Proceedings of the IMECE 2026
International Mechanical Engineering Congress & Exposition
IMECE2026
November 8-12, 2026, Vancouver, British Columbia, Canada

IMECE2026-191741

THE VIEWS EXPRESSED ARE THOSE OF THE AUTHOR AND DO NOT REFLECT THE
OFFICIAL POLICY OR POSITION OF THE DEPARTMENT OF WAR OR THE U.S. GOVERNMENT.

PHYSICS-CONSTRAINED NEURAL NETWORKS FOR ELECTRO-THERMAL BATTERY MODELING AND EMULATION

Connor Madden[1], Jarrett Peskar[1], Braden Holst[2], Malichi Flemming[1], Austin R.J. Downey[1,3,*], Kerry Sado[4], Ralph White[2], Jamil Khan[1]

[1]University of South Carolina, Department of Mechanical Engineering, Columbia, South Carolina
[2]University of South Carolina, Department of Chemical Engineering, Columbia, South Carolina
[3]University of South Carolina, Department of Civil and Environmental Engineering, Columbia, South Carolina
[4]University of South Carolina, Department of Electrical Engineering, Columbia, South Carolina

ABSTRACT

The maritime industry is becoming increasingly dependent on battery energy storage for the ability to buffer and supply pulsed loads to optimize efficiency of onboard power system. When designing electrified shipboard power systems, battery testing becomes expensive and time intensive, necessitating reliable models to emulate the electro-thermal response to loads. High fidelity physics-based methods can be computationally expensive and time consuming to parameterize. Thus, traditional physics-based models struggle with real-time digital twin applications, or hardware-in-the-loop testing. Alternatively, data driven methods have been explored, which are rapidly executable and less intensive to parameterize. However, they require obtaining larges sets of expensive (time and cost) experimental data and struggle when loads or conditions differ from training data. To address limitations of both methods, this work proposes a physics constrained data-driven electro-thermal battery model that is more generalizable across a wider portion of the battery performance space than pure machine learning methods; while also requiring less data and being simpler to develop than traditional physics-based approaches. This is accomplished by constraining neural network during training using both an equivalent circuit model represent the electrical components and a differential equations for temperature. The network learns how the internal states, including open circuit voltage, RC relaxation pairs, and DC internal resistance, and entropic heat coefficient evolve with time and current. During training, the model is fit to the measured voltage and temperature, while penalized to obey the governing equations. Furthermore, the model parameters are bound to physically reasonable values. Across a dynamic load profile the model shows a 65.8 mV RMSE in electrical potential, and 0.205°C RMSE in temperature. Model generalization was tested with a constant current load greater than any within the training dataset, where the model exhibited a 49.0 mV RMSE in electrical potential, and 1.02 °C RMSE in temperature.

Keywords: Deep neural networks, battery modeling, data-driven approach, hybrid pulse power characterization, online learning, electro-thermal prediction, activation functions, machine learning, energy storage, battery safety

1. INTRODUCTION

Battery systems are playing an increasingly central role in more-electrified platforms, often used to smooth load fluctuations, increase efficiency, and meet pulsed load demands in hybrid electric ships [1]. Across these applications, lithium-ion batteries are expected to operate not only as energy reservoirs, but also as dynamic components that interact continuously with power electronics, supervisory controllers, and thermal management systems [2]. As electrification expands into applications characterized by aggressive transients, pulse-power demands, and operation across broad environmental conditions, there is growing need for battery models that can reproduce both electrical and thermal behavior with sufficient fidelity for prediction, control development, safety analysis, and hardware or software validation [3]. In these settings, a model that captures terminal voltage alone is often insufficient. Voltage response, internal polarization, reversible and irreversible heat generation, and state evolution are tightly coupled, and modeling errors in one part of the system can propagate directly into erroneous predictions of power capability, thermal rise, or operational margin [4].

Equivalent-circuit models remain among the most widely used battery representations because they offer a practical compromise between physical interpretability and computational efficiency [5]. In contrast to high-fidelity electro-chemical models, equivalent-circuit models can often be evaluated rapidly enough for estimation, control, and hardware deployment applications

*Corresponding author: austindowney@sc.edu

Copyright © 2026 by ASME

while still preserving a structured relationship between internal states and terminal response [6]. However, conventional equivalent-circuit modeling workflows are often limited by the way parameters are identified and deployed. Often, parameter estimation is done through a sequence of independent fitting steps, lookup-table construction, manual interpolation, or hand-tuned parameter surfaces [7]. These procedures can become cumbersome when parameters vary strongly with state of charge and temperature, they may not generalize well when the battery is subjected to operating conditions outside the narrow subset used during identification. However, purely data-driven machine learning approaches have attracted increasing interest because of their ability to capture nonlinear relationships directly from experimental data [8]. Neural networks, in particular, can approximate highly nonlinear mappings between measured variables and battery response without requiring explicit derivation of every constitutive relationship [9]. Yet this flexibility comes with important drawbacks when applied without physical structure. A black-box model may fit training data well while producing nonphysical extrapolations, unrealistic internal trends, poor interpretability, or state evolutions that are difficult to reconcile with established battery behavior [10]. Model interpretability is a priority for many battery applications, especially those tied to estimation, control, and real-time deployment [11]. Quantities such as open-circuit voltage, ohmic resistance, entropic coefficient, and transient time scales carry direct physical meaning, and preserving that is vital for model trustworthiness and practical use.

Another challenge arises from parameter identifiability. Different experimental protocols excite different battery phenomena, and not all parameters are equally observable under all test conditions. Constant-current data may be well suited to reveal static voltage trends, baseline ohmic response, and thermal evolution over extended intervals, but it is not necessarily optimal for extracting fast and slow relaxation dynamics [12]. Conversely, pulse-power experiments such as hybrid pulse power characterization tests are specifically informative for transient polarization behavior and relaxation time scales, but are not the ideal source for establishing slower evolving or constant parameters of the model [13]. Attempting to learn all parameters simultaneously from all datasets can therefore blur the roles of individual parameters, reduce interpretability, and make optimization more sensitive to local minima or compensating errors between submodels.

These limitations motivate the need for a modeling framework that preserves the structure and interpretability of an electro-thermal equivalent-circuit representation while using machine learning only where it is most beneficial. To achieve this, a physics-constrained machine learning framework is developed to through neural approximators within a structured battery model. In this formulation, the neural networks serve as bounded maps for parameters that are expected to vary with state of charge and temperature, while the governing electrical and thermal equations remain explicitly defined. This approach allows the model to inherit the interpretability, deployability, and physical organization of a conventional equivalent-circuit model while benefiting from the expressiveness of machine learning.

The resulting model combines learned nonlinear submodels

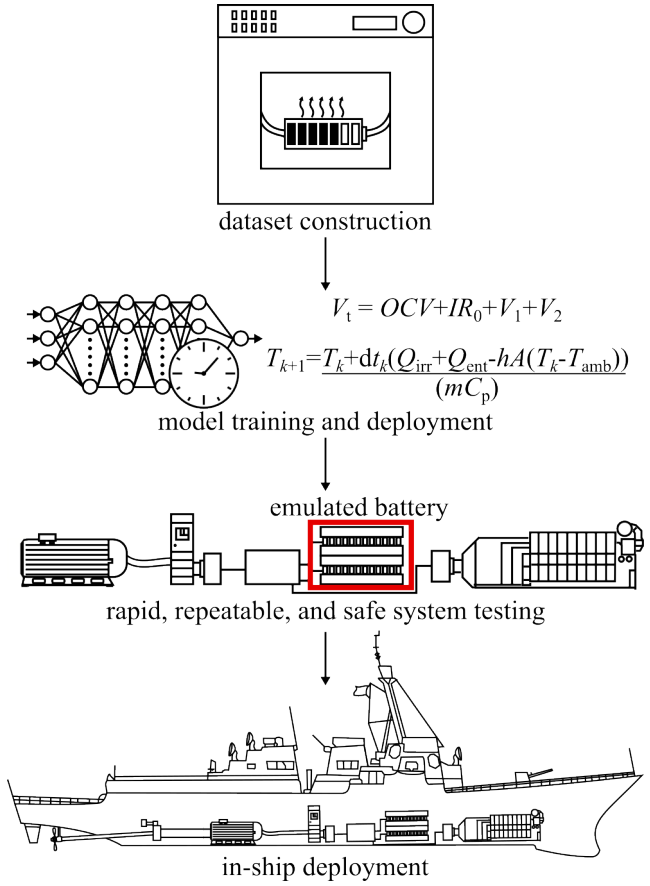


FIGURE 1: The process for developing and employing a battery emulator [14] through physics-constrained ML to improve system design and testing of advance shipboard architectures.

for open-circuit voltage, ohmic resistance, and entropic behavior with a low-order dynamic RC structure for transient response. Physical bounds and regularization are incorporated to enforce realistic behavior, and the electro-thermal formulation includes both irreversible and reversible heat generation. This produces a compact, interpretable model that can be deployed recursively under arbitrary current profiles while maintaining consistency between electrical and thermal predictions. Figure 1 illustrates the intended deployment workflow, beginning with experimental battery characterization, followed by physics-constrained model training, battery emulation, and eventual integration into shipboard system testing and design applications. The main contribution of this work is a physics-constrained machine learning framework for electro-thermal battery modeling that preserves the structure of an equivalent-circuit representation while learning nonlinear state-dependent behavior directly from experimental data. The framework emphasizes interpretability, staged identification, and deployment-oriented performance, providing a practical pathway for integrating data-driven modeling into system-level battery applications.

2. METHODOLOGY

This work develops a two-stage physics-constrained electro-thermal equivalent-circuit model for lithium-ion batteries. The framework combines bounded neural network submodels for static and thermodynamic quantities with a 2nd-order dynamic

Copyright © 2026 by ASME

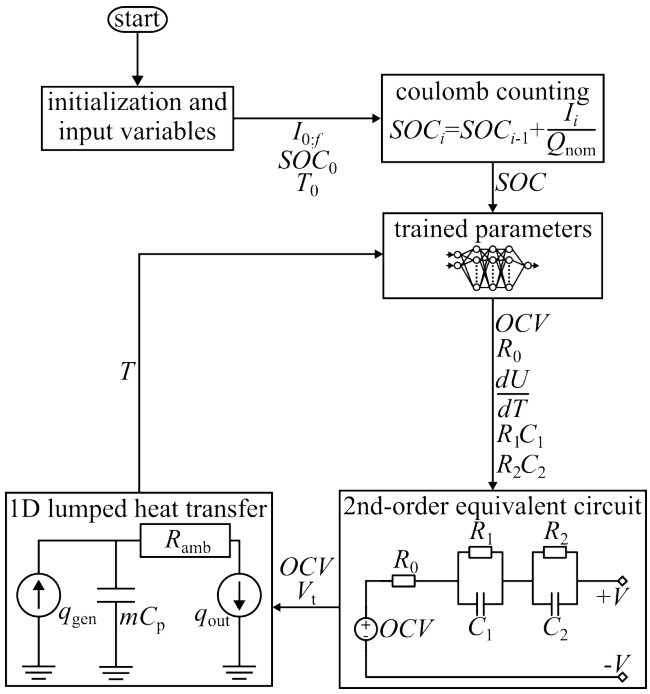


FIGURE 2: A diagram of the information exchanged between submodels of the electro-thermal physics-constrained ML model.

RC representation for transient behavior. The objective is to produce a model that remains interpretable, computationally efficient, and suitable for recursive deployment, while still capturing nonlinear dependence of battery behavior on state of charge and temperature. To achieve this, the electrical and thermal backbone is first identified from constant-current data, after which the dynamic polarization parameters are identified from Pulse Power Characterization (HPPC) data while the previously learned submodels are held fixed.

The public dataset used in this work [15] was previously used in a pure data driven modeling approach and is discussed in detail by Madden et al. [16]. This dataset consist of testing a single 3.6V 3000 mAh lithium nickel manganese cobalt (NMC) 18650 cell (Samsung 30Q) under constant current and HPPC tests. The cell was tested using a battery cycling system (NHR-9200) while held in a environmental chamber (IncuFridge Pro 328) to maintain a constant temperature. All sampling was done at 1 S/s.

2.1 Electro-thermal model structure and parameterization

Shown in figure 2, the deployed model is formulated as an electro-thermal equivalent-circuit representation driven by an input current profile. At each time-step, the model receives the current I_k, time-step duration Δt_k, ambient temperature $T_{\text{amb},k}$, and the current internal state. The internal state consists of the state of charge SOC_k, cell temperature T_k, and two RC polarization states, $v_{1,k}$ and $v_{2,k}$. The model then evaluates a set of learned and identified submodels to update voltage, thermal behavior, and state evolution.

The terminal voltage is expressed as

$$V_k = OCV(SOC_k, T_k) + I_k R_0(SOC_k, T_k) + v_{1,k} + v_{2,k}, \quad (1)$$

where $OCV(SOC, T)$ is the open-circuit voltage surface, $R_0(SOC, T)$ is the ohmic resistance surface, and v_1, v_2 are

TABLE 1: The global variables used in both the ECM and lumped thermal model, as well as the neural network hyperparameters.

physical values		network hyperparameters	
R_1	0.0011	$OCV(SOC,T)$ architecture	(64, 64, 16, 16)
C_1	3730	$R_0(SOC,T)$ architecture	(32, 32, 8, 8)
R_2	0.004	$dU/dT(SOC)$ architecture	(32, 32)
C_2	2225	learning rate	0.001
mC_p	45	training epochs	10000
hA	0.2	batch size	didk

polarization-state voltages associated with a second-order RC representation. The sign convention in (1) follows the implementation used throughout the framework and is retained consistently in training and deployment.

The polarization states evolve according to

$$a_1 = \exp\left(-\frac{\Delta t_k}{R_1 C_1}\right), \qquad a_2 = \exp\left(-\frac{\Delta t_k}{R_2 C_2}\right), \quad (2)$$

$$v_{1,k+1} = a_1 v_{1,k} + (1 - a_1)\, I_k R_1, \quad (3)$$

$$v_{2,k+1} = a_2 v_{2,k} + (1 - a_2)\, I_k R_2. \quad (4)$$

Here, R_1, C_1, R_2, and C_2 are identified during Stage 2 training and subsequently deployed as fixed global parameters. In this form, R_1 and R_2 govern the magnitudes of the polarization responses, while R_1C_1 and R_2C_2 determine the associated relaxation time scales. While polarization parameters are known to vary with state of charge and temperature, R_1, C_1, R_2, and C_2 are treated as global parameters in the present work to improve parameter identifiability and reduce coupling between the dynamic and electro-thermal submodels during training, and can be found in table table 1 along with other variables and hyperparameters.

The state of charge is advanced through coulomb counting,

$$SOC_{k+1} = SOC_k + \frac{I_k \Delta t_k}{Q\text{nom}}, \quad (5)$$

where Qnom is the rated cell capacity in coulombs. During recursive rollout, the state of charge is clipped to the range $[0, 1]$ to prevent drift outside the domain spanned by the training data. This ensures that the learned submodels are evaluated only within a physically meaningful range.

The thermal model is formulated as a lumped energy balance with irreversible and reversible heat generation,

$$q_{\text{irr},k} = |I_k|\,|OCV(SOC_k, T_k) - V_k|, \quad (6)$$

$$q_{\text{ent},k} = I_k T_k \frac{dU}{dT}(SOC_k), \quad (7)$$

$$T_{k+1} = T_k + \Delta t_k \frac{q_{\text{irr},k} + q_{\text{ent},k} - hA\,(T_k - T_{\text{amb}})}{mC_p}, \quad (8)$$

where mC_p is the effective thermal capacitance and hA is the lumped convective heat transfer coefficient. The term $dU/dT(SOC)$ is a learned entropic coefficient that captures the reversible thermal contribution as a function of state of charge. The dependence of dU/dT on SOC alone is a modeling simplification intended to reduce parameter coupling with $OCV(SOC, T)$

Copyright © 2026 by ASME

and $R_0(SOC, T)$, while retaining a physically interpretable representation of reversible heat generation.

Equations (1)–(8) define the complete deployed model. Given an initial condition $(SOC_0, T_0, v_{1,0}, v_{2,0})$ and an input current profile, the model may be iterated forward to predict terminal voltage, state of charge, cell temperature, and internal polarization-state evolution.

The deployed model depends on three learned submodels and four identified static parameters.

First, the open-circuit voltage submodel $OCV(SOC, T)$ is represented by a bounded neural network taking state of charge and normalized temperature as inputs. The architecture emphasizes state of charge as the dominant variable while allowing a smaller learned temperature correction. This choice reflects the expectation that OCV depends primarily on state of charge, with temperature introducing a secondary shift.

Second, the ohmic resistance submodel $R_0(SOC, T)$ is represented by a bounded neural network with the same inputs. As with the OCV model, the parameterization is structured so that SOC contributes the dominant effect, while temperature acts as a weaker correction. This allows the model to capture practical dependence of ohmic response on both electro-chemical state and thermal state.

Third, the entropic coefficient $dU/dT(SOC)$ is represented by a bounded neural network depending only on state of charge. Restricting this quantity to an SOC-only function reduces entanglement with the explicit temperature dependence already present in $OCV(SOC, T)$ and $R_0(SOC, T)$, while preserving a physically interpretable reversible heat term.

Each neural submodel uses a bounded output transform of the form

$$\hat{y} = y_{\min} + (y_{\max} - y_{\min})\,\sigma(z), \tag{9}$$

where z is the raw network output and $\sigma(\cdot)$ is the logistic sigmoid. This guarantees that each learned quantity remains within a prescribed physical range throughout training and deployment.

The dynamic polarization parameters R_1, C_1, R_2, and C_2 are not represented by neural networks. Instead, they are treated as static positive parameters identified from HPPC data. Positivity is enforced during optimization by learning unconstrained variables that are mapped through a softplus transform. Therefore, the resulting deployment model consists of nonlinear learned surfaces for static and thermodynamic behavior, combined with a conventional second-order RC transient structure for dynamic polarization behavior.

2.2 Training procedure

Outlined in figure 3, the training process is divided into two stages so that different experimental datasets inform different portions of the deployed model. Constant-current data is used to learn the static and thermodynamic submodels, while HPPC data is used to identify the dynamic polarization parameters. This is done because long constant-current experiments provide strong information for OCV behavior, ohmic response, and temperature evolution, whereas pulse-power experiments provide much stronger excitation of the dynamic RC states. Furthermore, this prevents the networks from attributing pulse behavior to false

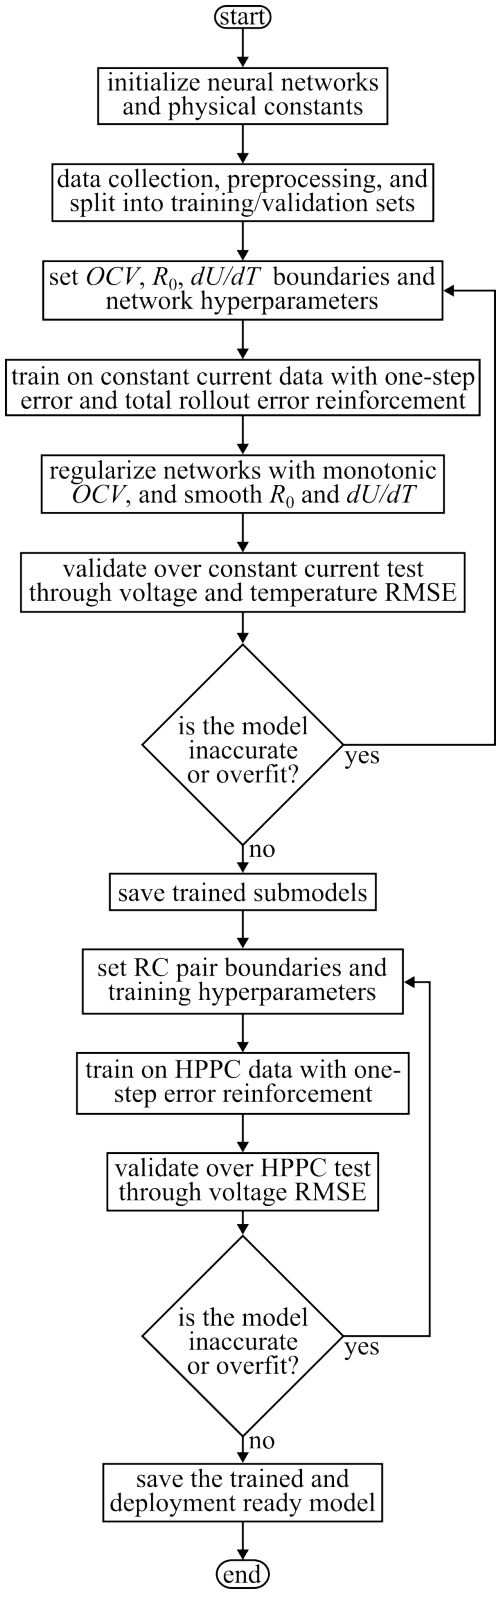


FIGURE 3: A flowchart explaining the training process of the electro-thermal physics-constrained ML battery model.

OCV drops or ohmic spikes, which is more likely when data is limited.

All datasets contain the columns current, voltage, temperature, and soc. If a time column is present, the time-step sequence

Copyright © 2026 by ASME

is inferred directly from it; otherwise a unit time-step is assumed. For each file, the current history, measured voltage, temperature, state of charge, and ambient temperature history are assembled into a sequence record.

Because experimental files can have different durations, all files in a stage are padded to the maximum sequence length in the batch. A binary mask $m_{k,b} \in \{0, 1\}$ is constructed to indicate whether time-step k of file b corresponds to a valid sample or a padded entry. This mask is applied to every data-fit loss, ensuring that padded values do not influence optimization. For a generic error sequence $e_{k,b}$, the masked mean for file b is computed as

$$\mathscr{L}_b = \frac{\sum_k m_{k,b} e_{k,b}}{\sum_k m_{k,b} + \varepsilon}, \tag{10}$$

where ε is a small positive constant for numerical stability. Stage losses are then formed by averaging these per-file quantities across files, so that long records do not dominate the optimization solely because they contain more time-steps, often due to lower currents.

For stage 1, the constant-current dataset is additionally split into a training subset and a held-out validation subset. The validation set is used only for total profile evaluation and model selection.

Stage 1 is designed to identify the electro-thermal backbone of the deployed model. The three trainable submodels in this stage are $OCV(SOC, T)$, $R_0(SOC, T)$, and $dU/dT(SOC)$. To improve both training stability and performance, stage 1 uses a hybrid objective consisting of teacher-forced, or one-step, losses and free-rollout, or total profile, losses.

In the teacher-forced branch, the measured state trajectories are used as inputs to the learned submodels. Specifically, the networks are evaluated on the measured state of charge and measured temperature histories,

$$\widehat{OCV}_{k,b} = OCV(SOC^{\text{meas}}_{k,b}, T^{\text{meas}}_{k,b}), \tag{11}$$

$$\hat{R}_{0,k,b} = R_0(SOC^{\text{meas}}_{k,b}, T^{\text{meas}}_{k,b}), \tag{12}$$

$$\widehat{\frac{dU}{dT}}_{k,b} = \frac{dU}{dT}(SOC^{\text{meas}}_{k,b}). \tag{13}$$

These are used to compute a teacher-forced voltage prediction,

$$V^{\text{teacher}}_{k,b} = \widehat{OCV}_{k,b} + I_{k,b}\hat{R}_{0,k,b}, \tag{14}$$

and a one-step teacher-forced temperature prediction,

$$T^{\text{teacher}}_{k+1,b} = T^{\text{meas}}_{k,b} + \Delta t_{k,b} \cdot \frac{q^{\text{teacher}}_{\text{irr},k,b} + q^{\text{teacher}}_{\text{ent},k,b} - hA\left(T^{\text{meas}}_{k,b} - T_{\text{amb},k,b}\right)}{mC_p}, \tag{15}$$

where the heat terms are evaluated using the teacher-forced states. This branch encourages accurate local behavior when the state trajectory is known.

However, one-step accuracy alone does not guarantee good recursive performance. To address this, stage 1 also includes a free-rollout branch. For each sequence, the simulation is initialized from the first measured state,

$$SOC_{0,b} = SOC^{\text{meas}}_{0,b}, \qquad T_{0,b} = T^{\text{meas}}_{0,b}, \tag{16}$$

and is then advanced according to (1)–(8) with $v_1 = v_2 = 0$. In other words, the model uses its own predicted SOC_k and T_k to evaluate OCV, R_0, and dU/dT at every subsequent time-step. This branch exposes the model to the same recursive error accumulation that arises during deployment.

The voltage and temperature data-fit terms are therefore constructed as

$$\mathscr{L}_{V,\text{teacher}} = \frac{1}{B}\sum_{b=1}^{B} \frac{\sum_k m_{k,b}\left(V^{\text{teacher}}_{k,b} - V^{\text{meas}}_{k,b}\right)^2}{\sum_k m_{k,b} + \varepsilon}, \tag{17}$$

$$\mathscr{L}_{T,\text{teacher}} = \frac{1}{B}\sum_{b=1}^{B} \frac{\sum_k m_{k,b}\left(T^{\text{teacher}}_{k,b} - T^{\text{meas}}_{k,b}\right)^2}{\sum_k m_{k,b} + \varepsilon}, \tag{18}$$

$$\mathscr{L}_{V,\text{free}} = \frac{1}{B}\sum_{b=1}^{B} \frac{\sum_k m_{k,b}\left(V^{\text{free}}_{k,b} - V^{\text{meas}}_{k,b}\right)^2}{\sum_k m_{k,b} + \varepsilon}, \tag{19}$$

$$\mathscr{L}_{T,\text{free}} = \frac{1}{B}\sum_{b=1}^{B} \frac{\sum_k m_{k,b}\left(T^{\text{free}}_{k,b} - T^{\text{meas}}_{k,b}\right)^2}{\sum_k m_{k,b} + \varepsilon}. \tag{20}$$

To encourage physically plausible submodels, several regularization terms are added. First, a monotonicity penalty is imposed on negative values of $\partial OCV/\partial SOC$,

$$\mathscr{L}_{\text{mono}} = \frac{\sum_i \max\left(0, -\frac{\partial OCV_i}{\partial SOC_i}\right)}{N + \varepsilon}, \tag{21}$$

where the derivatives are obtained by automatic differentiation over the flattened training batch. Second, a smoothness penalty is applied to the resistance surface,

$$\mathscr{L}_{R_0,\text{smooth}} = \frac{\sum_i \left[\left(\frac{\partial R_{0,i}}{\partial SOC_i}\right)^2 + 0.1\left(\frac{\partial R_{0,i}}{\partial T_i}\right)^2\right]}{N + \varepsilon}, \tag{22}$$

and a smoothness penalty is also applied to the entropic coefficient,

$$\mathscr{L}_{dU/dT,\text{smooth}} = \frac{\sum_i \left(\frac{\partial (dU/dT)_i}{\partial SOC_i}\right)^2}{N + \varepsilon}. \tag{23}$$

Finally, a variance-floor penalty discourages collapse of the OCV surface toward a constant output,

$$\mathscr{L}_{OCV,\text{var}} = \max(0, \delta - \text{Var}(OCV)), \tag{24}$$

where δ is a small target variance threshold.

The full stage-1 loss is then written as

$$\begin{aligned}\mathscr{L}_{\text{stage1}} = {} & w_{V,\text{teacher}}\mathscr{L}_{V,\text{teacher}} + w_{T,\text{teacher}}\mathscr{L}_{T,\text{teacher}} \\ & + w_{V,\text{free}}\mathscr{L}_{V,\text{free}} + w_{T,\text{free}}\mathscr{L}_{T,\text{free}} \\ & + \lambda_{\text{mono}}\mathscr{L}_{\text{mono}} + \lambda_{R_0}\mathscr{L}_{R_0,\text{smooth}} \\ & + \lambda_{OCV}\mathscr{L}_{OCV,\text{var}} + \lambda_{dU/dT}\mathscr{L}_{dU/dT,\text{smooth}}.\end{aligned} \tag{25}$$

Stage 1 training is performed in two phases. During the warmup phase, the teacher-forced terms retain substantial weight

Copyright © 2026 by ASME

to stabilize the learned constitutive surfaces. During the fine-tuning phase, the free-rollout terms are emphasized more strongly so that the model is optimized toward longer term deployment performance. The stage-1 parameters are selected using the held-out constant-current validation set, where total iteration voltage RMSE and temperature RMSE are computed periodically. The best stage-1 weights are retained according to this validation criterion.

Once stage 1 is complete, the learned submodels $OCV(SOC,T)$, $R_0(SOC,T)$, and $dU/dT(SOC)$ are no longer trainable. Stage 2 then uses HPPC data to identify the dynamic polarization parameters R_1, C_1, R_2, and C_2.

During stage 2, the frozen submodels are evaluated using the measured SOC and temperature histories from the HPPC data, yielding the quasi-static voltage and thermal quantities required by the full deployed model. The dynamic RC states are then rolled out according to (2)–(4) using candidate values of R_1, C_1, R_2, and C_2. The resulting terminal voltage and one-step thermal predictions are compared against the measured HPPC trajectories.

The stage-2 voltage and thermal losses are defined in the same per-file masked form as in stage 1, and a weak prior is also applied to discourage excessive departure from the initial RC guesses,

$$\mathcal{L}_{RC,\text{prior}} = \frac{(R_1 - R_{1,0})^2}{R_{1,0}^2 + \varepsilon} + \frac{(C_1 - C_{1,0})^2}{C_{1,0}^2 + \varepsilon} + \frac{(R_2 - R_{2,0})^2}{R_{2,0}^2 + \varepsilon} + \frac{(C_2 - C_{2,0})^2}{C_{2,0}^2 + \varepsilon}, \tag{26}$$

where $R_{1,0}$, $C_{1,0}$, $R_{2,0}$, and $C_{2,0}$ are the chosen initial guesses.

The full stage-2 loss is therefore

$$\mathcal{L}_{\text{stage2}} = w_{V,\text{stage2}}\mathcal{L}_{V,\text{stage2}} + w_{T,\text{stage2}}\mathcal{L}_{T,\text{stage2}} + \lambda_{RC}\mathcal{L}_{RC,\text{prior}}. \tag{27}$$

Only the four scalar RC parameters are optimized during Stage 2. Once identified, these parameters remain fixed during deployment. They are optimized using Adam, and early stopping is used to terminate the procedure once the stage-2 loss ceases to improve. The best identified parameter values are then saved together with the frozen stage-1 submodels to form the final deployable electro-thermal model.

3. RESULTS

To understand model reliability, the physics constrained networks were tested across three load types. First, the model is evaluated on a 4C constant current test, to assess the fit of the more impactful learned parameters including OCV, R_0, and dU/dT, as well as model generalizability, as only rates below 3C were used in training. The RC-pairs, R_1C_1 and R_2C_2 are assessed under an HPPC profile to display model performance under dynamic load.

Displayed in figure 4, model performance across a 4C, or 12 A, constant current discharge is shown. It can be seen that the model shows good agreement above 20% SOC, despite training on a maximum rate of 9 A. A significant source of error could stem from the profile extending beyond the temperature bounds of the training data, due to the discharge rate also exceeding the

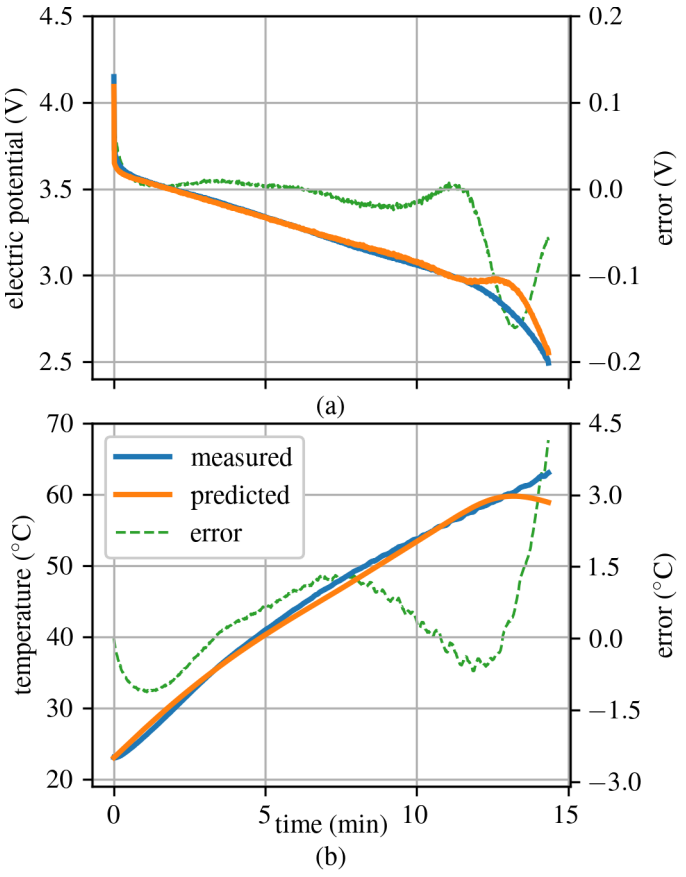


FIGURE 4: The modeled and experimental (a) voltage and (b) temperature across a 4C (12 A) constant current discharge at room temperature (~23°C ambient).

data trained on. The cell surpasses 55°C, the high temperature trained on, at ~12 min. This would align with 80% depth of discharge during a 4C constant current test, which is when the modeled voltage deviates from the experimental data. Despite the observed error at low SOC, 49.0 mV RMSE in electrical potential, and 1.02 °C RMSE in temperature is acceptable as battery behavior becomes significantly more difficult to model at high C-rates (>2C) [17]. The slight non-monotonic behavior observed near the end of discharge is attributed to extrapolation beyond portions of the training domain. As previously discussed, during the 4C discharge, the cell exceeds both the maximum current and temperature conditions observed during training, increasing uncertainty in the learned parameter surfaces and contributing to the observed voltage error at low SOC.

The model shows good electrical agreement over an HPPC profile seen in figure 5, containing a discharge and charge pulse at every 5% SOC, showcasing the dynamic capability of the model. During training, the model saw limited exposure to HPPC data, only 3 profiles, due to long training times, highlighting one limitation of the proposed training framework. However, the model shows a 65.8 mV RMSE in electrical potential, and 0.205°C RMSE in temperature, across a 1C HPPC profile at 30°C. It was assumed that all experimental data remained within the SOC range $[0,1]$. This could be a significant source of error, as experimental voltage drops below 2.5 V, signifying an

Copyright © 2026 by ASME

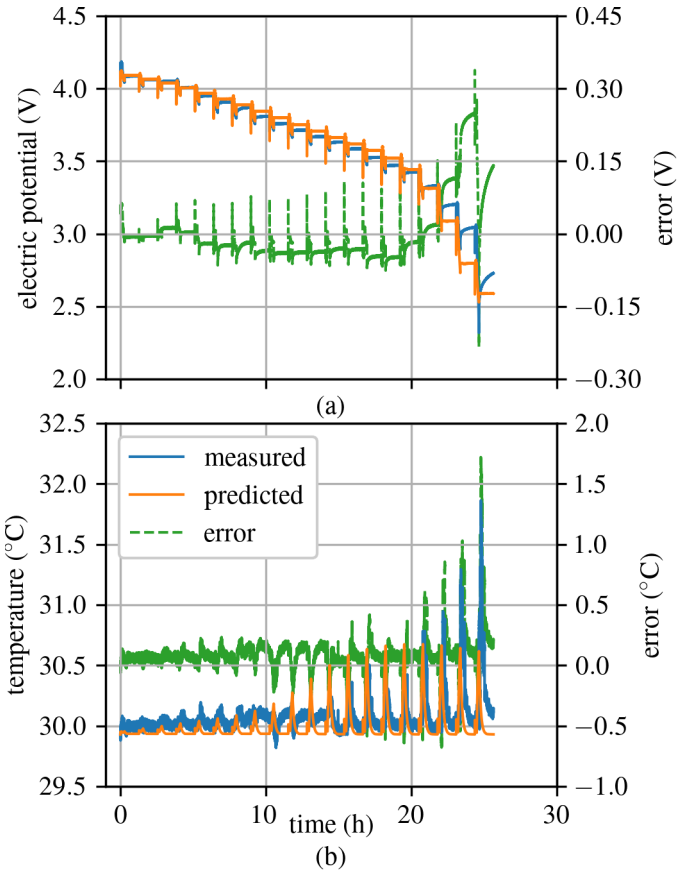


FIGURE 5: The modeled and experimental (a) voltage and (b) temperature across an HPPC profile in a 30°C ambient temperature chamber, containing 1C (3 A) pulses every 5% SOC, and C/3 (1 A) discharges between each pulse.

over discharge on the final pulse. Disparity between modeled and experimental SOC could contribute to the error in voltage noticeable below ~15% SOC. Furthermore, an over discharge would likely contribute to considerable heat generation during the final discharge, where the temperature error is the greatest.

A primary limitation of the proposed methodology lies in the separation of parameter identification across datasets and model components. While the staged approach improves observability, it implicitly assumes that these behaviors are decoupled and transferable across operating conditions. In practice, parameters such as polarization dynamics and thermal response can exhibit dependence on SOC, temperature, and even aging state. By enforcing static R_1, C_1, R_2, C_2 values and learning OCV, R_0, and dU/dT independently, the model may compensate for unmodeled coupling through the learned submodels, potentially leading to parameter bias or reduced generalization outside the training domain. Additionally, the reliance on bounded neural networks constrains outputs to physically reasonable ranges but may also limit the model's ability to capture edge-case behavior near operational limits, particularly under extreme currents or temperatures not well represented in the training data.

Another important source of error arises from the deployment structure and the underlying modeling assumptions. The lumped thermal model assumes spatially uniform temperature and simplified heat transfer via a constant hA, which neglects internal gradients and dynamic cooling effects present in real systems. Similarly, the use of coulomb counting for SOC propagation introduces cumulative error sensitivity, particularly if the initial SOC is uncertain or if current measurements contain bias. During training, the combination of one-step and total iteration losses mitigates but does not eliminate error accumulation, meaning small discrepancies in learned relationships can compound over long horizons. Finally, measurement noise, sensor calibration errors (especially in temperature), and inconsistencies between datasets (e.g., differing ambient conditions or instrumentation) can introduce discrepancies that the model may inadvertently learn as physical behavior, reducing predictive fidelity during deployment on unseen data.

4. CONCLUSION

This work presented physics-constrained ML for electro-thermal modeling of lithium-ion batteries. By embedding an equivalent-circuit representation and governing thermal equations directly into the neural network training process, the proposed framework preserves physical interpretability and computational efficiency while improving generalization beyond the training domain. Across a dynamic load profile the model shows a 65.8 mV RMSE in electrical potential, and 0.205°C °C RMSE in temperature. Model generalization was tested with a constant current load greater than any within the training dataset, where the model exhibited a 49.0 mV RMSE in electrical potential, and 1.02 °C RMSE in temperature. The two-stage identification strategy further enables robust separation of static, dynamic, and thermodynamic effects, making the model suitable for real-time applications such as digital twins and hardware-in-the-loop testing.

Future work includes model adaptation to new data as collected while under deployment through the development of an online learning framework. Additionally, effort will be made to minimize model training time and computation resources. Thus, encouraging tuning optimization to enable high fidelity performance. While the present work focuses on development of a physics-constrained machine learning framework built upon an electro-thermal ECM, future work will include systematic comparison against conventionally parameterized ECM approaches to evaluate differences in accuracy, computational cost, and generalization performance.

ACKNOWLEDGMENT

This material is based upon work supported by the Office of Naval Research (ONR) under contract numbers N00014-23-C-1012, and N00014-24-C-1301. Any opinions, findings conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the United States Navy. DISTRIBUTION STATEMENT A. Approved for public release: distribution is unlimited. Approved 06/29/2026, DCN# 2026-6-17-2410.

REFERENCES

[1] Yuan, Yupeng, Wang, Jixiang, Yan, Xinping, Shen, Boyang and Long, Teng. "A review of multi-energy hy-

Copyright © 2026 by ASME

brid power system for ships." *Renewable and Sustainable Energy Reviews* Vol. 132 (2020): p. 110081. DOI https://doi.org/10.1016/j.rser.2020.110081. URL https://www.sciencedirect.com/science/article/pii/S1364032120303725.

[2] Crespo, Elvis and Bazzi, Ali. "Mitigating Directed Energy Pulsed Load Effects Using Distributed Lithium-Ion Batteries and Load Control Models." *2025 IEEE Electric Ship Technologies Symposium (ESTS)*: pp. 275–282. 2025. DOI 10.1109/ESTS62818.2025.11152490.

[3] Brehler, Tobias, Diez, Claudius, Maisel, Sven, Nicoletti, Lorenzo, Kick, Michael, Grosu, Cristina and Lienkamp, Markus. "Comprehensive modeling of lithium-ion batteries at the system level to minimize experimental lifetime testing effort." *Journal of Energy Storage* Vol. 153 (2026): p. 120843. DOI https://doi.org/10.1016/j.est.2026.120843. URL https://www.sciencedirect.com/science/article/pii/S2352152X26005074.

[4] Farmann, Alexander and Sauer, Dirk Uwe. "A study on the dependency of the open-circuit voltage on temperature and actual aging state of lithium-ion batteries." *Journal of Power Sources* Vol. 347 (2017): pp. 1–13. DOI https://doi.org/10.1016/j.jpowsour.2017.01.098. URL https://www.sciencedirect.com/science/article/pii/S0378775317301088.

[5] Tran, Manh-Kien, Mathew, Manoj, Janhunen, Stefan, Panchal, Satyam, Raahemifar, Kaamran, Fraser, Roydon and Fowler, Michael. "A comprehensive equivalent circuit model for lithium-ion batteries, incorporating the effects of state of health, state of charge, and temperature on model parameters." *Journal of Energy Storage* Vol. 43 (2021): p. 103252. DOI https://doi.org/10.1016/j.est.2021.103252. URL https://www.sciencedirect.com/science/article/pii/S2352152X2100949X.

[6] Huang, Bo, Hu, Minghui, Chen, Lunguo, Jin, Guoqing, Liao, Shuiping, Fu, Chunyun, Wang, Dongyang and Cao, Kaibin. "A Novel Electro-Thermal Model of Lithium-Ion Batteries Using Power as the Input." *Electronics* Vol. 10 No. 22 (2021). DOI 10.3390/electronics10222753. URL https://www.mdpi.com/2079-9292/10/22/2753.

[7] Hu, Xiaosong, Li, Shengbo and Peng, Huei. "A comparative study of equivalent circuit models for Li-ion batteries." *Journal of Power Sources* Vol. 198 (2012): pp. 359–367. DOI https://doi.org/10.1016/j.jpowsour.2011.10.013. URL https://www.sciencedirect.com/science/article/pii/S0378775311019628.

[8] Wu, Yu-chen and Feng, Jun-wen. "Development and Application of Artificial Neural Network." *Wireless Personal Communications* Vol. 102 No. 2 (2018): pp. 1645–1656. DOI 10.1007/s11277-017-5224-x. URL https://doi.org/10.1007/s11277-017-5224-x.

[9] Zhao, Jingyuan, Feng, Xuning, Pang, Quanquan, Fowler, Michael, Lian, Yubo, Ouyang, Minggao and Burke, Andrew F. "Battery safety: Machine learning-based prognostics." *Progress in Energy and Combustion Science* Vol. 102 (2024): p. 101142. DOI https://doi.org/10.1016/j.pecs.2023.101142. URL https://www.sciencedirect.com/science/article/pii/S0360128523000722.

[10] Yu, Hanqing, Zhang, Zhengjie, Yang, Kaiyi, Zhang, Lisheng, Wang, Wentao, Yang, Shichun, Li, Junfu and Liu, Xinhua. "Physics-informed ensemble deep learning framework for improving state of charge estimation of lithium-ion batteries." *Journal of Energy Storage* Vol. 73 (2023): p. 108915. DOI https://doi.org/10.1016/j.est.2023.108915. URL https://www.sciencedirect.com/science/article/pii/S2352152X23023137.

[11] Peskar, Jarrett, Priddy, Braden, Sado, Kerry, Downey, Austin R. J., Booth, Kristen and Khan, Jamil. "Adaptive Agent-Based Control for Lithium-Ion Batteries in Naval Microgrids." *Volume 3A: 50th Design Automation Conference (DAC)*. 2024. American Society of Mechanical Engineers. DOI 10.1115/detc2024-143181. URL http://dx.doi.org/10.1115/detc2024-143181.

[12] Sun, Jinghua and Kainz, Josef. "Optimization of hybrid pulse power characterization profile for equivalent circuit model parameter identification of Li-ion battery based on Taguchi method." *Journal of Energy Storage* Vol. 70 (2023): p. 108034. DOI https://doi.org/10.1016/j.est.2023.108034. URL https://www.sciencedirect.com/science/article/pii/S2352152X23014317.

[13] Kampker, Achim, Späth, Benedikt, Song, Xiaoxuan and Wang, Datao. "Modelling of Battery Energy Storage Systems Under Real-World Applications and Conditions." *Batteries* Vol. 11 No. 11 (2025). DOI 10.3390/batteries11110392. URL https://www.mdpi.com/2313-0105/11/11/392.

[14] Madden, Connor, Peskar, Jarrett, Sado, Kerry, Downey, Austin R.J. and Khan, Jamil. "Electro-Thermal Hardware-in-the-Loop Battery Emulator for Shipboard Systems Testing." *2025 IEEE Electric Ship Technologies Symposium (ESTS)*: p. 47–52. 2025. IEEE. DOI 10.1109/ests62818.2025.11152433. URL http://dx.doi.org/10.1109/ests62818.2025.11152433.

[15] ARTS-Lab. "Dataset Electrothermal Deformation Characterization for Samsung 30Q Cell." GitHub. URL https://github.com/ARTS-Laboratory/dataset-electrothermal-deformation-characterization-for-samsung-30Q-cell. Accessed Jan 1st, 2026.

[16] Madden, Connor, Peskar, Jarrett, Downey, Austin, Sado, Kerry and Khan, Jamil. "Deep Neural Network-Based Modeling of Electro-Thermal Lithium-Ion Batteries Responses Leveraging Hybrid Pulse Power Characterization." Vol. 6: Energy: p. V006T09A009. 2025. DOI 10.1115/IMECE2025-168112. URL https://doi.org/10.1115/IMECE2025-168112.

[17] Li, Junfu, Wang, Dafang and Pecht, Michael. "An electrochemical model for high C-rate conditions in lithium-ion batteries." *Journal of Power Sources* Vol. 436 (2019): p. 226885. DOI https://doi.org/10.1016/j.jpowsour.2019.226885. URL https://www.sciencedirect.com/science/article/pii/S037877531930878X.

Copyright © 2026 by ASME