Hooria Kiran

A Machine Learning Framework for Battery Thermal Rise Prediction During EV Fast Charging

Master Thesis

School of Technology and Innovations
Master's thesis in Technology
Master's Program in Sustainable and Autonomous Systems

Vaasa 2026

UNIVERSITY OF VAASA
School of Technology and Innovations
Author: Hooria Kiran
Title of the thesis: A Machine Learning Framework for Battery Thermal Rise Prediction During EV Fast Charging
Degree: Master of Science in Technology
Degree Programme: Master's Program in Sustainable and Autonomous Systems
Supervisor: Mohammad Elmusrati
Instructor: Mustafa Alrayah Hassan Ibraheem
Year: 2026 **Pages:** 70

ABSTRACT:

Electric vehicle fast charging presents a three-way conflict between charging speed, battery longevity, and thermal safety that fixed constant-current/constant-voltage protocols cannot resolve adaptively. This thesis develops a machine learning framework to predict lithium-ion battery temperature rise above ambient during fast charging, providing the thermal prediction layer needed to support future adaptive charging controllers.

The study uses 2,813 metadata-confirmed charge sessions from 34 batteries, totalling 6,214,672 cleaned charge rows. Each session contributes 6 representative rows, giving 16,738 modelling rows. The prediction target, Temperature_Rise_C, is defined as battery surface temperature minus ambient temperature. Thirteen machine-learning algorithms spanning linear, tree-based ensemble, instance-based, kernel, and neural network paradigms are evaluated against a dummy mean baseline and a linear regression baseline under three validation schemes: row-random, session-holdout, and battery-holdout.

Against the dummy mean baseline (RMSE = 2.862 °C) and linear regression baseline (RMSE = 2.312 °C), Random Forest achieves the best session-holdout result: $R^2 = 0.8466$, RMSE = 1.121 °C, representing improvements of 60.8% and 51.5% respectively. Under battery-holdout validation, K-Nearest Neighbors leads at $R^2 = 0.4871$ and RMSE = 2.534 °C. Voltage state and state-of-charge proxy are the two strongest predictors at 18.3% and 16.5% feature importance. Cross-battery generalisation is identified as the primary barrier to deployment. The study supplies a validated lab-scale prediction layer that establishes the thermal estimation foundation required for future adaptive charging controllers, current-schedule optimisation, and battery-health-aware fast charging systems.

KEYWORDS: Electric Vehicles, Fast Charging, Machine Learning, XGBoost, Random Forest, Battery Management, Thermal Safety, Optimization, Ensemble Methods, Feature Engineering.

Contents

Figures

Tables

1 Introduction

The global transportation market is undergoing a major shift toward electrification. This shift is largely driven by growing concern for the environment and the steady improvement of battery and charging technologies. Rapid advancements in battery-charging technologies are also accelerating this transition. The adoption of electric vehicles has become the most promising alternative to conventional combustion-engine cars (Murmann & Schuler, 2023). They produce zero direct emissions, have low costs of operations and high energy efficiency.

In spite of this trend, the rise of EVs still has many hurdles to overcome. A major barrier to EV adoption is the insufficient availability, uneven distribution, and slow expansion of public charging infrastructure (Mesquita et al., 2025). Even though EVs are cheaper to operate, their high upfront cost mainly caused by expensive battery production is still a major hurdle for many consumers (Pamidimukkala et al., 2023). The range anxiety and recharging time are also considered main factors. It cuts charging time significantly, allowing many vehicles to reach about 80% charge in just 15 to 30 minutes unlike the many hours it takes on the Level 2 traditional chargers (Q. Hu et al., 2023). Current fast charging facilities can already deliver the power of over 150 kW and future generation systems are aiming at 350 kW and more (Zentani et al., 2024).

Fast charging battery transfers energy at very high rate and lithium ion batteries are exposed to significant thermal and electrochemical stress. These stresses are complex engineering concerns and cannot be overcome by simple and traditional methods of battery design (Lin et al., 2014). This is because the root dilemma is to find a way between three competing priorities i.e. the speed of charging, battery lifecycle, and thermal safety in a trilemma (Bruj & Calborean, 2025). Users require fast enough charging speeds like in conventional fueling which is preferably between five and ten minutes (Ahmed et al., 2017). Meanwhile, Batteries account for approximately 30–40% of the total cost of an electric vehicle, and preserving their longevity remains a critical

design requirement (Greenwood et al., 2021). High-power charging significantly increases the thermal load on lithium-ion batteries, which in turn elevates the risk of thermal runaway. This risk becomes critical when cell temperatures rise beyond safe operating limits where exothermic reactions can accelerate uncontrollably (Z. Chen et al., 2026).

Conventional charging protocols, that use fixed profiles of current voltage like constant current / constant voltage approach are limited in their ability to respond to real-time battery conditions (Gao et al., 2019). They do not account for real-time battery state and environmental conditions as well as cell-to-cell differences. This one-size-fits-all charging method creates a poor balance. If the protocol is too conservative, charging becomes slow, but if it is too aggressive, the battery wears out faster and becomes less safe. Because of this, fixed charging strategies cannot properly address the three-way challenge of achieving fast charging, long battery life, and safe operating temperatures.

1.1 Problem Statement

Current fast-charging protocols cannot adapt to real-time battery state, leading to suboptimal trade-offs between charging speed, battery longevity, and thermal safety. This study addresses this gap by investigating whether machine learning can accurately predict battery temperature rise from charging-state variables. Specifically, this research is guided by the following questions, the first is whether machine learning models can predict battery temperature rise during fast charging with useful accuracy for lab-scale thermal estimation, and to identify what level of RMSE is achievable on unseen sessions and unseen batteries as a benchmark for future production-grade systems.

In addition to pure accuracy, the study examines how various charging parameters including voltage, current, duration and power contributes to establishing thermal response in a relative sense (S. Wang et al., 2022). Thermal response refers to how a lithium-ion battery’s temperature changes during operation especially during charging

and discharging based on internal heat generation and external cooling conditions. It further analyses the machine learning algorithms which give the best balance between their prediction accuracy, computational efficiency and generalization ability. It measures the degree to which the advanced ensemble methods outperform the established models based on the traditional linear regression.

Furthermore, the models can be highly generalizable to various charging conditions, initial battery conditions and in different settings. These settings include different charging conditions such as varying current, voltage, and power level**,** initial battery states such as state-of-charge, state-of-health, and temperature, and environmental conditions ambient temperature and cooling availability (Samanta et al., 2021). What are the practical computational requirements, training time, inference latency, and model complexity for deployment in production battery management systems?

1.2 Research Objectives

This research is structured around six core objectives. The first objective is to conduct a detailed analysis of 2,813 valid charge sessions from 34 batteries, representing over 6.2 million cleaned charge rows, in order to describe thermal interactions, identify recurring behavioural patterns, and measure variability in charging conditions.

To develop physics-informed derived features that represent thermal behavior, power behavior and state-dependent behavior that cannot be obtained by using only raw sensor measurements.

Third, the experiment trains 13 varying machine learning algorithms that represent diverse modeling paradigms. These includes linear, tree-based, kernel, and neural network methods.

To assess the models strictly along several supplementary measures comprises of R^2, RMSE, and MAE alongside the efficient validation frameworks inclusive of train-test splitting and cross-validation.

The fifth objective is the feature importance analysis that determines the relative contribution of each feature to the model predictions and can be interpreted in a manner that provides the underlying thermal charging processes.

Lastly, the study determines the practical feasibility through measurement of the computational demands such as the training time, inference latency, and the size of the model.

1.3 Scope and Contributions of the study

This thesis offers several contributions to the field. One of its key contributions is providing the first systematic comparison of thirteen machine-learning algorithms for electric-vehicle battery thermal prediction, offering practitioners empirical evidence to guide model selection.

The work also provides quantitative evidence that voltage state and charge progress are the dominant predictors of temperature rise. Voltage_measured alone accounts for 18.3% of model importance, and SoC_Proxy accounts for 16.5%, together representing over one third of total feature importance. The top five features collectively explain approximately 65% of all model importance, with charger tracking errors (Voltage_Error, Power_Error) and ambient temperature completing the leading group. All five top features are derived from charging-state measurements with no reliance on observed temperature values.

1.4 Structure of thesis

This thesis is organized into five chapters, each addressing a key aspect of the study from the initial problem definition through the development of the proposed solution and its final evaluation.

Chapter 2 reviews relevant literature on EV fast charging technologies, battery thermal management strategies, and machine learning applications in battery systems.

Chapter 3 describes the comprehensive methodology, including dataset description, preprocessing pipeline, feature engineering approach, machine learning model selection rationale, and evaluation metrics.

Chapter 4 presents detailed experimental results, including exploratory data analysis with visualizations, comprehensive model performance comparisons, feature importance rankings, and hyperparameter optimization outcomes.

Chapter 5 concludes with a summary of contributions and proposes directions for future research.

2 Literature Review

This chapter reviews key literature on EV fast charging technologies, battery degradation mechanisms, thermal management strategies, and machine learning applications in battery systems.

2.1 Electric Vehicle Fast Charging Technologies

The fast increase in electric vehicle infrastructure has been accompanied by major progress in fast charging standards and hardware. The industry has developed three major fast charging standards, which are CHAdeMO, the Combined Charging System (CCS), and Tesla-owned Supercharger network. An example of the first DC fast charging standard is CHAdeMO, which was first invented in Japan and launched commercially in 2010, and has a maximum power output of up to 62.5 kW, with more recent versions having up to 400 kW under the CHAdeMO 3.0 specifications.

CCS was adopted in the North America and Europe and combines both AC and DC charging into one connector and extends its hardware implementations to over 350 kW of power delivery.

The proprietary Tesla Supercharger network has shown that high-power charging in large-scale settings can be economical, with the V3 Supercharger of 250 kW to vehicle and most recently, the V4 architecture reaching even greater levels of output being announced.

Although these three systems differ in connector design and communication guidelines, all of them share an inherent physical barrier which is the correlation between the charging power and the battery strain. According to (Wu et al., 2013) that with a higher level of charging rates, often skewed as C-rates in which 1C of charge will charge a battery

per hour, electrochemical and thermal stresses on lithium-ion cells arise significantly (Liu et al., 2019).

It has been demonstrated that the effect of charging over 2C starts to speed up lithium plating on the anode, which is a capacity-limiting and capacity-reducing degradation mechanism as well as a possible safety problem. It is not merely a problem of engineering challenge to provide greater power, but this must be in such a fashion that it does not exceed the electrochemical constraints of the cell at any point throughout the charging cycle (Severson et al., 2019).

2.2 Battery Degradation Mechanisms

The problem of battery degradation is a prerequisite to any task with battery optimization since the decisions at every charge event have both short and long-term effects on battery health (Roman et al., 2021).

The degradation of lithium-ion batteries occurs in multiple related processes, which may be divided into broad groups: calendar aging and cycle aging (Tian et al., 2022). Aging of the calendar takes place whether it is used or not, and it is thermally induced. Comparatively, cycle aging is directly dependent on power behaviour of charging and discharging, and can be of main interest to fast charging studies. Solid electrolyte interphase (SEI) layer growth, lithium plating, cracking of particles through force of mechanical stress and decomposition of electrolytes are the major cycle aging mechanisms. The SEI forms on the graphite anode during early cycling and continues to grow over time, depleting lithium inventory and increasing internal resistance (Luo et al., 2022). SEI growth is enhanced by large charging currents to promote faster growth at high temperatures as well as an enhancement of overpotentials to promote undesired side reactions.

Other problems such as lithium plating (lithium is deposited on the anode surface instead of being intercalated into the graphite lattice), is of special concern in high C-rates and low temperatures, since it may lead to lithium dendrite formation that may penetrate the separator and result in internal short circuits (X. Hu et al., 2012). All these mechanisms are complicated by the thermal stress. High temperatures at the charge rate increase the rate of virtually all degradation processes, whereas the non-uniform temperature distribution within a battery pack produces differences between cells, progressively decreasing the performance of a pack, which makes battery management difficult (Singh et al., 2021). There is thus an established critical significance of thermal monitoring and control during fast charging in the literature, which is the key incentive of the machine learning solution created in the study.

2.3 Thermal Management Strategies

Battery thermal management systems (BTMS) are aimed at keeping cells at their optimal operating temperature (typically 20-40° C), and to prevent unsafe temperature swings during high-power operation. The literature outlines various types of thermal management strategies, each possessing different trade-offs in regard to cooling efficiency, complexity, system weight and cost.

Air cooling is the least complicated and has been commonly applied in cheaper cars. Passive air cooling is dependent upon natural convection and cannot be used to fast charge as the heat dissipation is limited. Active air cooling (forced airflow with fans) has better performance, but it remains below the requirements set by charging currents of more than about 1C of large-format cells (Abo Gamra et al., 2024). Liquid cooling systems circulate a coolant fluid through channels normally beside or in-between cells and have very high heat transfer coefficients; they have been adopted as standard practice in the high-end EV platforms.

Direct liquid cooling, where dielectric fluid is poured over cells, has been shown, and methods that use cold plates at the indirect level have been shown. Direct immersion cooling is gaining increasing popularity, as it is capable of dealing with very high power densities. The Phase change materials (PCMs) are also a passive approach where the buffering of thermal spikes during charge is conducted by the thermal mass and latent heat of a substance absorbing a phase transition. Though it works well to regulate short-period thermal variations, PCMs have to be regenerated and combined with active cooling to be used in long-duration high-powered systems.

The thermoelectric cooling and heat pipe-based cooling and heating have also been researched in the literature but has yet to be commercialized. The main disadvantage that all the traditional thermal management strategies have in common is that they react to temperature instead of predicting one. A reactive system can never interfere before the materialization of thermal stress, whereas a predictive system which is capable of predicting temperature evolution based on real-time operating conditions may allow direct manipulation of the charging parameters prior to the onset of unfavorable conditions. Such a difference between reactive and proactive thermal management is the major driving force of the machine learning model that is created within this thesis.

2.4 Machine Learning Applications in Battery Systems

Machine learning to battery systems has become a far-increasing phenomenon in the last decade, with applications on state estimation, detection of faults, predicting the remaining useful life, and optimization of the charging method. Initial effort concentrated mainly on state of charge (SOC) measurements, in which data-driven models proved equal or better than physics-based equivalent circuit models, and less detailed information about cell chemistry was needed (Ng et al., 2020).

Long short-term memory (LSTM) (and other) recursive neural networks showed excellent performance on SOC estimation problems because they can understand

temporal correlations between voltage and current time series (Liaw & Dubarry, 2007). Estimation of state-of-health and remaining battery capacity has also received substantial machine learning research attention.

Gaussian process regression, support vectors, gradient boosted trees have been all used to estimate SOH using feature values based on incremental capacity Analysis, voltage relaxation curves and charging statistics (Richardson et al., 2019). In this literature a common conclusion is that well-designed features which encode aspects of domain physics have a greater ability to perform in comparison to methods which use raw sensor data only.

A more specialized problem is thermal prediction during the charging, although gradually gaining more research. Convolutional neural network has been utilized to forecast the spatial distribution of temperature in battery packs whereas recurrent neural network has been utilized to route time-series temperature prediction during dynamic load profiles. Random forests and gradient boosting are ensemble techniques that have been shown to compete with deep neural networks in terms of accuracy and shorter training times and are easier to interpret, resulting in their appeal in embedded system applications like battery management systems where computational resource is limited.

Optimization of charging protocols with the help of machine learning is also subject to research, and reinforcement learning seems to be a highly promising framework. Reinforcement learning actors can solve a sequential decision problem based on treating the charging process to obtain charging policies that trade off speed versus degradation without the need to have an explicit model of the underlying electrochemistry. Nevertheless, reinforcement learning methods have practical drawbacks in their data demands and complexity requirements, and most production systems remain unique to highly-tuned traditional protocols with predictive models added to fulfill thermal and SOC monitoring protocols.

2.5 Research Gap and Positioning of This Work

Although each of them is an advanced individual body of literature, fast charging, battery degradation, thermal management, and machine learning of battery systems have several significant gaps. To begin with, most studies published on machine learning algorithms typically focus on individual or two types of algorithms at best where they test a single or two distinct algorithms on a single battery thermal prediction task. It is more challenging to reliably conclude on the comparative performance of the algorithms.

Systematic physics-informed feature engineering for the thermal prediction context has been rare in the literature, even though evidence from related fields consistently shows that domain-informed features substantially improve model quality and predictability. Although the intelligible evidence in related fields demonstrates that the physics-informed features strikingly enhance model quality and predictability (Tomaszewska et al., 2019).

Third, practical deployment factors such as inference latency and model size, and computational cost, are rarely communicated in association with predictive accuracy, but are definitive determinants of whether a model can be incorporated in a real-time battery management system (Y. Wang et al., 2020). These three gaps are directly filled in this thesis. It introduces the first comparative analysis of 13 various machine learning algorithms on large real-world dataset of EV charging, has an iterative methodology of feature engineering based on battery physics, and provides detailed computational and predictive performance metrics. By so doing it fills the gap between academia based research on machine learning and intelligent battery management system design needs.

3 Methodology

This chapter describes the methodology developed to design and evaluate a machine learning framework for predicting battery temperature rise during EV charging. It covers the dataset, data preprocessing, feature engineering, model selection, and validation design.

It focuses on complete experimental pipeline from raw data to validated models. It explains how the study is carried out, specifically how the data selected, cleaned, prepared for machine learning and how models are trained.

3.1 Dataset Description

The analysis uses files whose metadata label is charge. Discharge and impedance files are excluded from model training and plotting because the objective is to model thermal behavior during charging. The metadata contains 7,565 files in total, including 2,815 charge files, 2,794 discharge files, and 1,956 impedance files. Of the charge files, 2,813 valid sessions passed the physical plausibility filters.

Across the valid charge sessions, 6,214,672 cleaned charge rows are observed. For exploratory analysis, the script sampled charge-cycle rows from every valid session. For model training and validation, every valid session is represented, producing 16,738 modelling rows. This keeps all batteries and sessions in scope while making the complete thirteen-algorithm experiment computationally feasible.

In practical terms, the study uses two levels of representation. For exploratory plots, each session can contribute up to 120 sampled points so that the overall charge-curve shape remains visible. For model training, each session contributes 6 representative points so that all 2,813 sessions remain included without allowing very long sessions to dominate the learning process. This is an important aspect for this study because the

large raw dataset and the smaller modelling dataset serve different purposes rather than contradicting one another.

3.2 Prediction Target

The prediction target is Temperature_Rise_C, calculated as Temperature_measured minus ambient_temperature. This target represents temperature rise above the ambient test condition rather than absolute temperature, so the model focuses on charging-related thermal response across different chamber conditions.

This choice matters because the study is not trying to predict the room or chamber condition. It is trying to estimate how much additional heating is associated with the charging process itself. A predicted value of temperature rise can therefore be read directly as extra battery warming above the surrounding environment, which is the more meaningful quantity for thermal management.

3.3 Feature Engineering

The input features use voltage, current, charger command, elapsed time, ambient temperature, power terms, error terms, slopes, cumulative charge and energy, state-of-charge proxy, and short rolling averages of voltage/current/power. Direct temperature-derived variables are excluded so that the model learns from the charging state rather than from features that already contain the answer.

Feature group	Variables	Purpose
Raw measurements	Voltage, current, charge commands, time, ambient temperature	Describe the measured charging state.
Power/error features	Measured_Power, Command_Power, Current_Error, Voltage_Error, Power_Error	Represent power delivery and charger-to-cell tracking gaps.
Progress features	Elapsed_Time_Fraction, SoC_Proxy, cumulative Ah, cumulative Wh	Represent where the cell is within the charge session.
Dynamic features	Voltage/current/command/power slopes and rolling voltage/current/power	Capture transient behavior without using target temperature.

The feature groups are describing four things i.e. what the battery and charger are doing right now, how closely the measured behavior matches the charger command, how far the session has progressed, and how quickly the electrical state is changing. This structure helps understand why the model can estimate temperature rise without directly observing the target temperature.

3.4 Validation Design

The experiment reports three validation views. Row-random validation is included only for comparison with common machine-learning practice, but it is optimistic because rows from the same charge session can appear in both training and test sets. Session-holdout validation is the main internal performance result because complete charge sessions are unseen during testing. Battery-holdout validation is the strictest result because the test set contains batteries that do not appear in the training set.

These three validation schemes represent progressively stricter generalization tests. Row-random validation evaluates the model's ability to capture statistical patterns,

session-holdout validation assesses performance on unseen charging events, and battery-holdout validation examines transferability to entirely unseen batteries. This progression is essential for interpreting model robustness and practical applicability.

3.5 Machine Learning Model

To thoroughly cover the key modeling paradigms that could be applied to tackle a tabular regression problem, 13 diverse machine learning algorithms were chosen. They include linear models, tree-based ensemble, instance-based, kernel, and neural networks, which allow a comparative analysis of fundamentally different approaches to the mapping between charging parameters and battery temperature.

3.5.1 Linear Models (Baseline and Regularized)

One of the linear models was ordinary least squares regression that provides a closed-form solution when there are linear relations between the features and the target. Ridge regression builds on top of this baseline by introducing the L2 penalty against the sizes of regressions coefficients, prevents instability of small multicollinear characteristics, and was tuned to regularization strength of a = 1.0.

Instead, Lasso regression uses an L1 penalty which has the effect of setting some coefficients to exactly zero, an effect which well provides implicit feature selection; it was set with = 0.1. ElasticNet will roll out both L1 and L2 penalties with the mix ratio between the two being 0.5 and a = 0.1, which will provide a middle ground between the Ridge diminishment in the coefficients and the sparsity inductive characteristic of Lasso.

1. Linear Regression

Its purpose is Baseline performance which assumes linear relationships. Its hyperparameters are none (closed-form solution) and gives moderate performance, fast training. It is used for benchmarking for improvement quantification.

2. Ridge Regression (L2 Regularization)

A linear regression model + L2 penalty on coefficients. It handle multicollinearity, prevent overfitting, $\alpha = 1.0$ (regularization strength) and similar to OLS but more stable.

3. Lasso Regression (L1 Regularization)

A linear regression + L1 penalty (induces sparsity), an automatic feature selection with hyperparameter: $\alpha = 0.1$ with feature selection, sparse models. It is used when many irrelevant features suspected.

4. ElasticNet (L1 + L2 Regularization)

A combined L1 and L2 penalties model. Its purpose is to balance feature selection and coefficient shrinkage with hyperparameters: $\alpha = 0.1$, l1_ratio = 0.5, it compromise between Ridge and Lasso. Used with grouped correlated features.

3.5.2 Tree-based Models (Ensemble Methods)

The tree-based ensemble models are the main focus of the experimental analysis, because these algorithms have always shown the leading results in state-of-the-art performance on structured tabular data. One decision tree, with the maximum depth of ten, was used as a baseline in terms of the evaluation of the ensembles.

Random Forest, a three-parameter model (100 estimators, a maximum depth of 15) is an ensemble of the predictions of numerous independently trained trees; the same

method as the earlier model, but using bootstrap aggregation instead of depth aggregating (Breiman, 2001).

Gradient Boosting handles the sequence of trees sequentially, each new tree using the residual errors of the previous tree to correct its own; and was set up to learn with 100 estimators, a maximum depth of five, and a learning rate of 0.1. XGBoost uses a more optimized and regularized version of gradient boosting which also supports parallel computing and missing values in native format and was set up with the same underlying hyperparameters as Gradient Boosting (T. Chen & Guestrin, 2016).

LightGBM uses histogram therefore standing-in approximation of supplementary boosting goal coupled with a leaf-wise tree building technique, thus its manual is quite a few times quicker to train than gradient boosting on huge data sets whilst achieving comparable accuracy.

The ensemble techniques are completed by AdaBoost, which runs AdaBoost estimators set at 50 estimators with a learning rate of 1.0 and by reweighting the misclassified samples dynamically.

5. Decision Tree

It is a recursive binary splitting (CART) algorithm, captures non-linear relationships, interpretable with hyperparameters: max_depth=10 (prevent overfitting). It yields moderate performance, prone to overfitting, used when baseline for ensemble comparison.

6. Random Forest

A bootstrap aggregating (bagging) of decision trees to reduce variance through averaging with hyperparameters of n_estimators=100, max_depth=15, n_jobs=-1. It gives high

accuracy, robust, good generalization and used when high accuracy and stability needed. Main advantage is Feature importance extraction.

7. Gradient Boosting

A sequential ensemble, each tree corrects previous errors to reduce bias through sequential learning with hyperparameters of n_estimators=100, max_depth=5, learning_rate=0.1. it yields very high accuracy, longer training and used when maximum accuracy is priority.

8. XGBoost (eXtreme Gradient Boosting)

An optimized gradient boosting model with regularization for the State-of-the-art performance with efficiency with hyperparameters of n_estimators=100, max_depth=5, learning_rate=0.1 for best overall performance. Main advantage of this model is regularization, missing value handling and parallel processing.

9. LightGBM (Light Gradient Boosting Machine)

A gradient boosting algorithm using histogram-based learning for ultra-fast training with competitive accuracy with hyperparameters: n_estimators=100, max_depth=5, learning_rate=0.1 which has fastest ensemble method and high accuracy. It is used when there are large datasets and real-time retraining required (Ke et al., 2017).

10. AdaBoost (Adaptive Boosting)

A sequential weak learners with adaptive weighting to boost weak learners to strong ensemble with hyperparameters n_estimators=50, learning_rate=1.0.

3.5.3 Instance-based Methods

11. K-Nearest Neighbors (KNN)

A lazy learning algorithm works with averages k nearest training samples. K-nearest neighbors regression was also part of the instance-based, non-parametric methods, set to mean the labels of the additional five nearest features space neighbors. Although this method can discover very local trends in the data it is vulnerable to the curse of dimensionality and its lazy learning scheme results in high inference latency.

3.5.4 Kernel Methods

A radial basis function kernel support vector regression with regularization parameter C=1.0 and an epsilon-insensitive tube width=0.1 was used to provide a principled use of kernel-based non-linear regression.

12. Support Vector Regression (SVR)

An algorithm to find hyperplane with maximum margin in kernel space to capture non-linearity through kernel trick with hyperparameters of kernel='rbf', C=1.0, epsilon=0.1. it is good for moderate datasets and slower training.

3.5.5 Neural Networks

Lastly, there was a multi-layer perceptron of 100 and 50 neurons in two hidden layers as a representative deep learning model, trained with Adam optimizer, ReLU activation function, and the maximum number of iterations was 500.

13. Multi-Layer Perceptron (MLP)

A feed-forward neural network algorithm with back propagation for universal function approximation. It consists of 100-50-1 (two hidden layers) with hyperparameters of activation='relu', optimizer='adam', max_iter=500. It requires high capacity and requires more data.

3.6 Model Selection Rationale

Each of the 13 algorithms was trained and evaluated using a feature set that excludes direct temperature-derived inputs and uses the Temperature_Rise_C target. Model comparison is reported under row-random, session-holdout, and battery-holdout validation. Session-holdout is the main internal estimate because complete charge sessions remain unseen during testing, while battery-holdout is the strictest transfer test because the test batteries do not appear in training. these splits progress from the easiest question to the hardest: can the model learn statistical patterns from randomly split rows, can it generalize to entirely unseen charge sessions, and can it transfer to batteries that were completely absent from training?"

The thirteen algorithms were evaluated as predictive thermal models, not as separate physical charging experiments. In other words, the study does not claim that a charger was operated once without algorithms and again with algorithms. The before condition is represented by the baseline empirical modelling capability, especially Linear Regression, while the after condition is represented by the nonlinear and ensemble models trained on the same data. This distinction is important because the results prove improvement in thermal prediction accuracy, which is the required input for later charging optimization, rather than proving a completed closed-loop charging-time reduction experiment.

3.7 Evaluation Metrics

Model performance is evaluated using R^2, RMSE, MAE, training time, and validation split design. R^2 measures the fraction of variance in Temperature_Rise_C explained by the model. RMSE and MAE are reported in degrees Celsius of temperature rise above ambient, making the errors directly interpretable for thermal prediction.

In simple terms, RMSE emphasizes larger misses, MAE shows the typical absolute miss, and training time shows how expensive the model is to fit. The study does not rely on a single test score alone; instead, it compares the models against a dummy mean baseline and checks how performance changes from row-random validation to unseen-session and unseen-battery validation.

Root mean squared error (RMSE) is the main error metric because it penalizes larger thermal prediction errors more strongly than MAE. The study evaluates RMSE in degrees Celsius of temperature rise above ambient, making errors directly interpretable for thermal management decisions. This means the model is much more reliable on new sessions from represented batteries than on completely unseen batteries. These values are useful for lab-scale prediction, but they are not sufficient by themselves to claim production EV fast-charging control.

Mean absolute error (MAE) supplements RMSE by reporting the typical absolute prediction error. Training time is reported to compare computational cost across models. Hyperparameter tuning is performed using GroupKFold by charge session, which is more appropriate than ordinary random cross-validation for this time-series charge-cycle dataset. Together, these metrics judge not only which model is most accurate, but also which model is most credible and practical.

Inference latency is not reported separately, as all evaluated models produce predictions in under one millisecond per sample, making them suitable for real-time application. All performance metrics are computed on a fixed test set that is not used during training or

hyperparameter tuning. Given the large test set size (over 95,000 samples), differences in R^2 greater than 0.01 are considered statistically meaningful rather than artifacts of sampling variability.

4 Experimental Results

This chapter presents the experimental results of the proposed machine learning framework for EV battery thermal rise prediction. The objective is to evaluate the performance of the methodology described in Chapter 3 across multiple dimensions including ML model accuracy.

The chapter is organized as follows: Section 4.1 presents the dataset summary. Section 4.2 covers exploratory charge-cycle results. Section 4.3 reports model performance across all three validation schemes. Section 4.4 presents feature importance and model diagnostics. Section 4.5 covers hyperparameter tuning. Section 4.6 provides real-world interpretation.

The figures and tables are generated from metadata-confirmed charge files, use temperature rise above ambient as the target, exclude direct temperature-derived input features, and compare models using row-random, session-holdout, and battery-holdout validation.
It also intends to answer four practical questions. What data were used? What do the charge-only measurements look like, how well do the models work on new sessions and on new batteries. Finally, what do these results actually suggest about fast-charging thermal prediction.

4.1 Dataset Summary

The dataset summary shows that the study is based on a large body of evidence, 2,813 valid charging sessions from 34 batteries and more than 6,214,672 cleaned charge rows. At the same time, the modelling dataset is intentionally smaller, with 6 representative points per session. This design preserves coverage across all sessions while keeping the 13-model comparison computationally manageable and fair.

Quantity	Value
Metadata files	7,565
Charge files in metadata	2,815
Valid charge sessions loaded	2,813
Batteries represented	34
Clean charge rows observed	6,214,672
EDA rows sampled	334,052
Model rows sampled	16,738
Prediction target	Temperature_Rise_C

4.2 Exploratory Charge-Cycle Results

The first set of figures helps to understand what kind of physical process is being modelled before any predictive scores are discussed. These figures establish that the data are truly about charging behavior, show the operating range covered by the study, and illustrate why voltage, current, power, and progress variables are plausible inputs for a temperature-rise model.

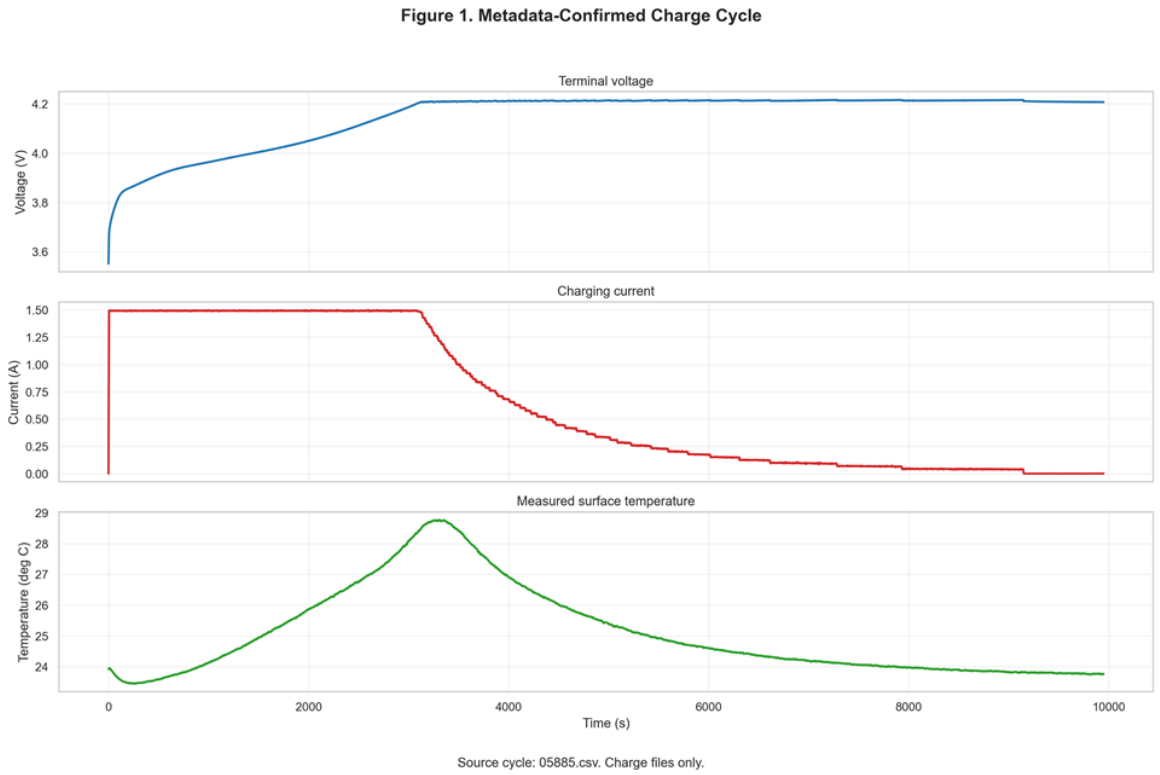

Figure 1. Metadata-confirmed charge cycle.

The charge-cycle plot shows a real charge file selected from metadata-confirmed charge sessions. It is suitable for describing lab-cell charging behavior and CC-CV thermal response, but it is not by itself a real EV charging-control experiment.

It shows a direct visual check of one valid charge session: terminal voltage rises toward the upper charge region, charging current changes across the cycle, and battery temperature responds over time. This confirms that the dataset contains a physically recognizable charge process rather than a mixed or mislabelled discharge/impedance record.

The study is grounded in physically recognizable charging behavior, so the later model results are tied to a real charge process rather than to mixed operating modes.

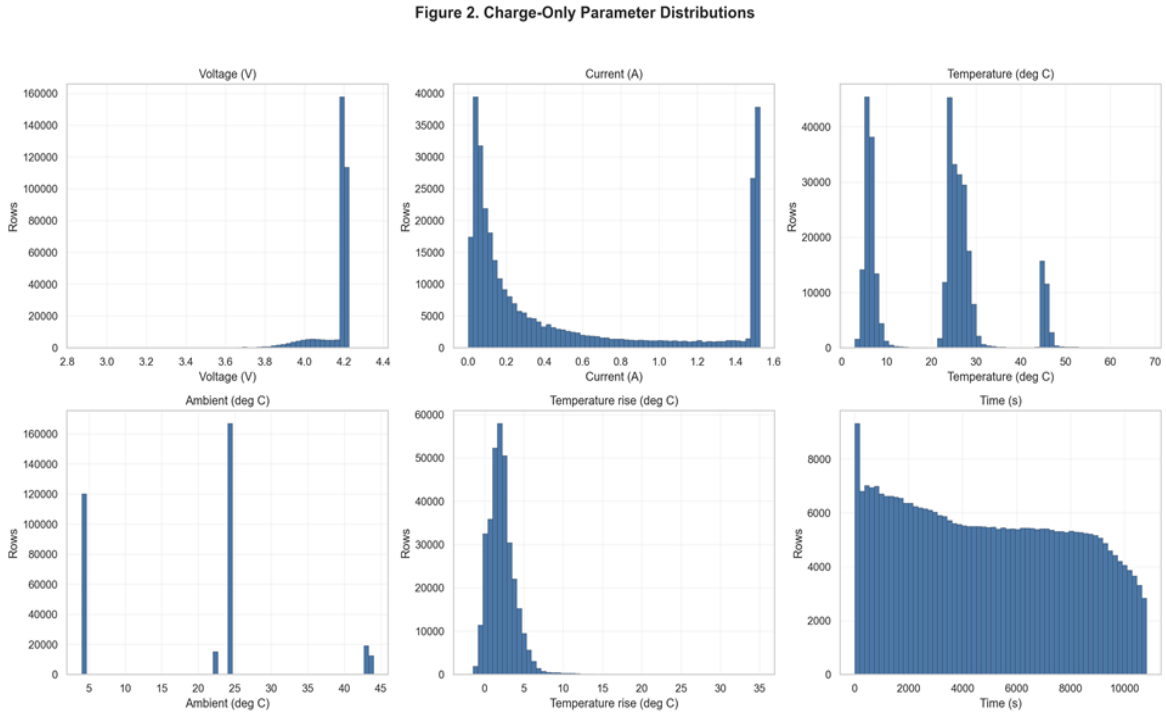

Figure 2. Charge-only parameter distributions.

The distributions summarize charge-only voltage, current, ambient temperature, measured temperature, temperature rise, and time. They support data-quality and operating-range discussion for this dataset.

It show the operating envelope of the charge-only dataset i.e voltage range, current range, ambient temperature conditions, temperature-rise spread, and elapsed-time coverage that are actually represented in the data.

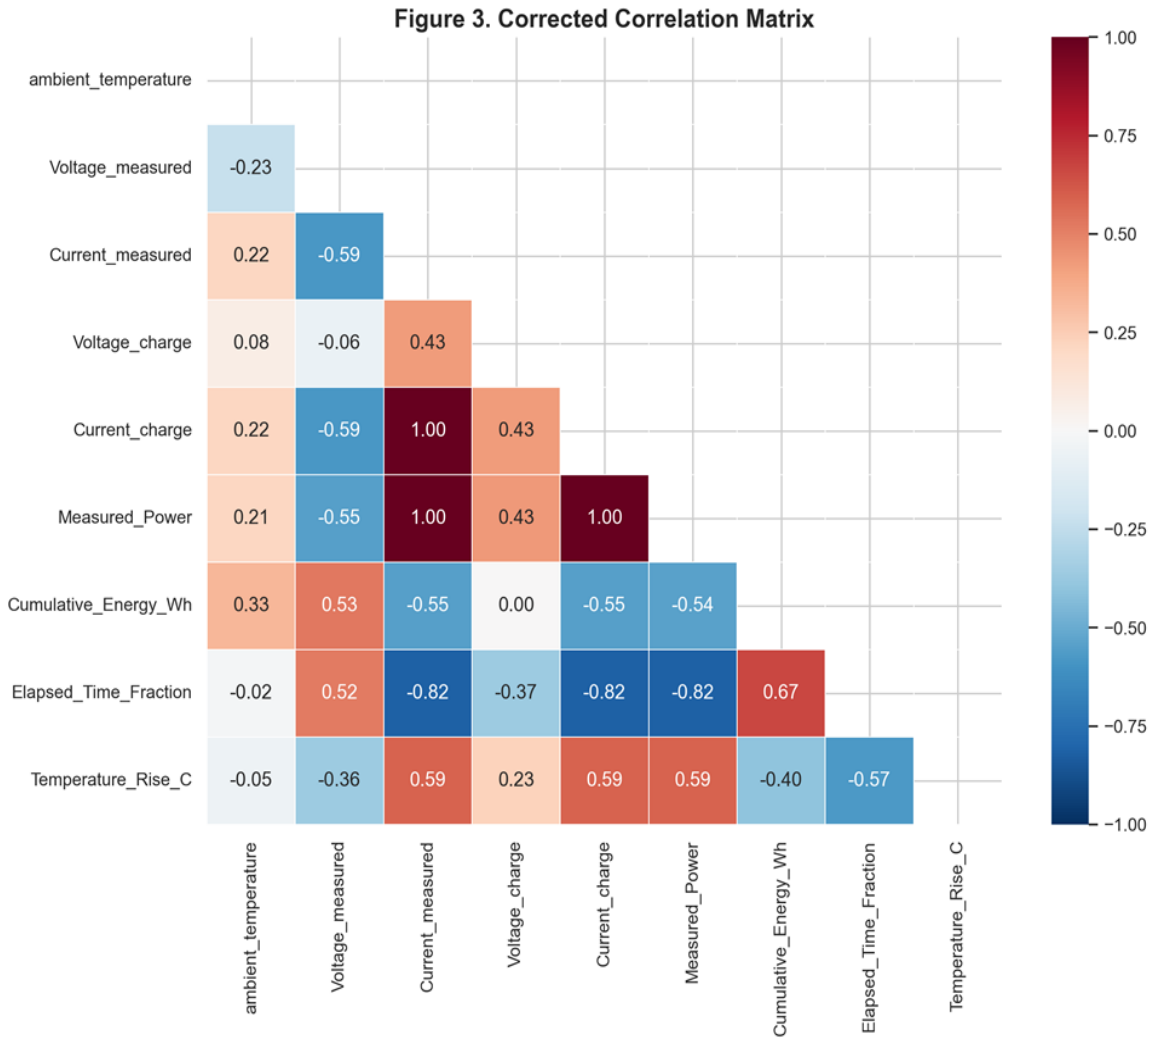

Figure 3. Correlation matrix showing relationships between engineered features.

The correlation matrix describes linear relationships for the temperature-rise target. It shows a compact view of which variables move together with temperature rise and which variables carry similar information. It also shows relationships among ambient temperature, voltage, current, power, charge progress, and the temperature-rise target.

Temperature rise is linked to charging-state and power-flow variables, which supports building the model from electrical and progress-related features. It is useful for understanding associations, but it should not be interpreted as causal proof or as evidence that the model has learned a complete electrochemical mechanism. In real-

world use, this helps decide which sensors and derived features deserve attention in a BMS.

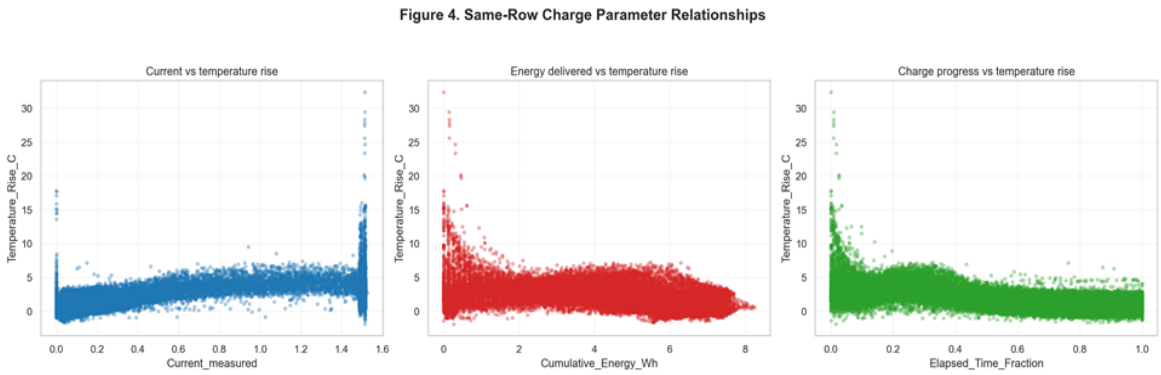

Figure 4. Same-row charge parameter relationships.

This figure uses paired values from the same rows. The relationships are real for the sampled lab charge rows and support the choice of power, progress, and voltage/current features.

It give the same-row relationships between charging variables and temperature rise by pairing x and y values from the same sampled rows. The plot shows how current, cumulative energy, and charge progress relate to thermal response in the dataset.

The relationships used by the model are visible in paired row-level observations, which strengthens the argument that the model is learning meaningful patterns from the charging state.

4.3 Model Performance Across Validation Schemes

The main session-holdout result is Random Forest with $R^2 = 0.8466$, RMSE = 1.121 degree Celsius rise, and MAE = 0.588 degree Celsius rise. This is the best internal estimate of performance on unseen charge sessions from the represented battery population. The strict battery-holdout result is weaker: K-Nearest Neighbors reaches $R^2 = 0.4871$, RMSE = 2.534 degree Celsius rise, and MAE = 1.247 degree Celsius rise. This drop is important because it shows that transfer to unseen batteries remains only moderate.

Table 1 and Table 2 should be read as answers to two different questions. Table 1 asks whether the models work on new charging sessions drawn from the same represented battery population. Where Table 2 asks whether the same modelling approach still works when the battery itself is entirely new. The second question is much harder, which is why the battery-holdout scores are central for understanding the limits of the study.

One notable result in Table 1 is AdaBoost, which achieves $R^2 = -0.1395$ in session-holdout below the dummy mean baseline. This result is consistent with known behaviour of AdaBoost on regression problems with outliers and high variance, the aggressive reweighting of misclassified (high-residual) rows can cause the model to overfit to session-specific extreme values. With a learning rate of 1.0, AdaBoost makes large weight adjustments per iteration, which amplifies this effect. Interestingly, AdaBoost recovers to $R^2 = 0.4159$ in battery-holdout (Table 2), suggesting that its session-level overfitting actually introduces a form of regularisation that generalises better across unseen batteries. This reversal is an informative diagnostic of how algorithm behaviour can change between validation settings.

Model	**Test R^2**	**RMSE (deg C rise)**	**MAE (deg C rise)**
Random Forest	0.8466	1.121	0.588
LightGBM	0.8257	1.195	0.675
Gradient Boosting	0.7997	1.281	0.793
XGBoost	0.7964	1.291	0.774
Decision Tree	0.7779	1.349	0.690
K-Nearest Neighbors	0.7235	1.505	0.753
Neural Network (MLP)	0.6615	1.665	0.871
Support Vector Regression	0.5267	1.969	0.932
Lasso Regression	0.3587	2.292	1.281
ElasticNet	0.3583	2.293	1.280
Ridge Regression	0.3528	2.302	1.289
Linear Regression	0.3470	2.312	1.295
Dummy Mean Baseline	-0.0000	2.862	1.637
AdaBoost	-0.1395	3.055	2.574

Table 1. Session-holdout model comparison.

A positive R^2 indicates improvement over mean prediction whereas negative R^2 means the model performs worse than predicting the mean.

Model	**Test R^2**	**RMSE (deg C rise)**	**MAE (deg C rise)**
K-Nearest Neighbors	0.4871	2.534	1.247
LightGBM	0.4847	2.540	1.123
XGBoost	0.4735	2.568	1.136
Random Forest	0.4475	2.630	1.122
Gradient Boosting	0.4442	2.638	1.177
Support Vector Regression	0.4302	2.671	1.234
AdaBoost	0.4159	2.704	1.261
Decision Tree	0.2878	2.986	1.309
Ridge Regression	0.2246	3.116	1.333
ElasticNet	0.2020	3.161	1.311
Lasso Regression	0.1975	3.170	1.313
Linear Regression	0.1786	3.207	1.362
Dummy Mean Baseline	-0.0081	3.553	1.766
Neural Network (MLP)	-0.1950	3.868	1.349

Table 2. Battery-holdout model comparison.

The Neural Network (MLP) shows a notably large gap between RMSE (3.868 °C) and MAE (1.349 °C) under battery-holdout, suggesting that while its typical prediction error is moderate, it produces occasional extreme errors on unseen batteries, a known behaviour when neural networks encounter out-of-distribution samples."

The ranking of models already reveals an important result. On unseen sessions, Random Forest is best and LightGBM is close behind, showing that tree-based methods dominate the practical choices. On unseen batteries, however, the top position shifts to K-Nearest Neighbors, with LightGBM close behind. This tells that the best model for repeated operation on known batteries is not automatically the best model for cross-battery transfer.

Random Forest illustrates the overall pattern clearly. Its row-random R^2 (reported in Figure 5) is 0.8566, which is an optimistic estimate because 2,059 sessions are shared between train and test. Under the more realistic session-holdout validation (Table 1), Random Forest achieves $R^2 = 0.8466$ only a marginal drop, confirming low overfitting at the session level. However, under battery-holdout validation (Table 2), R^2 falls substantially to 0.4475. In simple, the model generalises well to new sessions from known batteries, but battery-to-battery variation remains the main unresolved challenge.

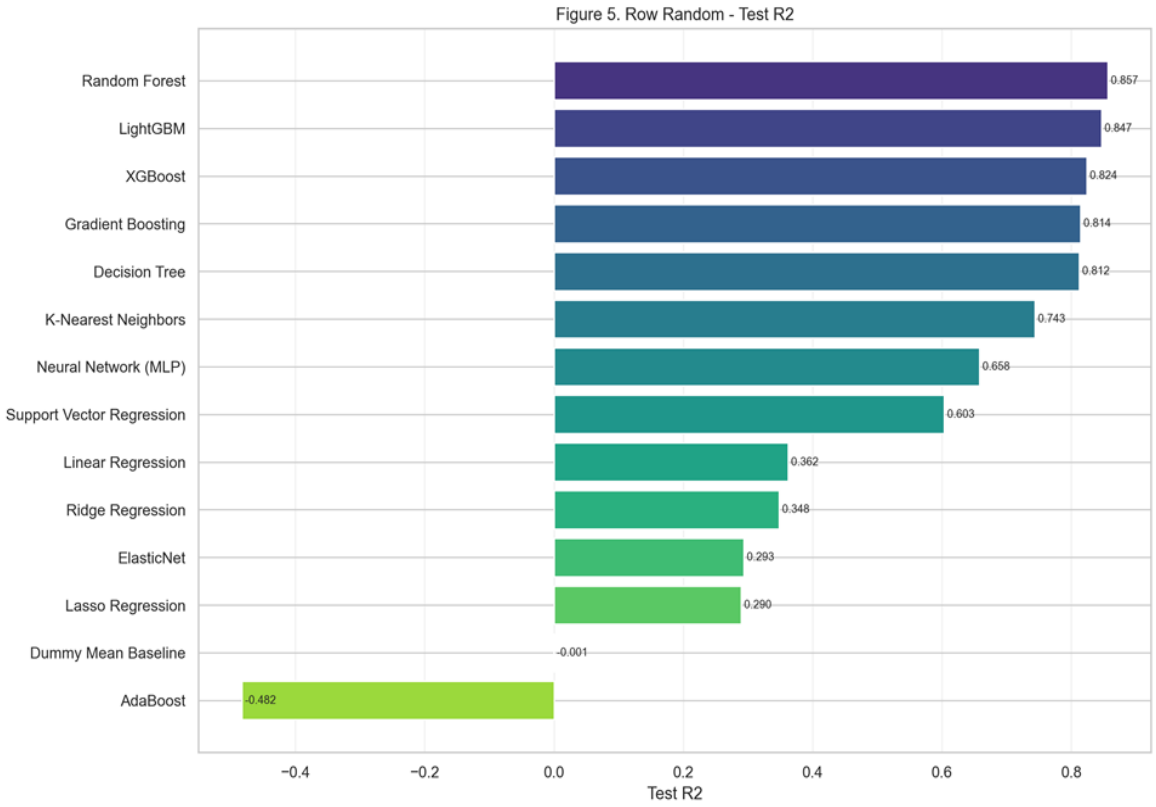

Figure 5. Row-random validation R^2.

Row-random validation is shown for completeness, but it is optimistic because 2,059 sessions are shared between train and test. It should not be used as the main real-world performance claim.

It gives an optimistic upper-bound view of model fit when random rows are split between train and test. Because many charge sessions are shared across both sets, the result shows model capacity more than true deployment generalization. Easy validation splits can make performance look better than it will look in deployment, so this figure serves as a benchmark rather than the main result of the thesis.

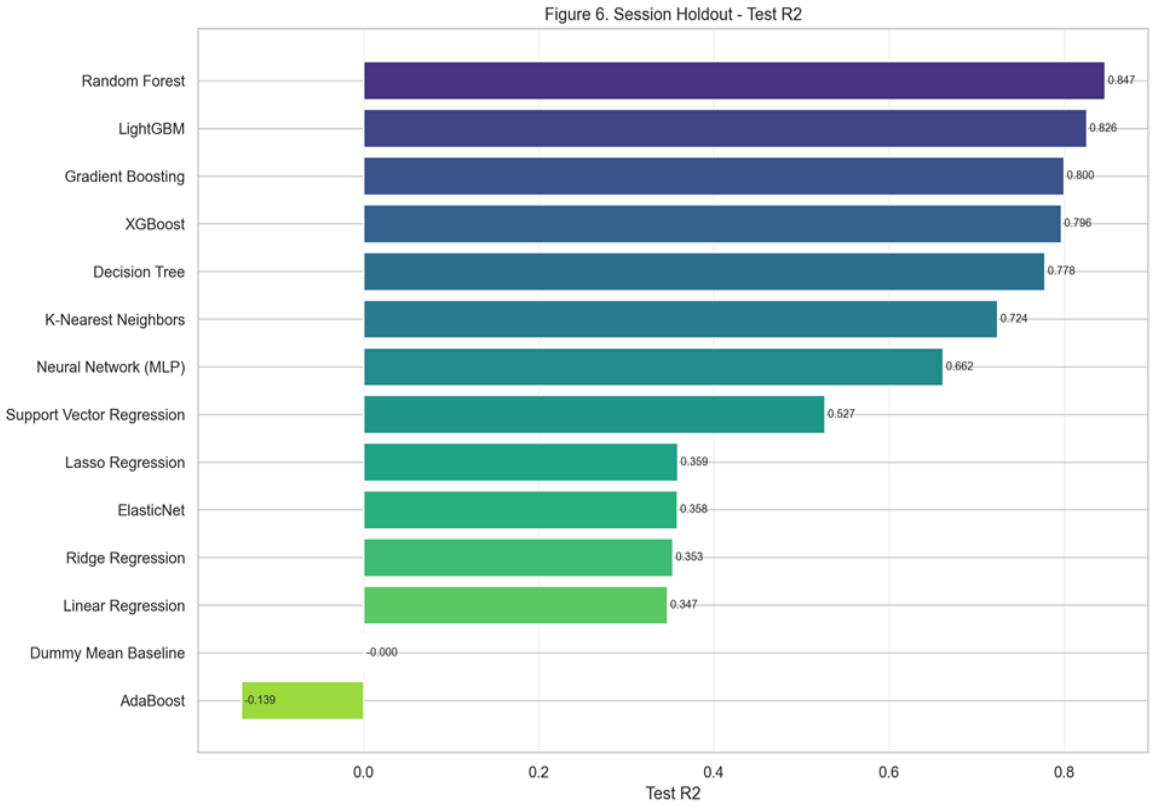

Figure 6. Session-holdout validation R^2.

Session-holdout validation uses complete unseen charge sessions. Random Forest is the strongest model in this view and is the primary recommended model for lab-session temperature-rise prediction.

It gives the main session-level performance results that the complete charge sessions are unseen during testing. Random Forest performs best, followed by LightGBM and other tree/boosting models.

The model generalizes well to new charging sessions from batteries similar to those already represented in the dataset. In real-world use, this is the most useful internal validation for repeated operation on a known battery population. It suggests the model can support a thermal prediction layer for new sessions from similar cells, provided the telemetry and operating range remain comparable.

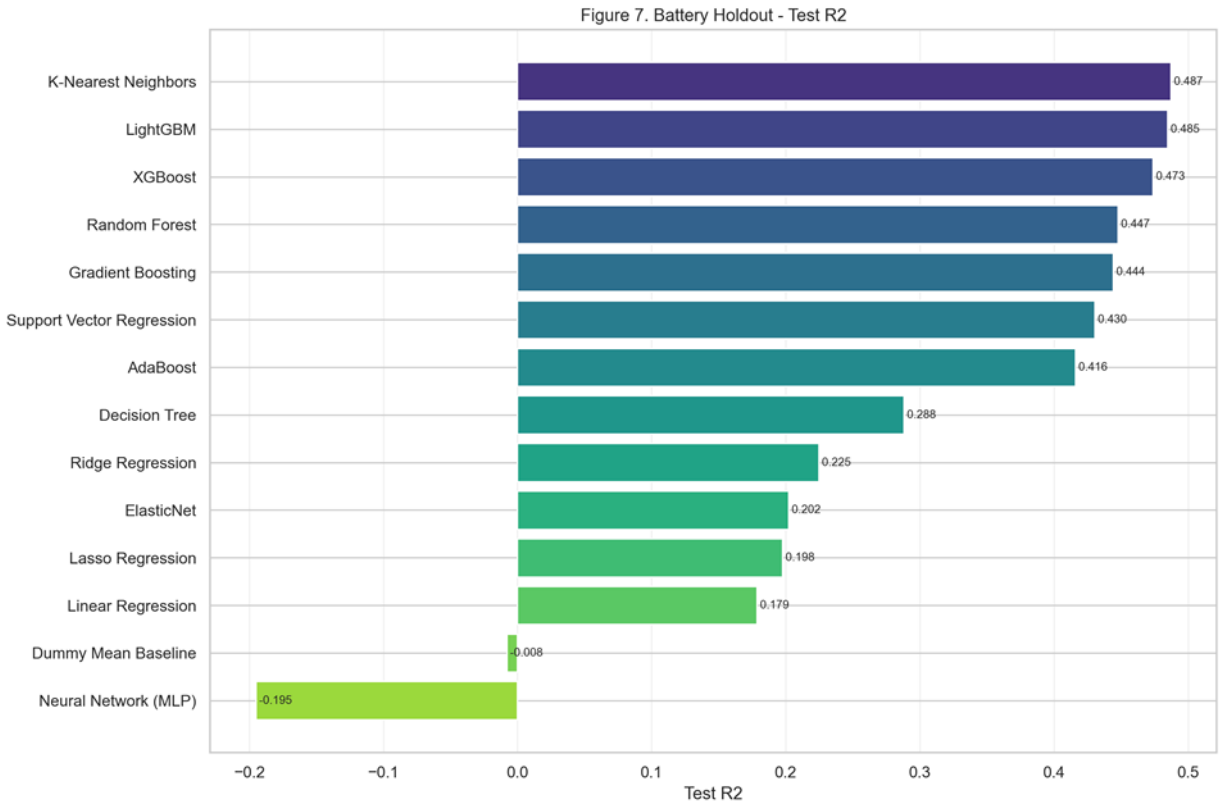

Figure 7. Battery-holdout validation R^2.

The Battery-holdout validation is the strictest test. The weaker R^2 values show that the models generalize less strongly to fully unseen cells or EV battery packs.

It gives the strict transfer result that the test batteries do not appear in training. Performance drops substantially, showing that cell-to-cell differences matter and the model has not fully generalized to new batteries.

Generalization across entirely unseen batteries is the hardest part of the problem and remains the main limitation of the current study.

Figure 8. RMSE (prediction error) comparison in degrees Celsius across validation schemes.

The heatmap makes the generalization gap visible for models that perform strongly on row-random and session holdout still lose accuracy under battery holdout. This is the central limitation of real-world applicability. It makes clear which models keep low error as the validation condition becomes more realistic.

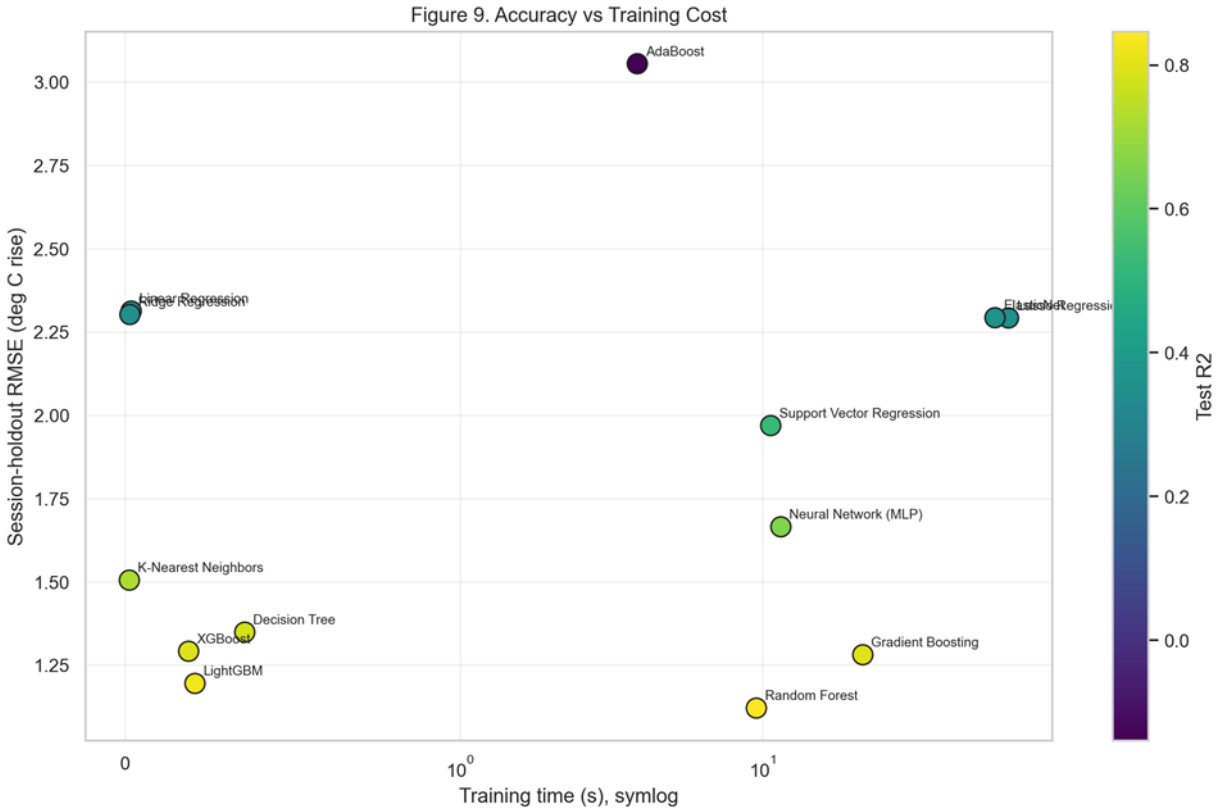

Figure 9. Accuracy versus training cost.

Random Forest gives the best session-holdout accuracy, while LightGBM gives a strong speed-accuracy trade-off. This supports model-selection discussion but does not replace hardware validation.

It shows the trade-off between prediction accuracy and training cost. Random Forest gives the best session-holdout accuracy, while LightGBM is close and trains much faster.

This supports practical model selection. If the model is retrained often or must run in a lightweight engineering workflow, LightGBM may be preferred if accuracy is the priority and retraining cost is acceptable, Random Forest is stronger.

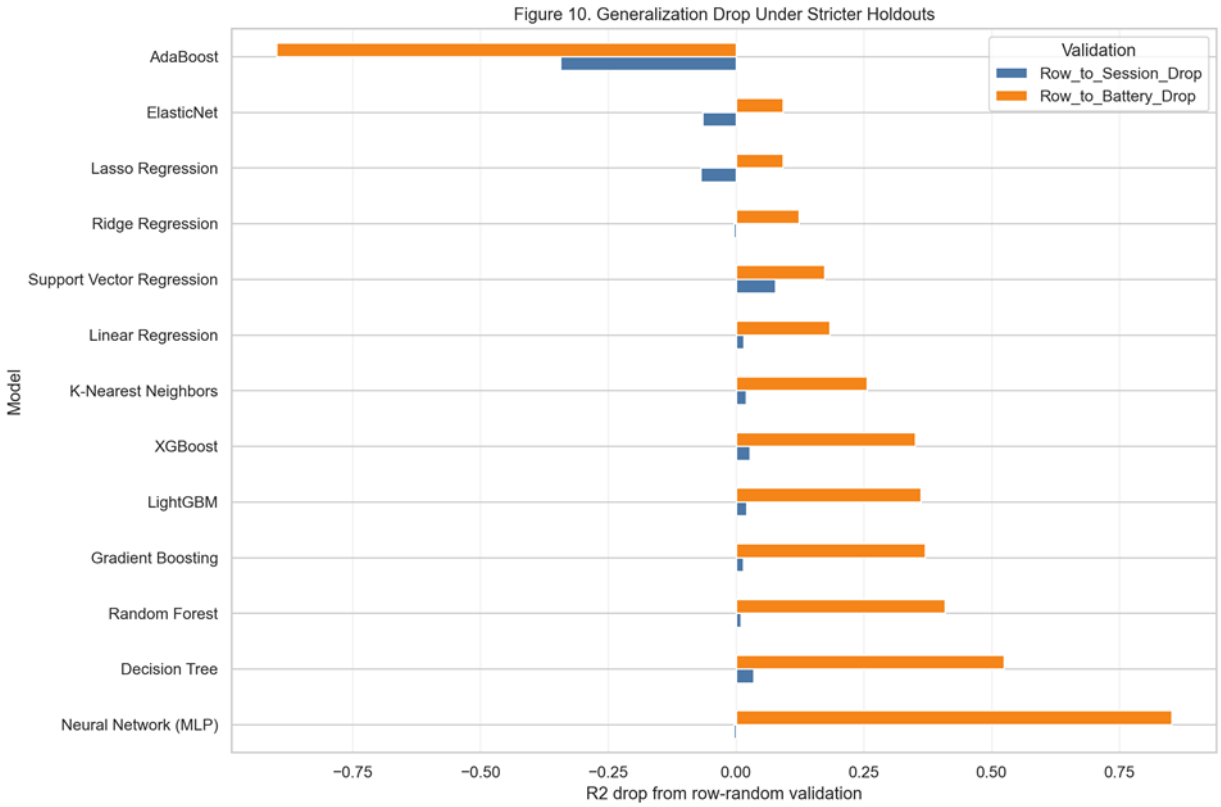

Figure 10. Generalization drop under stricter holdouts.

This figure shows how performance changes as the validation setting becomes more realistic. The drop from row-random to battery-holdout performance shows why lab-session prediction and cross-battery transfer should be discussed separately.

It shows the amount of performance lost when moving from easy row-random validation to stricter session and battery holdouts. This is one of the most important plots because it exposes overoptimistic claims. The size of the performance drop under stricter holdouts is itself one of the most important findings of the study.

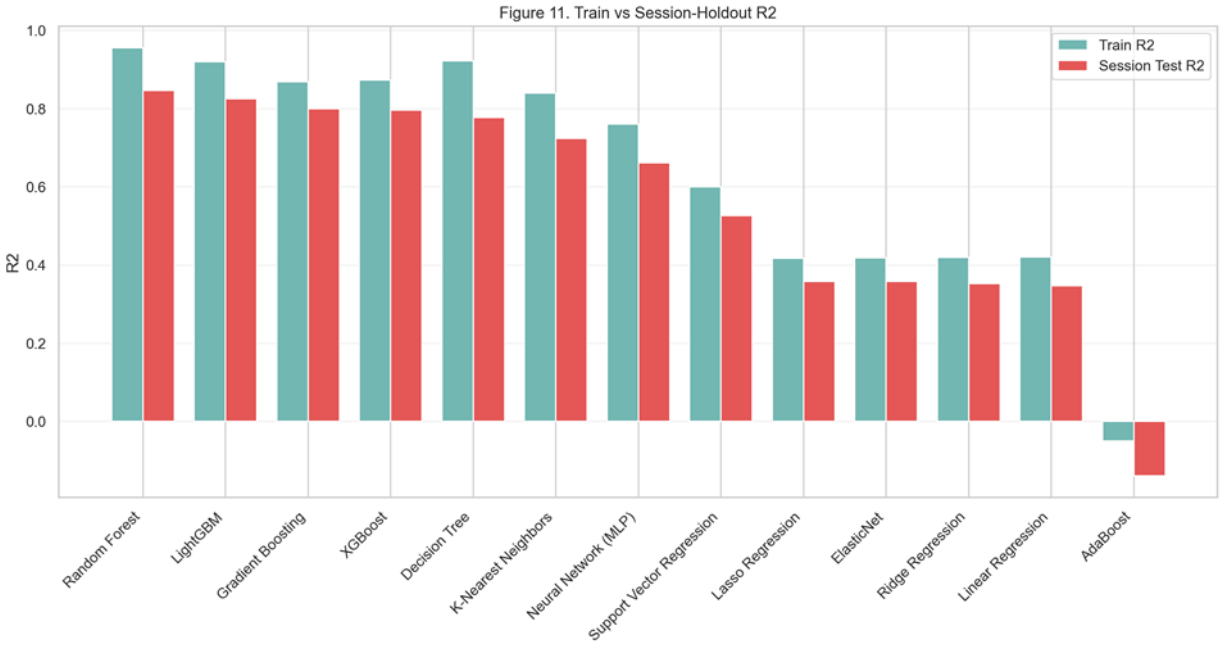

Figure 11. Train versus session-holdout R^2.

The train-test comparison shows which models overfit and which maintain session-level generalization. Random Forest and LightGBM remain the most reliable high-performing choices. It shows a comparison between training fit and session-holdout fit. Models with high training performance but weaker test performance are learning patterns that may not transfer reliably.

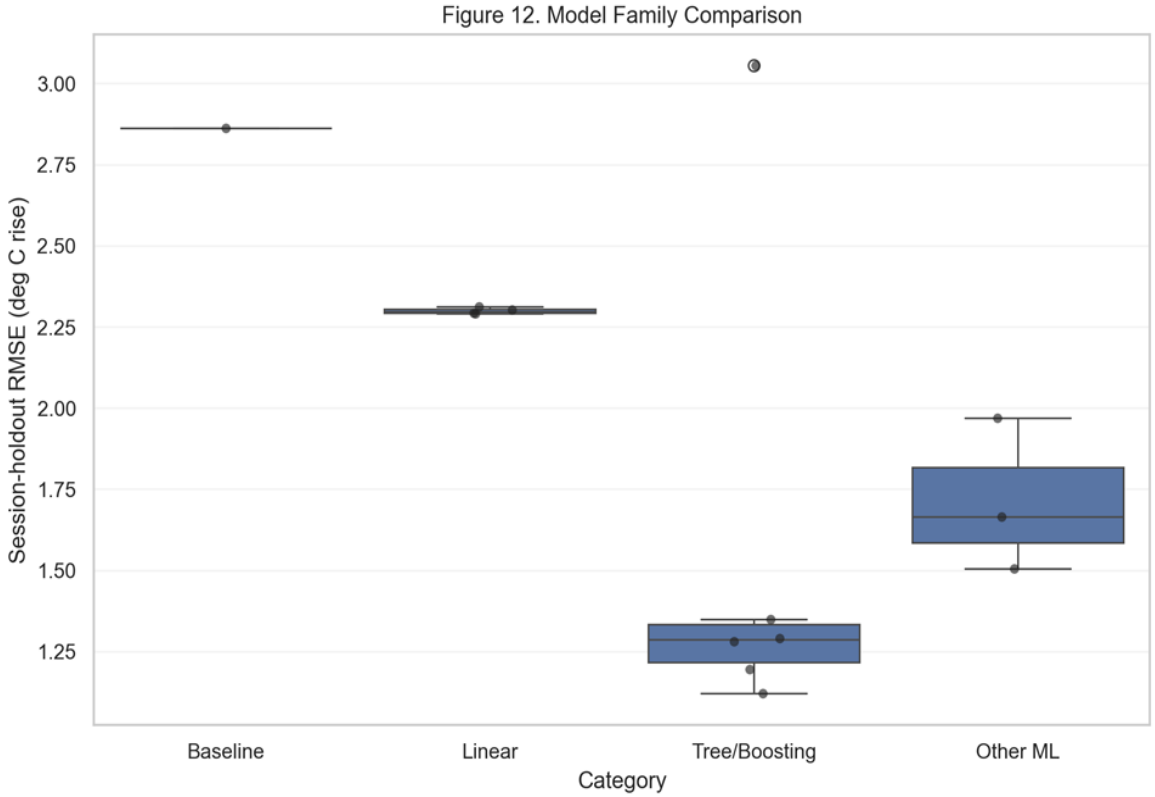

Figure 12. Model family comparison.

Tree and boosting models generally outperform linear, kernel, and neural-network alternatives for this feature set and target. This supports the model-family conclusion without claiming universal superiority.

Tree-based and boosting methods are generally stronger than linear, kernel, and neural-network alternatives for this feature set and dataset size. For this dataset and target, tree-based structure matches the problem better than the tested linear, kernel, and neural alternatives.

Figure 13. Top 5 Session-Holdout Models

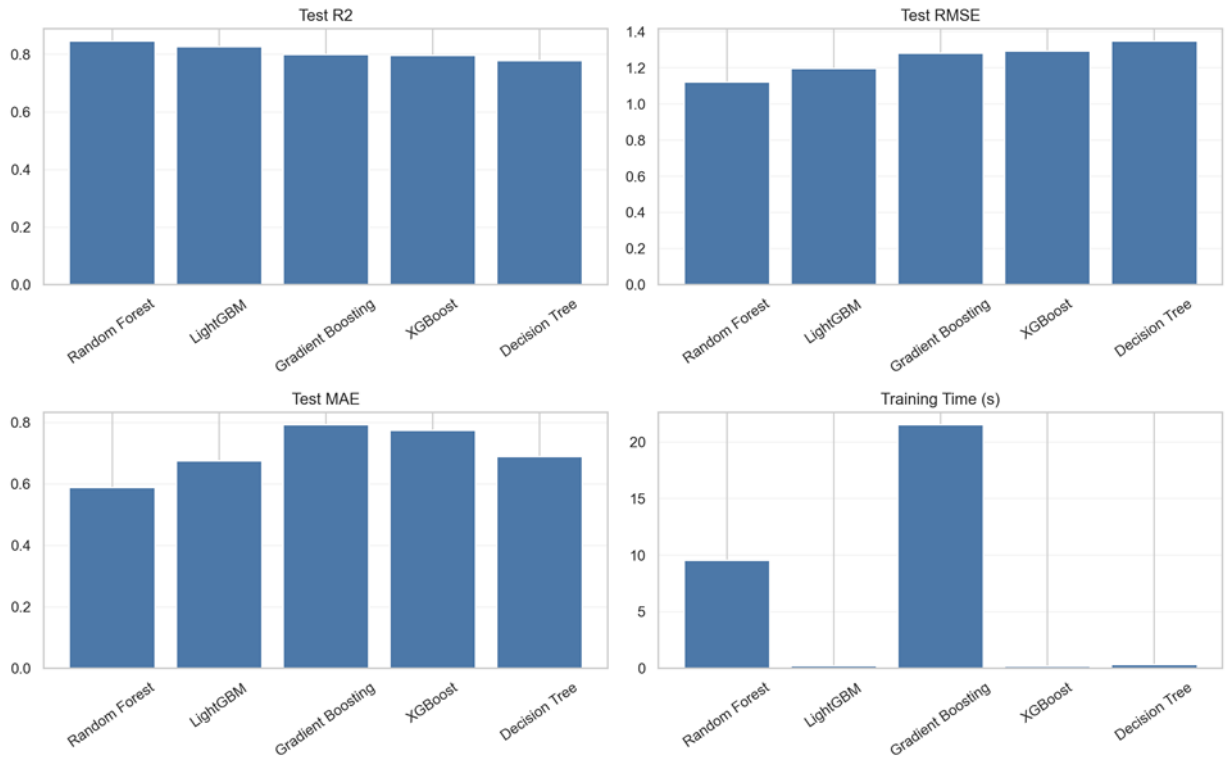

Figure 13. Top five session-holdout models.

The top-five comparison shows the practical trade-off between error and training cost. Random Forest is most accurate, while LightGBM is faster with only modestly higher error.

It delivers a detailed view of the top-performing session-holdout models across accuracy and training cost. The figure shows that the best model is not automatically the fastest or simplest model.

The final recommendation is a trade-off decision, not a single-score contest, which is why both accuracy and cost are compared. This supports a design trade-off decision. If accuracy is the priority and retraining cost is acceptable, Random Forest is the stronger

choice. For frequent retraining or tighter compute budgets, LightGBM is the more practical option.

4.4 Feature Importance And Model Diagnostics

After identifying which models perform best, the next question is why they work. The feature-importance and diagnostic results help to see which measured signals drive the prediction and whether the model's errors remain small enough to be useful in practice.

Feature	**Importance**
Voltage_measured	0.1827
SoC_Proxy	0.1647
ambient_temperature	0.1129
Voltage_Error	0.0999
Elapsed_Time_Fraction	0.0939
Power_Error	0.0841
Voltage_charge	0.0748
Rolling_Voltage_10	0.0350
Resistance_Proxy	0.0289
Command_Power	0.0250

Table 3. Top feature importances.

The feature-importance ranking highlights voltage/state-of-charge terms, ambient temperature, voltage error, elapsed charge fraction, and power-error terms. Direct temperature-derived variables are intentionally absent so that the ranking reflects information available from charging-state signals. Importance values are Random Forest mean decrease in impurity, normalised to sum to 1.0

Table 3 reports the ten highest-ranked features from the Random Forest importance ranking. The remaining eight features each contribute less than 2% of total importance individually and are all visible in Figure 14, which shows the complete 18-feature ranking.

The five most important individual features are Voltage_measured, SoC_Proxy, ambient_temperature, Voltage_Error, and Elapsed_Time_Fraction. Taken together, these show that the model learns mainly from three kinds of information: where the battery is in the charge process, what ambient condition surrounds the cell, and how closely the measured electrical behavior matches the charger command. That is a physically interpretable result.

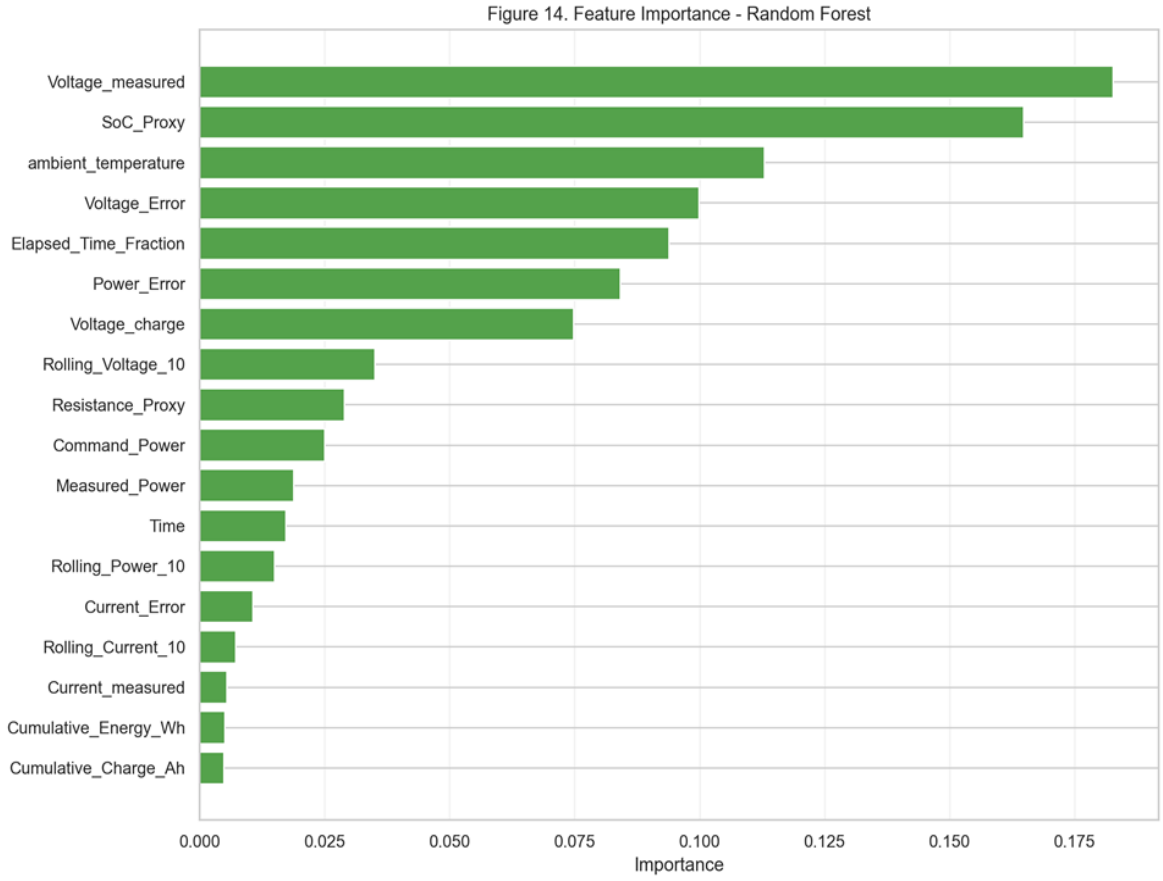

Figure 14. Feature importance.

The feature ranking supports physically meaningful interpretation without direct target leakage. Voltage and charge-progress features hold the top two positions (Voltage_measured: 18.3%, SoC_Proxy: 16.5%), while ambient temperature ranks third at 11.3%. The relatively high importance of ambient temperature is consistent with its role as a baseline thermal condition: it does not determine how fast the battery heats up during charging, but it sets the starting point from which temperature rise is measured. Its non-trivial importance score reflects that charging sessions conducted in different ambient conditions produce systematically different temperature-rise trajectories, even when electrical inputs are comparable.

It shows the feature ranking after excluding direct temperature-derived inputs. Voltage-related features, state-of-charge proxy, ambient temperature, voltage error, elapsed charge fraction, and power-error terms carry the most useful information.

The strongest predictive information comes from charge-state, ambient-condition, and error-tracking signals rather than from hidden temperature shortcuts. This helps decide which measurements must be reliable. Voltage, current, charger command, ambient temperature, and timing/progress signals should be carefully validated in a BMS because errors in these inputs can directly reduce prediction quality.

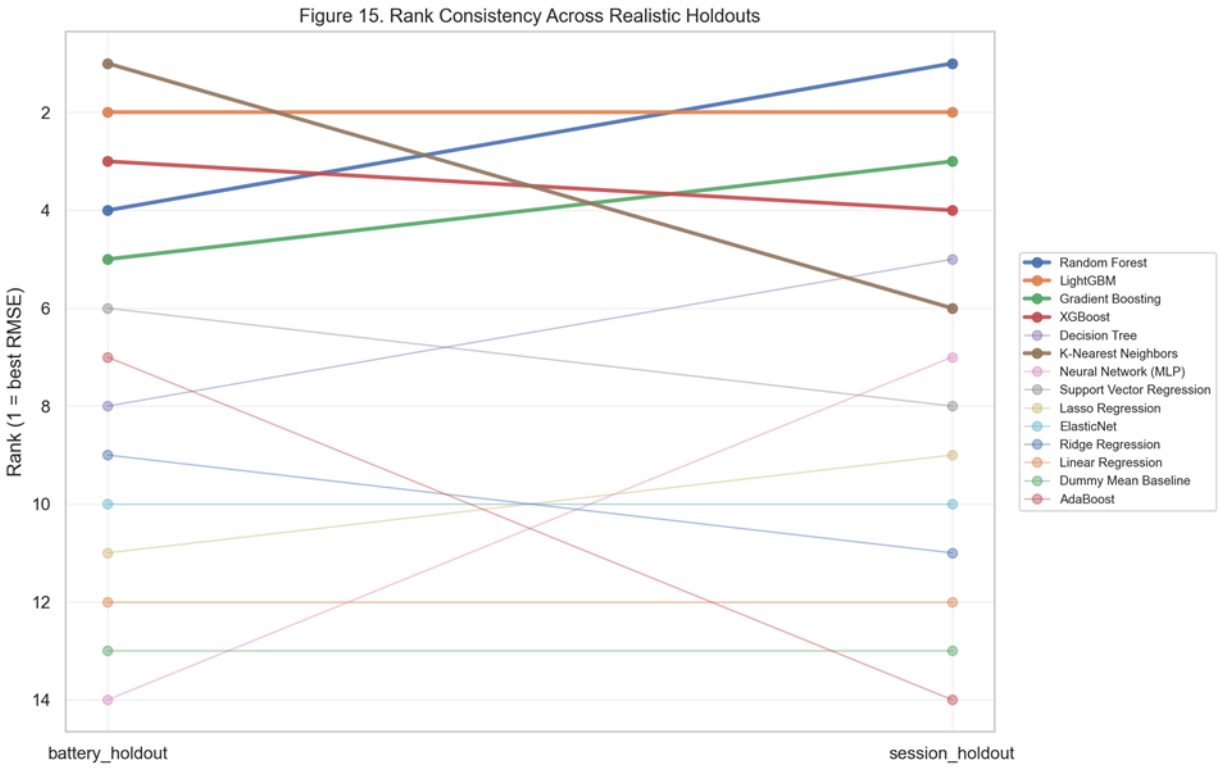

Figure 15. Rank consistency across realistic holdouts.

Rank consistency shows which models remain strong when validation becomes stricter. Random Forest and LightGBM are the most consistent high-performing recommendations across the validation settings.

A robustness view of model rankings under realistic holdouts. A model that remains near the top across validation schemes is more trustworthy than a model that wins only one easy metric. A model that remains strong across several validation schemes is more trustworthy than a model that wins only one easy test.

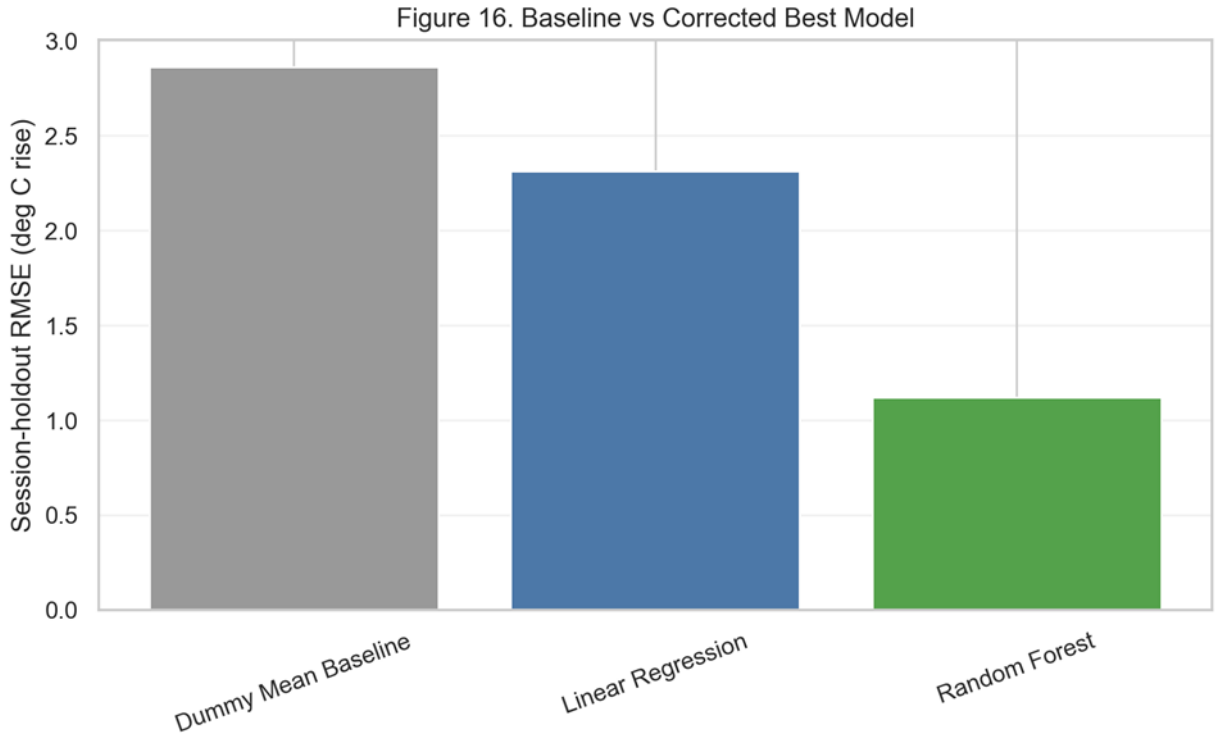

Figure 16. Baseline versus best model.

The baseline comparison shows that machine learning improves over a mean-temperature-rise baseline, but the improvement should be interpreted in degrees Celsius of temperature rise, not as proof of charging-time gains.
It shows a concrete improvement over simple baselines. The best model reduces session-holdout error substantially compared with predicting the average temperature rise or using a simple linear model.

Machine learning adds measurable predictive value beyond naive and simple linear baselines, which justifies the modelling effort. The model provides extra predictive information beyond raw monitoring and simple regression. That information can become a decision input for current limiting, cooling pre-activation, and thermal-margin checks.

Figure 17. Session-Holdout Prediction Diagnostics

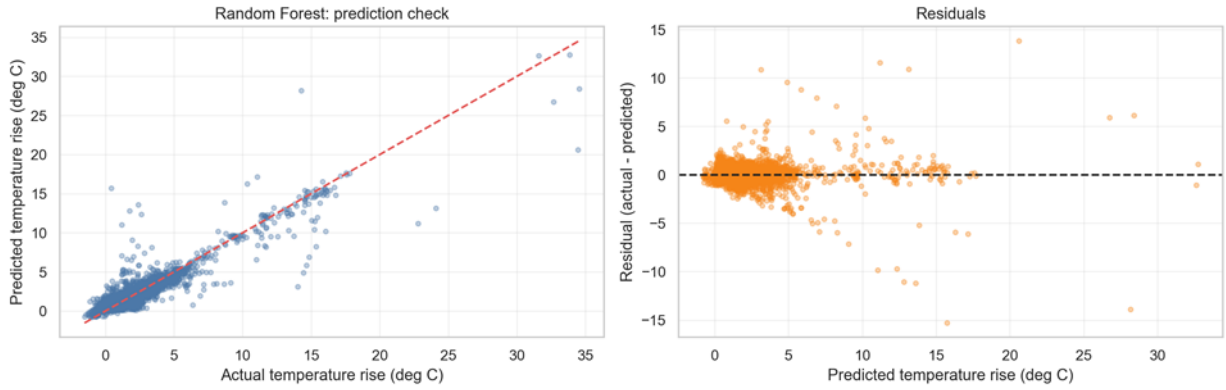

Figure 17. Session-holdout prediction diagnostics.

The prediction and residual plots show useful accuracy on unseen sessions while preserving visible residual spread. This is a credible diagnostic result, not a perfect-control claim.

It shows a predicted-versus-actual behavior and residual spread on unseen sessions. The model is useful, but the residuals also show it is not perfect and should not be used without safety margins.

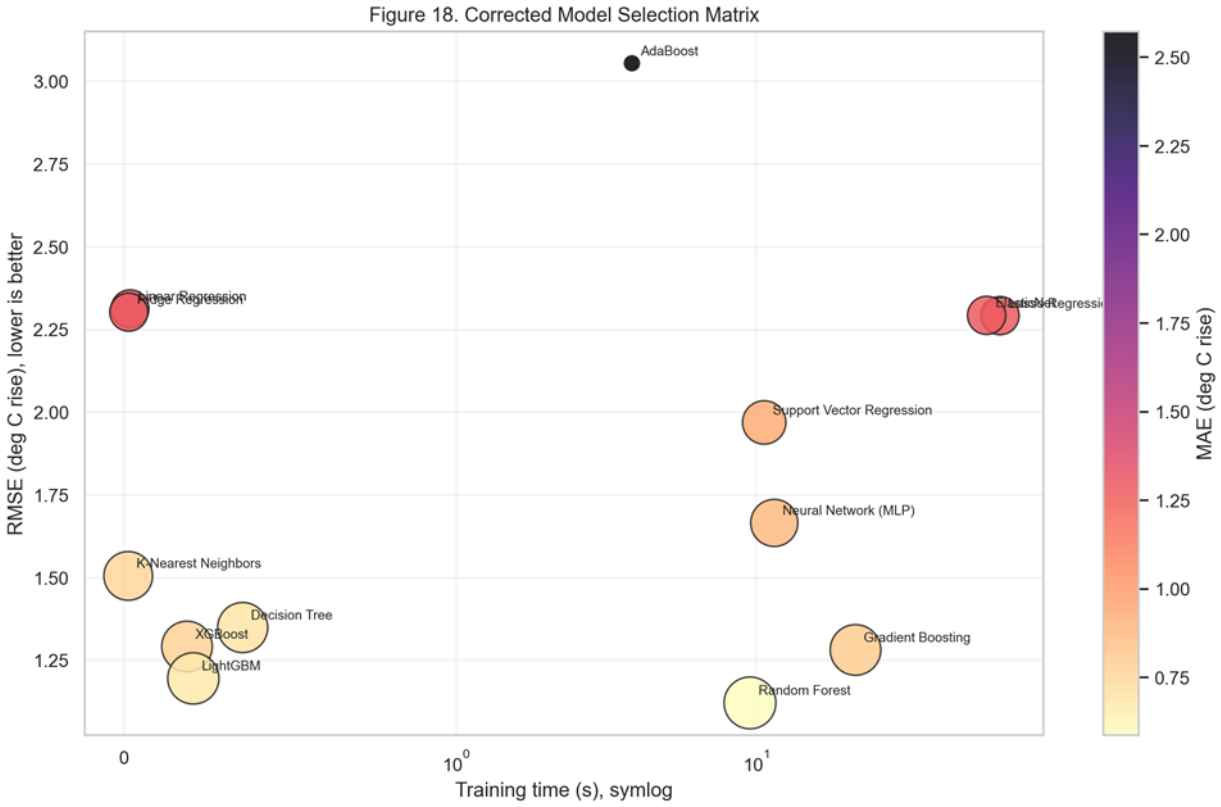

Figure 18. Model selection matrix.

The model-selection matrix supports two practical choices: Random Forest when accuracy is preferred, and LightGBM when speed and retraining cost matter more.

It delivers a combined view of error, training time, and model choice. It translates the raw performance table into an engineering selection map.

4.5 Hyperparameter Tuning

Model	Pre-tuning RMSE (°C)	Best grouped-CV RMSE (°C)	Best parameters
Random Forest	1.121	1.070	{"max_depth": 18, "min_samples_leaf": 3, "n_estimators": 220}
LightGBM	1.195	1.087	{"learning_rate": 0.08, "n_estimators": 220, "num_leaves": 63}
XGBoost	1.291	1.224	{"learning_rate": 0.08, "max_depth": 4, "n_estimators": 220}
Gradient Boosting	1.281	1.228	{"learning_rate": 0.08, "max_depth": 4, "n_estimators": 220}

Table 4. Grouped hyperparameter tuning results.

The hyperparameter tuning was performed only for the top four session-holdout models. The remaining nine models were evaluated using their default configurations as specified in Section 3.5.

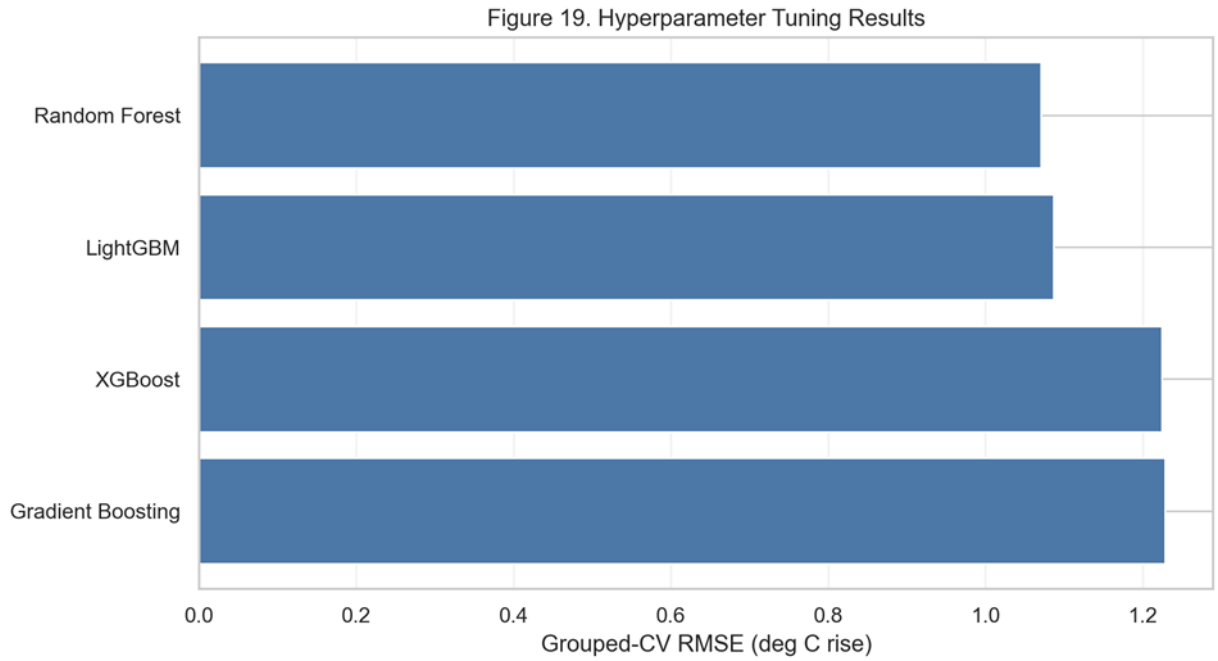

Figure 19. Hyperparameter tuning results.

Grouped tuning confirms Random Forest as the best tuned model under session-grouped cross-validation, with LightGBM close behind and much faster to train.

It shows a grouped-CV tuning results for the strongest deployable model families. Random Forest is best after grouped tuning, while LightGBM remains close with much lower tuning/training cost. The final model choice remains strong even after grouped tuning that respects session structure.

4.6 Real-World Interpretation

These experimental results support a temperature-rise prediction framework for lithium-ion charge-cycle data. They show that the model can provide useful thermal estimates for unseen sessions within the represented battery population, while also

making clear that broader deployment will require future-temperature prediction, control optimization, and validation on unseen cells, packs, and hardware.

The most important practical aspect is that the study achieved a useful prediction layer, not a completed fast-charging controller. The model has earned the ability to estimate thermal risk better than simple baselines, especially for new sessions from similar batteries. What it has not yet earned is proof of safe closed-loop control, faster charging in hardware, or universal transfer across unseen batteries.

Real-world need	**What the model provides**	**Evidence**	**How it can help**	**What is still required**
Avoid overheating during fast charging	Estimated temperature rise above ambient for a charging state	Session-holdout Random Forest RMSE = 1.121 deg C rise	A BMS can compare predicted rise against a thermal margin before choosing current.	Forward prediction and charger-in-the-loop testing.
Avoid overly conservative charging	A state-aware estimate instead of a fixed safety rule	RMSE reduced by 60.8% vs dummy mean and 51.5% vs Linear Regression on unseen sessions	When predicted thermal rise is low, a controller may allow less conservative current limits.	A real optimization policy and measured charge-time experiment.
Select a deployable model	Accuracy-cost comparison across 13 algorithms	Random Forest best accuracy; LightGBM close with faster training	Engineers can choose Random Forest for accuracy or LightGBM for	Embedded inference benchmark on target hardware.

			retraining speed.	
Generalize to unseen cells	Battery-holdout validation result	Best unseen-battery R^2 = 0.4871 and RMSE = 2.534 deg C rise	Shows some transfer exists but exposes the main deployment weakness.	More batteries, cell chemistries, pack-level data, and calibration.
Support battery-health-aware charging	Thermal-risk input, not SOH prediction	Temperature is a key driver of battery degradation	Can be one input to future aging-aware control because temperature affects degradation.	Capacity fade, resistance growth, or SOH labels and aging validation.

Table 5. Real-world usefulness of the prediction layer.

Claim	**Status**	**Reason**
Lab-cell temperature-rise prediction	Yes	All evaluated models are trained and tested on charge-cycle temperature-rise data.
Prediction on unseen charge sessions	Yes, within the represented battery population	Session-holdout test has zero shared sessions and Random Forest reaches R^2 = 0.8466.
Prediction on completely unseen batteries	Partly	Battery-holdout test has zero shared batteries, but best R^2 falls to 0.4871.
Charging-time reduction	Not directly proven	No closed-loop current schedule or hardware

		charge-time experiment was run.
Battery-health improvement	Not directly proven	The model does not predict SOH, capacity fade, resistance growth, or remaining useful life.
Production EV BMS deployment	Not yet	Pack-level data, embedded testing, safety validation, and controller validation are still required.

Table 6. Deployment-readiness interpretation.

5 Conclusion

This thesis presents a machine-learning framework for predicting lithium-ion battery temperature rise during charge-cycle operation. The study uses metadata-confirmed charge files only, excludes direct temperature-derived features, predicts temperature rise above ambient rather than absolute temperature, and reports both session-holdout and battery-holdout validation.

It shows that charge-state information contains enough structure to estimate thermal rise on new charging sessions with useful accuracy, while also showing that transfer to completely unseen batteries is still much more difficult.

In the experimental setting, the available charging information is essentially reactive: voltage, current, time, charger command, ambient temperature, and measured battery temperature are observations of what has already happened. A controller using only raw monitoring has no learned estimate of temperature rise under the current charging state, so it must either act conservatively or wait until temperature has already increased. For reference, the machine-learning models are compared against a dummy mean predictor and a Linear Regression model.

Condition	**Model or state**	**R^2**	**RMSE**	**MAE**	**Meaning**
Direct monitoring only	Raw monitoring only	Not scoreable	Not scoreable	Not scoreable	No learned forecast of temperature rise.
Simple baseline	Dummy mean predictor	-0.0000	2.862	1.637	Predicts the average temperature rise only.

Classical baseline	Linear Regression	0.3470	2.312	1.295	Captures only a simple linear mapping.
Best session-generalization model	Random Forest	0.8466	1.121	0.588	Best model on unseen charge sessions.
Best battery-transfer model	K -Nearest Neighbors	0.4871	2.534	1.247	Best model on unseen batteries.

Table 7. Baseline and transfer-performance summary.

The amount of improvement is substantial for unseen charge sessions. Compared with the dummy mean baseline, the Random Forest model reduces session-holdout RMSE by 60.8% and MAE by 64.1%. Compared with Linear Regression, it reduces RMSE by 51.5% and MAE by 54.6%. In practical terms, this changes the result from a rough average-temperature-rise estimate into a state-aware thermal prediction layer that can inform later charging-current and cooling decisions. The 28.7% improvement for KNN is calculated against the battery-holdout dummy baseline (RMSE=3.553, Table 2), not the session-holdout dummy shown in this table.

This improvement is one of the strongest results of the thesis. It means the modelling approach is not just fitting noise or reproducing an average curve. It is extracting useful information from the measured charging state and turning that information into a materially better estimate of thermal rise.

The stricter unseen-battery result is more conservative. Relative to the dummy mean baseline, the best battery-holdout model reduces RMSE by 28.7% and MAE by 29.4%. This is helpful because it shows that some information transfers to unseen cells, but the

lower battery-holdout R^2 also shows that the model is not yet strong enough to claim universal EV-pack deployment without additional data and hardware validation.

This more conservative result should not be seen as a failure. It is a useful scientific finding because it identifies where the real difficulty lies. The study therefore learns two things at once: session-level prediction is achievable, but cross-battery generalization remains the key barrier to broader deployment.

The practical value of this result is therefore an information gain for control, not a directly measured charging-time gain. A future BMS or charger controller could use this temperature-rise estimate to choose charging current with better knowledge of expected thermal response, avoid unnecessarily conservative current limits when the predicted rise is low, and reduce overheating risk when the predicted rise is high. The present study supplies the prediction layer needed for that controller; it does not yet test the controller itself.

Controller step	**How this study helps**	**Covered aspects / Current status**	**Practical next requirement**
Read charging state	Uses voltage, current, charger command, time, ambient temperature, and derived power/progress features.	Yes	Connect feature calculation to real BMS telemetry.
Estimate thermal response	Predicts temperature rise above ambient with 1.121 deg C RMSE on unseen sessions.	Yes for lab sessions	Predict future temperature over the next control horizon, not only current-state response.

Choose current limit	Prediction can warn whether a candidate current is likely to exceed the thermal margin.	Not yet	Add optimizer that tests candidate current schedules against temperature constraints.
Actuate charger/cooling	Prediction can become an input to current tapering or cooling pre-activation.	Not yet	Run charger-in-the-loop and thermal-management experiments.
Protect battery life	Lower thermal prediction error can reduce thermal overstress decisions.	Indirectly only	Add aging labels such as capacity fade, resistance growth, or SOH.

Table 8. Practical controller pathway enabled by the study.

The strongest session-holdout model is Random Forest, with $R^2 = 0.8466$, RMSE = 1.121 degree Celsius rise, and MAE = 0.588 degree Celsius rise. This supports practical use as a lab-scale thermal prediction layer for unseen charge sessions drawn from the represented battery population.

The strictest battery-holdout result is more limited. The best battery-holdout model, K-Nearest Neighbors, reaches $R^2 = 0.4871$ and RMSE = 2.534 degree Celsius rise. This shows that transfer to completely unseen batteries remains moderate and must be improved before the method can be claimed as production-ready.

Grouped hyperparameter tuning identifies Random Forest as the best tuned model, with grouped-CV RMSE = 1.070 degree Celsius rise. LightGBM remains a strong alternative when retraining speed is more important than the last increment of accuracy.

First, battery temperature rise during charging is predictable from charging-state variables with useful accuracy on unseen sessions. Second, tree-based models provide the strongest overall performance, with Random Forest leading on session-holdout accuracy and LightGBM offering an attractive speed-accuracy balance. Third, battery-to-battery variation remains the dominant challenge for any claim beyond the represented battery population.

The study therefore demonstrates a validated lab-data temperature-rise prediction framework that can support future fast-charging optimization research. It does not yet demonstrate a deployed charger-in-the-loop optimization policy, a direct charging-time reduction, or measured battery-health extension. Future work should build a forward-looking temperature predictor, optimize candidate current schedules under thermal and aging constraints, and validate the controller on unseen cells, EV-scale packs, and real charging hardware.

References

Abo Gamra, K., Bilfinger, P., Schreiber, M., Kröger, T., Allgäuer, C., & Lienkamp, M. (2024). Unlocking the full potential of electric vehicle fast-charging over lifetime through model-based aging adaptation. *Journal of Energy Storage*, *99*, 113361. https://doi.org/10.1016/j.est.2024.113361

Ahmed, S., Bloom, I., Jansen, A. N., Tanim, T., Dufek, E. J., Pesaran, A., Burnham, A., Carlson, R. B., Dias, F., Hardy, K., Keyser, M., Kreuzer, C., Markel, A., Meintz, A., Michelbacher, C., Mohanpurkar, M., Nelson, P. A., Robertson, D. C., Scoffield, D., … Zhang, J. (2017). Enabling fast charging – A battery technology gap assessment. *Journal of Power Sources*, *367*, 250–262. https://doi.org/10.1016/j.jpowsour.2017.06.055

Breiman, L. (2001). Random Forests. *Machine Learning*, *45*(1), 5–32. https://doi.org/10.1023/A:1010933404324

Bruj, O., & Calborean, A. (2025). Lifecycle Evaluation of Lithium-Ion Batteries Under Fast Charging and Discharging Conditions. *Batteries*, *11*(2), 65. https://doi.org/10.3390/batteries11020065

Chen, T., & Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, 785–794. https://doi.org/10.1145/2939672.2939785

Chen, Z., Zhang, J., Liu, C., Yang, C., & Chen, S. (2026). Thermal Runaway in Lithium-Ion Batteries: A Review of Mechanisms, Prediction Approaches, and Mitigation Strategies. *Batteries*, *12*(3), 88. https://doi.org/10.3390/batteries12030088

Gao, Y., Zhang, X., Cheng, Q., Guo, B., & Yang, J. (2019). Classification and Review of the Charging Strategies for Commercial Lithium-Ion Batteries. *IEEE Access*, *7*, 43511–43524. https://doi.org/10.1109/ACCESS.2019.2906117

Greenwood, M., Wentker, M., & Leker, J. (2021). A bottom-up performance and cost assessment of lithium-ion battery pouch cells utilizing nickel-rich cathode active materials and silicon-graphite composite anodes. *Journal of Power Sources Advances*, *9*, 100055. https://doi.org/10.1016/j.powera.2021.100055

Hu, Q., Amini, M. R., Wiese, A., Seeds, J. B., Kolmanovsky, I., & Sun, J. (2023). Electric Vehicle Enhanced Fast Charging Enabled by Battery Thermal Management and Model Predictive Control. *IFAC-PapersOnLine*, *56*(2), 10684–10689. https://doi.org/10.1016/j.ifacol.2023.10.721

Hu, X., Li, S., & Peng, H. (2012). A comparative study of equivalent circuit models for Li-ion batteries. *Journal of Power Sources*, *198*, 359–367. https://doi.org/10.1016/j.jpowsour.2011.10.013

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., & Liu, T.-Y. (2017). LightGBM: A Highly Efficient Gradient Boosting Decision Tree. *Advances in Neural Information Processing Systems*, *30*. https://proceedings.neurips.cc/paper/2017/hash/6449f44a102fde848669bdd9eb6b76fa-Abstract.html

Liaw, B. Y., & Dubarry, M. (2007). From driving cycle analysis to understanding battery performance in real-life electric hybrid vehicle operation. *Journal of Power Sources, Hybrid Electric Vehicles*, *174*(1), 76–88. https://doi.org/10.1016/j.jpowsour.2007.06.010

Lin, X., Perez, H. E., Mohan, S., Siegel, J. B., Stefanopoulou, A. G., Ding, Y., & Castanier, M. P. (2014). A lumped-parameter electro-thermal model for cylindrical batteries. *Journal of Power Sources*, *257*, 1–11. https://doi.org/10.1016/j.jpowsour.2014.01.097

Liu, K., Li, K., Peng, Q., & Zhang, C. (2019). A brief review on key technologies in the battery management system of electric vehicles. *Frontiers of Mechanical Engineering*, *14*(1), 47–64. https://doi.org/10.1007/s11465-018-0516-8

Luo, C., Zhang, Z., Qiao, D., Lai, X., Li, Y., & Wang, S. (2022). Life Prediction under Charging Process of Lithium-Ion Batteries Based on AutoML. *Energies*, *15*(13), 4594. https://doi.org/10.3390/en15134594

Makki, M., Lee, C. W., & Ayoub, G. (2023). Stress Distribution Inside a Lithium-Ion Battery Cell during Fast Charging and Its Effect on Degradation of Separator. *Batteries*, *9*(10), 502. https://doi.org/10.3390/batteries9100502

Mesquita, A. R., de Abreu, V. H. S., Poyares, C. N., & Santos, A. S. (2025). Barriers to Electric Vehicle Adoption: A Framework to Accelerate the Transition to Sustainable Mobility. *Sustainability*, *17*(18), 8318. https://doi.org/10.3390/su17188318

Mishra, S., Choubey, A., Reddy, B. A., & Misra, R. (2024). Enhancing EV lithium-ion battery management: Automated machine learning for early remaining useful life prediction with innovative multi-health indicators. *The Journal of Supercomputing*, *80*(14), 20813–20860. https://doi.org/10.1007/s11227-024-06264-w

Murmann, J. P., & Schuler, B. A. (2023). Exploring the structure of internal combustion engine and battery electric vehicles: Implications for the architecture of the automotive industry. *Industrial and Corporate Change*, *32*(1), 129–154. https://doi.org/10.1093/icc/dtac049

Ng, M.-F., Zhao, J., Yan, Q., Conduit, G. J., & Seh, Z. W. (2020). Predicting the state of charge and health of batteries using data-driven machine learning. *Nature Machine Intelligence*, *2*(3), 161–170. https://doi.org/10.1038/s42256-020-0156-7

Pamidimukkala, A., Kermanshachi, S., Rosenberger, J. M., & Hladik, G. (2023). Evaluation of barriers to electric vehicle adoption: A study of technological, environmental, financial, and infrastructure factors. *Transportation Research Interdisciplinary Perspectives*, *22*, 100962. https://doi.org/10.1016/j.trip.2023.100962

Richardson, R. R., Osborne, M. A., & Howey, D. A. (2019). Battery health prediction under generalized conditions using a Gaussian process transition model. *Journal of Energy Storage*, *23*, 320–328. https://doi.org/10.1016/j.est.2019.03.022

Roman, D., Saxena, S., Robu, V., Pecht, M., & Flynn, D. (2021). Machine learning pipeline for battery state-of-health estimation. *Nature Machine Intelligence*, *3*(5), 447–456. https://doi.org/10.1038/s42256-021-00312-3

Samanta, A., Chowdhuri, S., & Williamson, S. S. (2021). Machine Learning-Based Data-Driven Fault Detection/Diagnosis of Lithium-Ion Battery: A Critical Review. *Electronics*, *10*(11), 1309. https://doi.org/10.3390/electronics10111309

Severson, K. A., Attia, P. M., Jin, N., Perkins, N., Jiang, B., Yang, Z., Chen, M. H., Aykol, M., Herring, P. K., Fraggedakis, D., Bazant, M. Z., Harris, S. J., Chueh, W. C., & Braatz,

R. D. (2019). Data-driven prediction of battery cycle life before capacity degradation. *Nature Energy*, *4*(5), 383–391. https://doi.org/10.1038/s41560-019-0356-8

Singh, A., Feltner, C., Peck, J., & Kuhn, K. I. (2021). *Data Driven Prediction of Battery Cycle Life Before Capacity Degradation* (arXiv:2110.09687). arXiv. https://doi.org/10.48550/arXiv.2110.09687

Tian, J., Li, S., Liu, X., Yang, D., Wang, P., & Chang, G. (2022). Lithium-ion battery charging optimization based on electrical, thermal and aging mechanism models. *Energy Reports*, *8*, 13723–13734. https://doi.org/10.1016/j.egyr.2022.10.059

Tomaszewska, A., Chu, Z., Feng, X., O'Kane, S., Liu, X., Chen, J., Ji, C., Endler, E., Li, R., Liu, L., Li, Y., Zheng, S., Vetterlein, S., Gao, M., Du, J., Parkes, M., Ouyang, M., Marinescu, M., Offer, G., & Wu, B. (2019). Lithium-ion battery fast charging: A review. *eTransportation*, *1*, 100011. https://doi.org/10.1016/j.etran.2019.100011

Wang, S., Wu, T., Xie, H., Li, C., Zhang, J., Jiang, L., & Wang, Q. (2022). Effects of Current and Ambient Temperature on Thermal Response of Lithium Ion Battery. *Batteries*, *8*(11), 203. https://doi.org/10.3390/batteries8110203

Wang, Y., Tian, J., Sun, Z., Wang, L., Xu, R., Li, M., & Chen, Z. (2020). A comprehensive review of battery modeling and state estimation approaches for advanced battery management systems. *Renewable and Sustainable Energy Reviews*, *131*, 110015. https://doi.org/10.1016/j.rser.2020.110015

Wu, B., Yufit, V., Marinescu, M., Offer, G. J., Martinez-Botas, R. F., & Brandon, N. P. (2013). Coupled thermal–electrochemical modelling of uneven heat generation in

lithium-ion battery packs. *Journal of Power Sources*, *243*, 544–554. https://doi.org/10.1016/j.jpowsour.2013.05.164

Xiong, R., Cao, J., Yu, Q., He, H., & Sun, F. (2018). Critical Review on the Battery State of Charge Estimation Methods for Electric Vehicles. *IEEE Access*, *6*, 1832–1843. https://doi.org/10.1109/ACCESS.2017.2780258

Zentani, A., Almaktoof, A., & Kahn, M. T. (2024). A Comprehensive Review of Developments in Electric Vehicles Fast Charging Technology. *Applied Sciences*, *14*(11), 4728. https://doi.org/10.3390/app14114728