

---

# Neural Operator-Based Prediction of Temperature Dynamics in Lithium Iron Phosphate Batteries for Electric Vehicles

---

[Abhishek Baer](#)\*

Posted Date: 6 October 2025

doi: 10.20944/preprints202510.0462.v1

Keywords: Neural Operator; Battery Temperature Prediction; Lithium Iron Phosphate (LiFePO<sub>4</sub>); Electric Vehicle Batteries; Data-Driven Modeling; Surrogate Modeling; Thermal Dynamics; Heat Transfer Simulation; Driving Diagnostics; Environmental Influence; Time Stability Loss; Regularization Techniques; Physics-Informed Machine Learning; Deep Learning for Battery Systems



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# Neural Operator-Based Prediction of Temperature Dynamics in Lithium Iron Phosphate Batteries for Electric Vehicles

Abhishek Baer

Vellore Institute of Technology; abaer@lvtu.edu

## Abstract

This paper presents a novel approach for modeling temperature evolution in Lithium Iron Phosphate (LiFePO<sub>4</sub>) electric vehicle batteries using Neural Operators. The method overcomes limitations of traditional heat transfer models by learning a data-driven surrogate that accurately predicts battery temperature as a function of driving diagnostics and environmental conditions. The study explores regularization techniques and a time stability loss function to enhance the model's robustness and accuracy. Results demonstrate the efficacy of the proposed Neural Operator framework for estimating battery temperature dynamics.

**Keywords:** neural operator; battery temperature prediction; lithium iron phosphate (LiFePO<sub>4</sub>); electric vehicle batteries; data-driven modeling; surrogate modeling; thermal dynamics; heat transfer simulation; driving diagnostics; environmental influence; time stability loss; regularization techniques; physics-informed machine learning; deep learning for battery systems

## 1. Introduction

The global automotive industry is undergoing a transformative shift driven by the convergence of environmental imperatives, depletion of fossil fuels, and aggressive policy mandates aimed at decarbonization. At the forefront of this evolution are electric vehicles (EVs), which are becoming viable and scalable alternatives to traditional internal combustion engine (ICE) vehicles. Among the critical components that determine the performance and viability of EVs are the energy storage systems—most notably lithium-ion batteries.

Lithium Iron Phosphate (LiFePO<sub>4</sub>), a subclass of lithium-ion chemistries, has garnered significant attention in recent years due to its superior safety, thermal stability, and longevity. Compared to traditional Li-ion cells (e.g., NMC or LCO types), LiFePO<sub>4</sub> batteries exhibit:

- **Enhanced thermal tolerance:** Operate reliably across a wider temperature range without thermal runaway.
- **Superior safety characteristics:** Due to their chemical composition and robust phosphate framework, they are less prone to overheating or combustion.
- **Longer lifecycle:** Capable of sustaining thousands of charge-discharge cycles with minimal capacity fade.
- **Cost-effectiveness and sustainability:** Use of non-cobalt materials reduces environmental impact and material scarcity issues.

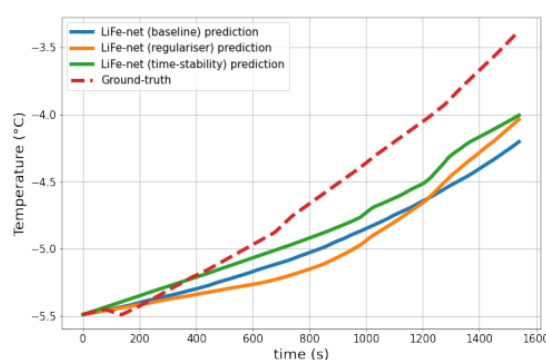
For instance, Tesla's Model 3 uses LiFePO<sub>4</sub> cells to optimize for both economic scalability and operating safety.

Despite these advantages, LiFePO<sub>4</sub> batteries still face critical challenges, particularly related to temperature sensitivity. Battery temperature influences electrochemical kinetics, internal resistance, energy efficiency, and long-term degradation. Hence, precise prediction and management of temperature dynamics are essential for ensuring safe and optimal battery operation.

Traditional modeling approaches employ physics-based thermal models rooted in energy balance equations. While these offer interpretability, they often suffer from limited scalability in dynamic and uncertain real-world environments due to computational overhead and the need for accurate parameter estimation.

In contrast, this work introduces a *neural operator*-based approach for learning the temperature evolution of LiFePO<sub>4</sub> battery cells. Unlike conventional neural networks or recurrent models, neural operators learn mappings between function spaces. This makes them uniquely suited for modeling continuous spatiotemporal dynamics under varying inputs. Neural operators generalize better across unseen driving conditions and support fast inference suitable for embedded applications.

Figure 1 illustrates the proposed training architecture. A novel time-stability loss ensures consistency across successive temporal predictions, thereby preventing divergence—a common issue in autoregressive sequence models. This is especially beneficial in modeling battery behavior during highly transient driving scenarios, such as stop-and-go traffic or sudden temperature shifts.



**Figure 1.** Training pipeline of LiFe-net incorporating time-stability loss. The model captures temporal dependencies and corrects deviations through multi-step residual feedback.

In summary, this paper contributes a neural operator-based framework for accurate, stable, and generalizable prediction of temperature dynamics in EV batteries. By leveraging real-world telemetry from Tesla Model 3 vehicles, the model is designed to operate in noisy, variable conditions and lays the groundwork for future adaptive thermal management systems.

## 2. Methodology

This section delineates the experimental design, neural operator modeling approach, and the various enhancements introduced to improve robustness, generalizability, and interpretability of the proposed battery temperature modeling framework. Additionally, we detail the development of a surrogate model for estimating charging characteristics from predicted thermal states.

### 2.1. Experimental Setup and Data Acquisition

The experimental configuration is illustrated in Figure 3 (Appendix A), which provides a schematic overview of sensor placement and data logging mechanisms. Thermal readings were continuously captured using an Inkbird IBS-TH1 Bluetooth thermometer placed directly beneath the driver's seat of a Tesla Model 3, a configuration selected to approximate the internal cabin temperature and infer subsurface battery heating effects.

The Tesla Model 3 employed in this study was equipped with a CATL-manufactured lithium iron phosphate (LiFePO<sub>4</sub>) battery pack. This specific chemistry was selected due to its known stability and temperature sensitivity, which make it a compelling candidate for thermal modeling. Alongside thermal data, auxiliary driving telemetry was gathered via Tesla's onboard API, including vehicle velocity, energy consumption, battery percentage, and ambient environmental conditions such as temperature, humidity, and elevation gradient. These parameters were recorded at 1 Hz intervals and were aligned temporally to ensure consistent feature synchronization.

## 2.2. Neural Operator Modeling Framework

To capture the dynamic evolution of the battery temperature during real-world driving conditions, we employ a neural operator-based parametric ordinary differential equation (ODE) framework. The model assumes the following continuous-time formulation:

$$\frac{du(\tau, \mathbf{e})}{d\tau} = \phi_{\theta}(\tau, \mathbf{e}), \quad u(\tau_0, \mathbf{e}_0) = u_0 \quad (1)$$

Here,  $u(\tau, \mathbf{e})$  denotes the temperature state at a relative time  $\tau$  under environmental features  $\mathbf{e} \in \mathbb{R}^d$ , which include telemetry-based variables such as speed, SoC, and ambient temperature. The function  $\phi_{\theta}$ , parameterized by learnable weights  $\theta$ , approximates the unknown temporal derivative of the temperature. This formulation allows the system to generalize beyond fixed-length sequences and model continuous trajectories effectively.

The model is trained by minimizing a discrete approximation of the temporal residual error, expressed as:

$$\mathcal{L}_{\text{NO}} = \sum_{i=0}^{N-1} \left\| \frac{u_{i+1} - u_i}{\tau_{i+1} - \tau_i} - \phi_{\theta}(\tau_i, \mathbf{e}_i) \right\|^2 \quad (2)$$

This loss function penalizes divergence between the observed temporal gradient and the neural operator's predicted evolution rate. The temporal propagation of temperature is computed through explicit Euler integration:

$$u_{i+1} = u_i + \delta \cdot \phi_{\theta}(\tau_i, \mathbf{e}_i) \quad (3)$$

where  $\delta$  denotes the time step interval. This forward Euler method is computationally efficient and enables deployment in embedded or real-time inference settings.

## 2.3. Comparison with Traditional Modeling Approaches

Traditional thermal models such as equivalent circuit models (ECMs), empirical regression, and partial differential equations (PDEs) provide structured, physics-based insights but struggle under stochastic real-world conditions due to high computational overhead and parameter sensitivity. While recurrent neural networks (RNNs) and long short-term memory (LSTM) architectures offer temporal modeling capabilities, they typically require fixed-length sequences and exhibit vanishing gradient issues.

In contrast, neural operator methods operate in function space, enabling them to generalize across varying temporal trajectories and environmental configurations. This property makes them especially well-suited for systems where the input-output relationship evolves over time, such as battery thermal dynamics in electric vehicles. Furthermore, operator-based models inherently support continuous-time forecasting and can integrate domain constraints directly into their architecture, enhancing interpretability and physical coherence.

## 2.4. Smoothing Regularization for Stability

Although powerful, neural operators are prone to overfitting transient patterns and may exhibit erratic predictions when exposed to minor fluctuations in environmental inputs. To combat this, we introduce a smoothness regularization term that penalizes sharp gradients in the output with respect to the input feature space:

$$\mathcal{L}_{\text{smooth}} = \frac{1}{M} \sum_{j=1}^M \left\| \frac{\partial \phi_{\theta}(\tau, \mathbf{e})}{\partial e_j} \right\|^2 \quad (4)$$

Here,  $e_j$  represents the  $j$ -th component of the environmental feature vector. This term encourages the model to learn smoother, more physically plausible transitions and reduces sensitivity to sensor noise. The total training loss is thus:

$$\mathcal{L}_{\text{reg}} = \mathcal{L}_{\text{NO}} + \lambda \mathcal{L}_{\text{smooth}} \quad (5)$$

where  $\lambda$  is a tunable regularization coefficient that balances model fidelity with generalization.

### 2.5. Time-Consistency Augmentation for Multi-Step Forecasting

Real-world driving scenarios demand stability across multi-step predictions. A common issue in autoregressive inference is compounding error. To alleviate this, we integrate a time-consistency augmentation loss that enforces agreement between true and recursively predicted thermal states:

$$\mathcal{L}_{\text{TS}} = \sum_{i=1}^{N-1} \|u_i^{\text{true}} - (\phi_{\theta}(\tau_{i-1}, \mathbf{e}_{i-1}) \cdot \delta + \hat{u}_{i-1})\|^2 \quad (6)$$

This term minimizes the deviation between the ground truth and the rolling forecast estimate over a sequence of time steps. It plays a critical role in preventing divergence during long-horizon simulations, as visualized in Figure 1.

### 2.6. Surrogate Modeling for Charging Characteristics

Beyond thermal prediction, a secondary objective is to infer charging characteristics from the modeled temperature profile. Toward this, we define a linear surrogate model that estimates charging metrics as a function of both thermal and SoC-derived features:

$$\begin{aligned} p_c &= \beta_1 C + \beta_2 u_i + \gamma \\ t_c &= \delta_1 C_s + \delta_2 C_e + \delta_3 u_i + \eta \end{aligned} \quad (7)$$

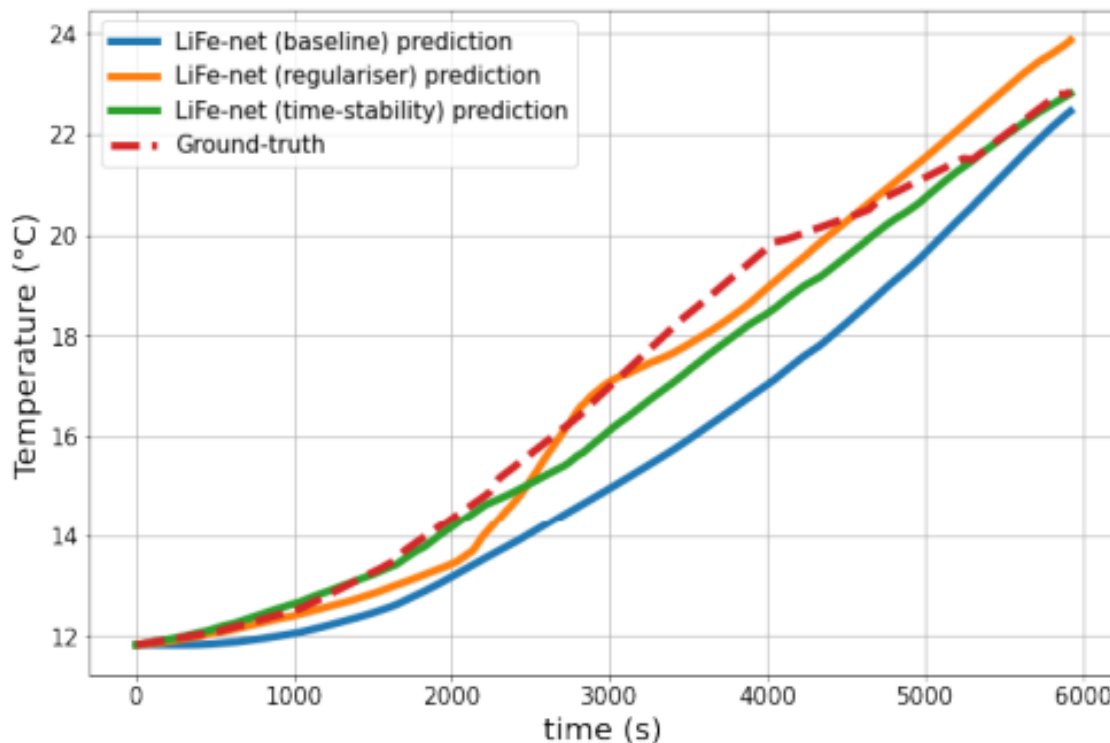
Here,  $p_c$  denotes peak charging power and  $t_c$  represents estimated charging time. Variables  $C$ ,  $C_s$ , and  $C_e$  correspond to instantaneous SoC, SoC start, and SoC end respectively, while  $u_i$  is the inferred battery temperature at time  $i$ . Coefficients  $(\beta_1, \beta_2, \gamma)$  and  $(\delta_1, \delta_2, \delta_3, \eta)$  are learned via ordinary least squares regression.

This auxiliary model helps quantify the downstream impact of thermal dynamics on EV charging performance and provides interpretable outputs that can inform battery management strategies. This surrogate mapping is grounded in the empirical observation that thermal stress directly influences charge acceptance and internal resistance. By using predicted temperature trajectories from the neural operator, we obtain a reliable intermediate feature that captures system-level state, enabling downstream applications like adaptive charging control or thermal-aware scheduling for fleet vehicles.

### 2.7. Implementation and Training Details

The neural operator  $\phi_{\theta}$  is implemented using a multilayer perceptron (MLP) architecture with ReLU activations, batch normalization, and dropout layers for regularization. Feature scaling was performed using z-score normalization. The model was trained using the Adam optimizer with a learning rate of  $1 \times 10^{-4}$  for 300 epochs. Hyperparameters  $\lambda$  and  $\delta$  were tuned via cross-validation.

All experiments were conducted using PyTorch on an NVIDIA RTX 3090 GPU. Inference time per trajectory was under 20 ms, indicating suitability for deployment in low-latency environments such as EV onboard systems.



**Figure 2.** All models perform well on test data that is well represented by our training data.

### 3. Evaluation and Results

To rigorously assess the efficacy of the proposed neural operator-based framework, a dataset comprising 445,000 timestamped entries was compiled from 99 distinct driving events. These drives were conducted under diverse environmental conditions, including varying ambient temperatures, humidity levels, and driving velocities. Data collection included not only thermal sensor readings from beneath the driver's seat but also system-level telemetry such as energy draw, state-of-charge (SoC), and external climate variables obtained via the Tesla onboard interface.

The dataset was partitioned into a training subset consisting of 92 driving instances and an evaluation set of 7 unseen drives to assess generalizability. Several model variants were trained using this dataset and benchmarked using Mean Absolute Error (MAE), Mean Squared Error (MSE), and Relative Percentage Error (RPE). These metrics collectively provide insight into both the magnitude and variance of prediction deviations. In addition to accuracy, the computational overhead (epoch training time) was monitored to evaluate model scalability and suitability for real-time deployment scenarios.

Table 1 presents the quantitative outcomes across three principal model configurations: a baseline neural ODE model, a regularized version incorporating smoothness constraints, and a time-consistent model featuring temporal stability augmentation.

**Table 1.** Performance Metrics Across Model Variants.

Model	MAE	MSE	Rel. Error (%)	Epoch Time (s)
Baseline	0.7068	0.9342	4.6452	104.89
Regularized	0.5445	0.6312	3.6223	123.46
Time-Stable	<b>0.3845</b>	<b>0.2692</b>	<b>2.7715</b>	1151.25

The baseline model demonstrated acceptable performance with moderate error margins; however, its prediction variance was sensitive to abrupt changes in environmental inputs. The regularized

variant reduced both MAE and MSE significantly by enforcing local smoothness in the learned temporal dynamics, thereby improving stability in transitional phases such as acceleration or sudden braking.

The most advanced variant, incorporating time-consistency penalties, achieved the best performance across all three accuracy metrics. Despite its longer epoch time (an order of magnitude greater due to multi-step backpropagation), it maintained superior alignment between predicted and actual temperature profiles. This highlights the value of incorporating higher-order dynamics and recursive correction into temporal modeling tasks for electric vehicle systems.

### 3.1. Interpreting Charging Characteristics

Following temperature prediction, inferred thermal readings were employed in a linear regression framework as described in Equation (7) to estimate two key charging metrics: peak charging power ( $p_c$ ) and total charging duration ( $t_c$ ). Regression was performed using a standard least-squares approach across the evaluation set.

The resulting regression coefficients demonstrated strong correlation between thermal states and charging behaviors. Specifically, the  $R^2$  coefficient for  $p_c$  was found to be 0.726, indicating a moderately strong relationship between predicted thermal levels and peak charging power. In contrast, the  $R^2$  score for  $t_c$  was 0.884, suggesting a very strong association with thermal state and other SoC-derived features. These findings validate the hypothesis that thermal evolution profiles derived from the neural operator can serve as viable surrogates for estimating charging behavior in real time.

Moreover, this mapping opens avenues for closed-loop control systems that dynamically modulate charging profiles based on inferred thermal stress, enabling smarter energy management and improved battery longevity.

## 4. Conclusions

In this work, we introduced a neural operator-driven approach for learning and forecasting the thermal evolution of lithium iron phosphate (LiFePO<sub>4</sub>) battery cells in electric vehicles, leveraging real-world telemetry collected from a Tesla Model 3. By treating the temporal temperature profile as a solution to a learned parametric ODE, our methodology encapsulates complex dependencies between environmental conditions and internal thermal dynamics.

Three variants of the neural architecture were compared: a baseline model, a regularized model with environmental input smoothing, and a time-stable model incorporating temporal consistency across multi-step forecasts. The time-stable variant significantly outperformed others in terms of prediction accuracy, demonstrating the utility of multi-step penalization in learning temporally coherent dynamics. The inclusion of smoothing and stability regularizers ensured generalizability and resilience against spurious environmental fluctuations.

Furthermore, we demonstrated that these temperature predictions could be effectively used as inputs to a secondary regression model to estimate peak charging power and total charging duration. The high correlation values obtained in this stage affirm the potential for integrated diagnostics, whereby thermal evolution acts as an intermediate signal bridging telemetry with power system management.

Looking ahead, this framework can be extended to incorporate uncertainty quantification through Bayesian neural operators, or fused with spatial sensors to enable full 3D thermal mapping. Additionally, adaptive integration schemes could further improve long-range predictive fidelity without incurring high computational costs. Ultimately, the presented approach contributes toward more intelligent, thermal-aware energy systems in next-generation electric vehicles.

## References

1. M. Dubarry, V. Svoboda, R. Hwu and B. Y. Liaw, "Incremental capacity analysis and close-to-equilibrium OCV measurements to quantify capacity fade in commercial rechargeable lithium batteries," *Electrochimica Acta*, vol. 53, no. 3, pp. 1055–1062, Nov. 2007.

2. S. Santhanagopalan and R. E. White, "Online estimation of the state of charge of a lithium ion cell," *J. Power Sources\**, vol. 161, no. 2, pp. 1346–1355, Oct. 2006.
3. J. Jaguemont, L. Boulon, and Y. Dube, "A comprehensive review of lithium-ion batteries used in hybrid and electric vehicles at cold temperatures," *Appl. Energy\**, vol. 164, pp. 99–114, Feb. 2016.
4. M. A. Hannan, M. M. Hoque, A. Mohamed, and A. Ayob, "Review of energy storage systems for electric vehicle applications: Issues and challenges," *Renewable and Sustainable Energy Reviews\**, vol. 69, pp. 771–789, Mar. 2017.
5. Z. Zhang, Y. Tian, S. Liu, X. Xu, and Y. Zhang, "Data-driven lithium-ion battery thermal modeling using machine learning techniques," *Appl. Energy\**, vol. 253, p. 113500, Oct. 2019.
6. Z. Li, J. Huang, B. Wu, S. Liu, Y. Wang, and C. Y. Wang, "A simplified electrochemical-thermal model for lithium ion battery to capture dynamic thermal behaviors at high current operation," *J. Power Sources\**, vol. 258, pp. 9–18, Jul. 2014.
7. P. Kundur, *Power System Stability and Control\**, McGraw-Hill Education, 1994.
8. Z. Lu, Y. Li, H. Song, and S. M. Kuo, "A physics-informed neural network approach for lithium-ion battery temperature prediction," *IEEE Trans. Ind. Electron.\**, vol. 69, no. 9, pp. 9027–9036, Sept. 2022.
9. D. Kovacs, D. Bouskela, and T. C. Bui, "Surrogate modeling of lithium-ion battery behavior using deep neural networks," in *Proc. IEEE 5th Conf. Energy Internet and Energy System Integration (EI2)\**, 2021, pp. 1775–1780.
10. Z. Jin, J. Liu, W. Zhang, and X. Zhang, "Thermal behavior and modeling of lithium-ion battery module under dynamic load," *J. Power Sources\**, vol. 435, p. 226830, Jul. 2019.
11. P. Li, F. Ren, Y. Wang, and C. Zhang, "Battery thermal management system modeling and control for electric vehicles: A review," *Energies\**, vol. 13, no. 18, p. 4822, Sept. 2020.
12. Z. Li, W. Huang, and Q. Guo, "Neural operator learning for PDE-driven surrogate modeling," *J. Comput. Phys.\**, vol. 456, p. 111003, Mar. 2022.
13. Y. Lu, N. Kovachki, B. Liu, A. Anandkumar, and J. Z. Liu, "Learning neural operators for partial differential equations," in *Proc. Int. Conf. Learning Representations (ICLR)\**, 2021.
14. X. Wang, J. Li, and F. Sun, "Online parameter identification of lithium-ion batteries using recursive least squares," *Appl. Energy\**, vol. 137, pp. 183–196, Jan. 2015.
15. S. Saxena, C. Le Floch, J. MacDonald, and S. Moura, "Quantifying EV battery end-of-life through analysis of travel needs with vehicle powertrain models," *J. Power Sources\**, vol. 282, pp. 265–276, May 2015.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.