text{comm}} u_{\text{comm}}(t) \right] \tag{5}$$

3.2. Diurnal Integration and Sliding-Window Calibration

Between one hour and one day the state equation integrates at hourly steps with the capacity held at its daily value:

$$\mathrm{SOC}(h) = \mathrm{SOC}(h-1) - \frac{1}{Q_{\text{rated}}(d)} \int_{h-1}^{h} P_{\text{discharge}}(\tau)\, d\tau \tag{6}$$

Updating the coefficients at every sample tracks the load most closely and costs the most: it requires a current measurement the platform does not always expose, and it slows the loop on a device that is also doing the work being measured. The calibration runs on a window instead. A learning-rate update corrects the power coefficient η from the residual of the previous step, and a moving average of the same length supplies the parameter value used in the

integration:

$$\begin{cases} \hat{\eta}(t) = \eta(t-\Delta t) + \alpha \dfrac{P_{\text{act}}(t) - P_{\text{pred}}(t)}{P_{\max}} \\ \bar{Y}(t) = \dfrac{1}{N}\sum_{i=1}^{N} Y(t - i\Delta t) \end{cases} \tag{7}$$

Here P_{act} and P_{pred} are the measured and predicted power, $P_{\max}$ the normalising power, Y the coefficient being smoothed, $\bar{Y}$ its window average, N the window length and α the learning rate. A short window adapts within the hour and passes measurement noise into the coefficient; a long window smooths the coefficient and lags a change of regime. The learning rate stays small enough that the update does not oscillate when the residual changes sign between steps.

3.3. Aging-Coupled Capacity Fade

Aging enters through the denominator. Between one day and one month the correction is carried by an aging coefficient $k_{\text{aging}}(d)$ that reduces the usable capacity of day d, with $E_{\text{discharge}}(d)$ the energy drawn on that day:

$$\text{SOC}_{\text{end}}(d) = \text{SOC}_{\text{start}}(d) - \frac{E_{\text{discharge}}(d)}{Q_{\text{rated}}(d)\left[1 - k_{\text{aging}}(d)\right]} \tag{8}$$

The coefficient is calibrated on the day-scale decline of usable capacity.

Beyond one month the capacity trajectory dominates, because a fluctuation of the instantaneous drain averages out over a month while the loss of cyclable lithium does not. The aggregate drain becomes a smooth power law in time, $\bar{k}t^{q}$ with $\bar{k}$ a rate constant and $q \geq 0$ the intensification exponent.

Capacity fade follows the interphase-growth argument. The film thickens by diffusion of the reacting species through the film itself, so the growth rate is inversely proportional to the thickness, and the thermally activated diffusivity converts the result into a square-root law:

$$\frac{dL}{dt} = \frac{D}{L}, \qquad L(t) = \sqrt{2Dt}, \qquad D = D_0 \exp\left(-\frac{E_a}{RT}\right) \tag{9}$$

Here L is the film thickness, D the diffusivity, D_0 its pre-exponential factor, E_a the activation energy and R the gas constant. With the loss of cyclable lithium proportional to the film thickness, the rated capacity becomes

$$Q_{\text{rated}}(t,T) = Q_0 - k_0 t^{p} \exp\left(-\frac{E_a}{2RT}\right) \tag{10}$$

Where Q_0 is the initial capacity, $k_0 = \kappa_0\sqrt{2D_0}$ collects the geometric and diffusivity constants, and $p \approx 0.5$ is the growth exponent set by the diffusion mechanism. Linearising the Arrhenius factor about a reference temperature and substituting the aggregate drain and the capacity trajectory into the state equation gives the unified form:

$$\frac{d\text{SOC}}{dt} = \frac{\bar{k}\,t^{q}}{Q_0 - k_0 t^{p} \exp\left(-\dfrac{E_a}{2RT_{\text{ref}}}\right)\left[1 + \dfrac{E_a}{2RT_{\text{ref}}^{2}}\left(T - T_{\text{ref}}\right)\right]} \tag{11}$$

The numerator carries usage intensification and the denominator carries degradation. The two exponents measure different things: q rises when the device is used harder as time goes on and sits at zero for a settled routine, while p is pinned near one half by diffusion. Because $p < 1$, the capacity term stays finite over years; because the numerator grows as t^{q}, a sustained increase in usage eventually outruns the fade.

4. Validation on Public Datasets

4.1. Datasets and Runs

Sixteen datasets were assembled. Eight of them are the principal sources of coefficients and usage statistics: the HIRF battery set and the battery aging data for the discharge and fade parameters, the user behaviour dataset for the load statistics, the Samsung phone and smartphone specification data for capacity and panel power, the Antutu and smartphone CPU statistics for the processor utilisations, and the EV predictive maintenance records for cell-level degradation. Nine of the sixteen carried the parameters the model needs and served as validation runs, and the remainder supplied estimates for parameters without direct measurement. Each run contributes one hundred samples, the error is reported in SOC percentage points, and the averages over the nine runs are an MAE of 6.0786 and an RMSE of 7.0300.

Table 1. Per-run validation metrics in SOC percentage points

Dataset	Samples	RMSE	MAE	Mean observed	Mean predicted
User behaviour	100	11.50	9.94	67.23	63.03
Smartphone list	100	6.11	5.40	69.21	63.89
Samsung phones	100	6.70	5.77	69.93	64.25
ML benchmarks	100	6.15	5.24	69.08	63.96
Antutu benchmarks	100	5.91	5.12	68.77	63.80

4.2. Structure of the Error

Table 1 gives five of the nine runs, each labelled by the dataset the run draws on. In every run the predicted mean sits below the observed mean; across the five tabulated runs the gap is 4.20 to 5.68 points, a relative bias between 6.2% and 8.1%. In four of them the offset accounts for 97% to 99% of the MAE. The residual is a displaced level, not point-to-point scatter. The user-behaviour run departs from the pattern: the offset there explains 42% of the MAE and the rest comes from scatter.

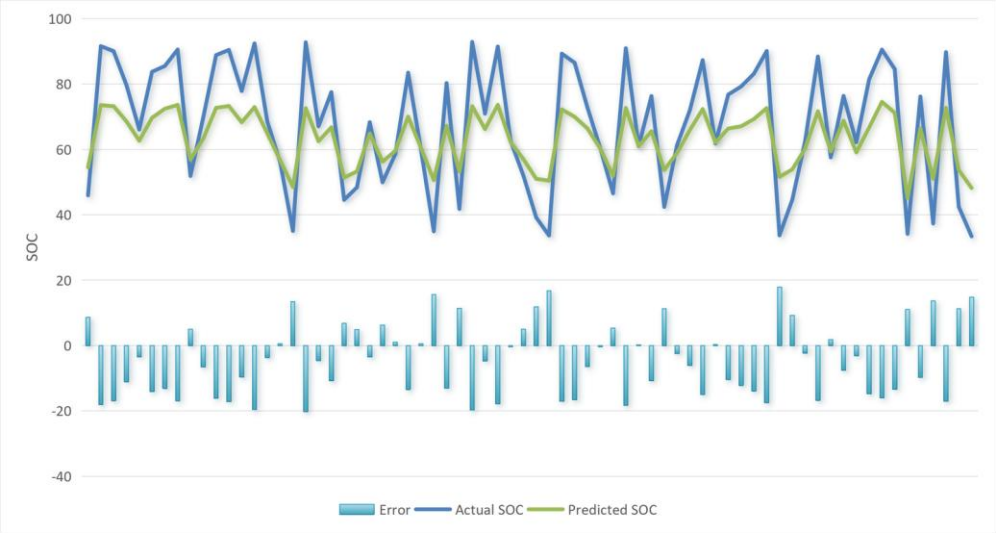


Figure 1. Measured and predicted state of charge over one validation run, with the pointwise residual below

Figure 1 shows that run over a full day. The predicted trajectory follows the measured one through the charge and discharge swings, stays below it for the whole of the record, and the residual bars oscillate between about -20 and +16 points about a level near -4. Two features of the record stand out: a displacement that holds across the whole of it, and a spread of comparable size. Where the first dominates, a single coefficient set applied to a heterogeneous fleet produces a level error; where the second dominates, the window calibration has not resolved the transients of that device.

4.3. What the Offset Implies

A level offset is the signature of a coefficient set that was identified elsewhere. The nine runs span device classes with different idle draw and different rated capacity, and one set of coefficients cannot reproduce all of their means. Per-device recalibration of the baseline power and the capacity needs one discharge record, and it should remove most of the error before any change of the model structure is considered.

What the offset leaves unexplained sits in the high-frequency band. A model calibrated on a window returns a displaced level and misses transients of a few samples, which is what the residual bars of Figure 1 show. The window length therefore sets the resolution of the model: endurance is an hour-scale quantity, and a minute-level prediction needs a circuit representation with a current measurement rather than a windowed power coefficient.

5. Endurance Prediction

5.1. Across Six Usage Settings

Endurance in what follows is the time at which the state reaches a 20% cutoff, the level at which a device enters a saving mode. The six settings of Table 2 fix the inputs for the calculation.

Table 2. Setting Vectors for the Six Usage Scenarios

Scenario	Vector
Normal Usage	[0.5, 0.4, 0.3, 0.1, 0.3, 0.5, 25]
Heavy Gaming	[0.9, 0.8, 0.9, 0.8, 0.6, 0.8, 30]
Video Streaming	[0.8, 0.7, 0.4, 0.3, 0.7, 0.3, 28]
Outdoor Navigation	[0.7, 0.6, 0.5, 0.4, 0.8, 0.6, 35]
Low Power Mode	[0.3, 0.2, 0.2, 0.1, 0.2, 0.2, 22]
Extreme Usage	[1.0, 1.0, 1.0, 0.9, 0.9, 0.9, 40]

Entries 1 to 6 are the normalised brightness, CPU utilisation, GPU utilisation, communication utilisation, location intensity, which weights the location classes introduced with the radio term, and application load, all on the unit interval. Entry 7 is the ambient temperature in degrees Celsius. Figure 2 plots the predicted endurance against the initial charge for the six settings. At full charge the endurance runs from 3.0 h under extreme usage and 4.0 h under heavy gaming, through 5.0 h for outdoor navigation and 6.1 h for video streaming, to 6.6 h under normal usage and 11.9 h in low-power mode. The spread between the extreme and the low-power setting is a factor of 4.0, which is the range a user can move within without changing hardware.

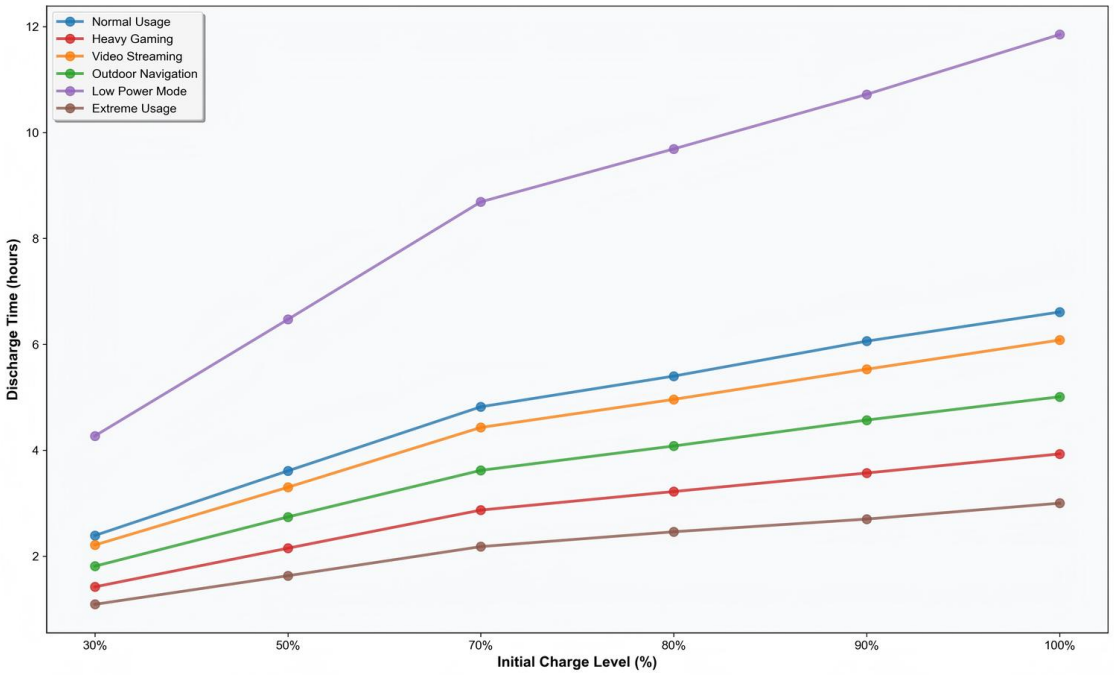


Figure 2. Predicted time to the 20% level against initial charge for the six usage scenarios

5.2. Shape of the Curves

Close to straight lines in the initial charge, they have a slope between about 0.3 h and 1.1 h per ten percentage points depending on the scenario. Near-linearity follows from the model: none of the terms in the decomposition depends on the state except through the damping of the load, which is small over the horizon plotted. In planning, the figure settles the question that matters: a charge of 30% buys about 2.4 h under normal usage and about 1.0 h under extreme usage, so the same remaining charge supports a different plan depending on the setting.

6. Conclusions

Across the six settings the predicted endurance spans 3.0 h to 11.9 h at full charge, a factor of 4.0, and the model keeps the error at 6.08 in mean absolute terms across nine validation runs. The four bands share one state equation, with usage intensification in the numerator and interphase-driven capacity fade in the denominator, and the rated capacity is what carries the aging axis from a constant into a trajectory.

The error that remains sits in the calibration, and one discharge record per device should remove most of it before the structure of the model is changed. Two bounds follow from the assumptions. The thermal term is a second-order polynomial and stops describing the drop in usable capacity at low temperature, and the capacity trajectory is a prior taken from published campaigns, not a measurement on the validated devices. Within those bounds, transfer to another device class rests on the coefficient set, since the decomposition, the integrator and the fade term carry over unchanged.

References

[1] Albakri, M., & Darwish, A. (2026). A review of modelling, state of charge estimation and management methods of EV lithium-ion batteries. Batteries, 12(3), 92. https://doi.org/10.3390/batteries12030092

[2] Peng, J., Qian, X., Xu, X., et al. (2025). State of charge estimation techniques for lithium-ion batteries: A review. Energy Technology, 13(12), 2500542. https://doi.org/10.1002/ente.202500542

[3] Wang, K., Jia, T., Wang, M., et al. (2026). An adaptive timescale decoupled parameter identification equivalent circuit model for lithium-ion battery capacity estimation. Journal of Energy Storage, 161, 122024. https://doi.org/10.1016/j.est.2026.122024

[4] Zhu, W., Zhou, P., Ren, D., et al. (2022). A mechanistic calendar aging model of lithium-ion battery considering solid electrolyte interface growth. International Journal of Energy Research, 46(11), 15521–15534. https://doi.org/10.1002/er.8249

[5] Karger, A., Schmitt, J., Kirst, C., et al. (2023). Mechanistic calendar aging model for lithium-ion batteries. Journal of Power Sources, 578, 233208. https://doi.org/10.1016/j.jpowsour.2023.233208

[6] Samanta, A., & Williamson, S. (2023). Machine learning-based remaining useful life prediction techniques for lithium-ion battery management systems: A comprehensive review. IEEJ Journal of Industry Applications, 12(4), 563–574. https://doi.org/10.1541/ieejjia.22004793

[7] Ding, M., Wang, T., & Wang, X. (2021). Establishing smartphone user behavior model based on energy consumption data. ACM Transactions on Knowledge Discovery from Data, 16(2), 1–40. https://doi.org/10.1145/3461459

[8] Choi, Y., Ha, R., & Cha, H. (2018). Fully automated OLED display power modeling for mobile devices. Pervasive and Mobile Computing, 50, 41–55. https://doi.org/10.1016/j.pmcj.2018.07.006

[9] Ali, J., & Altamimi, M. (2022). Energy consumption model for data transfer in smartphone. Computer Communications, 182, 13–21. https://doi.org/10.1016/j.comcom.2021.10.014

[10] Ma, S., Jiang, M., Tao, P., et al. (2018). Temperature effect and thermal impact in lithium-ion batteries: A review. Progress in Natural Science: Materials International, 28(6), 653–666. https://doi.org/10.1016/j.pnsc.2018.11.002

[11] Alqahtani, A. H., Fahmy, H. M., Hasanien, H. M., et al. (2024). Parameters estimation and sensitivity analysis of lithium-ion battery model uncertainty based on osprey optimization algorithm. Energy, 304, 132204. https://doi.org/10.1016/j.energy.2024.132204