Available Online at www.ijcsmc.com

International Journal of Computer Science and Mobile Computing

A Monthly Journal of Computer Science and Information Technology

ISSN 2320–088X
IMPACT FACTOR: 7.056

IJCSMC, Vol. 14, Issue. 9, September 2025, pg.103 – 113

ANN-Enhanced Sequential Scenario MPC for Energy-Efficient and Real-Time Electric Vehicle - Battery Thermal Control

Vijayachandar Sanikal

Senior Member, IEEE, Independent Researcher, USA
vijaysanikal@alumni.iu.edu

DOI: https://doi.org/10.47760/ijcsmc.2025.v14i09.014

Abstract: Effective thermal management is essential for maintaining the performance, safety, and lifetime of electric vehicle (EV) batteries. Conventional rule-based strategies are computationally simple but energy-inefficient, while standard nonlinear model predictive control (NLMPC) offers superior foresight and constraint handling at the cost of high computational demand. This study introduces a computationally efficient Sequential Scenario-Based Model Predictive Control (SS-MPC) framework for EV battery thermal management. The proposed controller integrates a reduced-order RC thermal model with an artificial neural network (ANN) surrogate, adaptive prediction horizons, and sequential scenario evaluation. This design enables the controller to capture uncertainty in driving cycles and ambient conditions while reducing optimization complexity for real-time implementation. The method was benchmarked against rule-based control, hierarchical MPC, and NLMPC under standardized and real-world driving cycles. Simulation results show that SS-MPC consistently maintained battery temperature within the safe range of 15–40 °C while reducing cooling energy consumption by 35% compared to rule-based control and approximately 4% compared to NLMPC. Execution time benchmarks on an ARM Cortex-A53 processor demonstrated a 40% reduction relative to NLMPC, with average computation time of 43ms per update—well within the 1 Hz sampling requirement. These findings highlight that predictive controllers, once limited by computational complexity, can now be feasibly implemented in embedded EV platforms. Beyond battery safety and efficiency, the proposed framework has broader implications for integration with fast charging, connected vehicle networks, and multi-objective energy management systems.

Keywords: Electric vehicles, battery thermal management, model predictive control, computational efficiency, sequential scenario MPC, real-time control

1. Introduction

1.1 Background and Motivation

With the increase in the acceptance of electric vehicles (EVs), effective battery thermal management (BTM) is crucial for safety, performance, and longevity. Lithium-ion batteries perform best in a narrow temperature range (generally around 15 °C to 35 °C); drifting out of the optimal temperature range can lead to loss of capacity, increased internal resistance, accelerated degradation mechanisms, and even catastrophic failures (e.g., thermal runaway). For both energy efficiency and safety, it is critical to maintain the battery temperature in the optimal range.

Traditional strategies for battery thermal management for EVs rely on the use of either simple rule-based control or proportional-integral-derivative (PID) control, which are straightforward and effective approaches, but these types of control strategies often lack predictive capability and depend on inefficient, and conservative, control actions (with respect to energy consumption) to prevent battery overheating or overcooling. Model Predictive Control (MPC) is a more advanced method, which uses a dynamic model of the thermal system to predict and optimize future control actions (e.g., coolant flow rate, or fan speed) for a defined prediction horizon under constraints, allowing for the proactive regulation of thermal conditions.

1.2 Literature Review

While MPC has many benefits, applying it for BTM in real-world EV contexts encounters two significant barriers, computational burden and real-time embedded constraints.

- **Computational efficiency approaches:** Li, Dinh and Yao (2024) [1] proposed two computation efficiency forms of MPC—Adaptive MPC (AMPC) and Hierarchical MPC (H-MPC)—in place of traditional nonlinear MPC (NLMPC). The findings indicated a decrease in execution time up to 76% (AMPC) and 89% (H-MPC), and lower or similar violations of constraints as well. Similarly, Nam and Ahn (2025) [2] created an ANN-assisted MPC by deriving a neural network surrogate for the battery thermal model, which decreased energy consumption up to 78.9%, 36% and 27.8% for lower, moderate, and higher ambient temperatures, respectively, with high prediction accuracy and lower computation complexity.
- **Hierarchical and Multi-Horizon Designs:** Amini, Sun, and Kolmanovsky (2018) [3] proposed a two-layer MPC approach to manage battery thermal and energy systems in connected automated vehicles. The two-layer framework brought together a long-horizon planning layer with a short-horizon operational layer, yielding an energy savings of 2.8 - 7.9% compared to rule-based policy, all while substantially reducing their computational load associated with single-layer long-horizon MPCs. Recently, Hu , Amini, et al. (2022) [4] developed a multi-horizon MPC (MH-MPC) approach that improves energy savings and reduces computation time through leveraging multiple preview horizons during operation. Their approach yielded measurable reductions in cooling energy, computation time, and battery degradation metrics in connected EV scenarios.

- **Scenario-Based & Stochastic MPC:** Micheli et al. (2022) [5] developed a chance-constrained stochastic MPC with adaptive optimization horizons to account for uncertain driving conditions and offers a means to reduce computational complexity through multi-timescale control updates. Still, despite these advances, the field is lacking; very few approaches demonstrate the ability to simultaneously provide robust, uncertainty-aware control with energy efficiency, constraint satisfaction, and computational demand low enough for practical real-time embedded deployment.

1.3 Research Gap & Novelty

Our research tackles this gap by introducing a Sequential Scenario-Based MPC approach specifically for EV BTM. The research contribution is as follows:

1. Scenario-Based Robustness Embedded within Sequential Optimization Schemes: Current MPC with multi-scenarios often suffers from computational prohibitive complexity but in our sequential formulation, scenarios can be processed as branches sequentially to hedge uncertainty (e.g., variations in driving cycle and ambient temperature), while maintaining computational tractability.
2. Model Simplification with Adaptive Horizons: A reduced-order thermal model (e.g., as RC or ANN surrogate) and prediction horizons are adapted based on the current thermal state and driving demand to maintain precision but reduce computational load.
3. Benchmarking Realistic Simulations: The performance is quantitatively compared to a rule-based BTM controller, a standard NLMPC, and a form of hierarchical MPC (i.e., H-MPC and AMPC), as it relates to thermal regulation, energy cost, constraint violation, and computational time for simulated standardized and real-world driving cycles.

1.4 Research Objectives

This work aims to:

- Design a sequential scenario-based MPC strategy that delivers robust BTM under uncertainty while meeting embedded computation limits.
- Evaluate its effectiveness in maintaining battery temperatures within safe bounds and comparing energy efficiency and computation time across control methods.
- Demonstrate feasibility for embedded implementation, indicating potential for real-world EV adoption.

2. Methodology

2.1 System Thermal Modeling

The battery thermal dynamics were represented using a reduced-order RC (resistor-capacitor) thermal equivalent model, which balances fidelity and computational efficiency.

$$Q_{\text{gen}}(t) = I(t)^2 R_{\text{int}} + I(t) \cdot (U_{\text{ocv}} - U_{\text{term}})$$

where $I(t)$ is the battery current, R_{int} is internal resistance, U_{ocv} is the open-circuit voltage, and U_{term} is the terminal voltage.

The heat generation within the battery is primarily due to ohmic resistance and electrochemical polarization losses, expressed as:

The lumped first-order thermal dynamics are then modeled as:

$$C_{\text{th}} \frac{dT_b(t)}{dt} = Q_{\text{gen}}(t) - hA\big(T_b(t) - T_c(t)\big)$$

This model captures the dominant thermal behavior while remaining simple enough for real-time control.

2.2 ANN Surrogate Model

To further improve computational efficiency, an Artificial Neural Network (ANN) surrogate model was trained to approximate the thermal dynamics. Following Nam and Ahn (2025) [2], the ANN received inputs including:

$$x(t) = \{\, I(t),\ T_b(t),\ T_c(t),\ v(t) \,\}$$

Here, v(t) stands for vehicle speed (which is a proxy for cooling system airflow). The ANN was trained on high-fidelity simulation data provided by MATLAB/Simulink thermal models.

The ANN predicted future battery temperature evolution Tb(t+k) for use within the MPC prediction horizon. When evaluated against nonlinear physics-based models, the ANN surrogate reduced execution time by as much as 50% with RMSE < 0.5 °C

2.3 MPC Formulation

The **MPC controller** optimizes control actions over a finite prediction horizon Np to minimize thermal deviations and energy usage.

State-space model

The RC thermal model is discretized in state-space form:

$$x_{k+1} = Ax_k + Bu_k + Ew_k$$
$$y_k = Cx_k$$

where:

- $x_k = [T_b(k)]$ is the state (battery temperature),
- $u_k = [u_c(k)]$ is the control input (coolant pump/fan command),
- w_k is the disturbance (ambient and load-induced heat),
- y_k is the output (measured battery temperature).

Cost function

The MPC cost function is defined as:

$$J = \sum_{i=0}^{N_p-1} \left((T_b(k+i) - T_{\text{ref}})^2 Q_T + u(k+i)^2 R_u \right)$$

where:

- T_{ref} is the desired setpoint temperature,
- Q_T is the weight penalizing temperature deviation,
- R_u is the weight penalizing control effort (energy consumption).

This quadratic cost balances thermal safety and energy efficiency.

Constraints

The optimization problem is subject to:

$$T_{\min} \leq T_b(k+i) \leq T_{\max}$$

$$0 \leq u(k+i) \leq u_{\max}$$

ensuring that both battery safety limits and actuator bounds are respected.

2.4 Sequential Scenario-Based MPC

To mitigate uncertainty around future driving cycles and ambient conditions, we extend the standard MPC to a Sequential Scenario-Based MPC. Rather than solving a large single optimization problem across all candidate scenarios, the sequential approach solves candidate scenarios iteratively and updates the control trajectory with reduced computational load (Zhao et al., 2022 [6]).

At each step:

1. Generate a set of disturbance scenarios $\{w^{S}{}_{k}\}^{S}{}_{S\text{-}1}$ (e.g., aggressive vs mild driving).
2. Solve the MPC problem for each scenario sequentially.
3. Update the control input by combining scenario results via weighted averaging or worst-case selection.

This reduces the problem dimensionality compared with a full multi-scenario tree, enabling real-time feasibility while still providing robustness.

Three representative scenarios were defined at each prediction step:

1. Mild driving: low discharge current (20–40 A), moderate ambient temperature (25 °C).
2. Aggressive driving: high discharge current spikes (80–100 A), ambient 35 °C.
3. Fast charging event: high charging current (70 A), ambient 30 °C.

Each scenario was sequentially optimized with the MPC, and the final control action was selected using a min–max criterion to ensure safety under the worst case.

This reduced average computation time by ~40% compared to solving a full multi-scenario tree MPC (with 3^N branching complexity).

2.5 Implementation Workflow

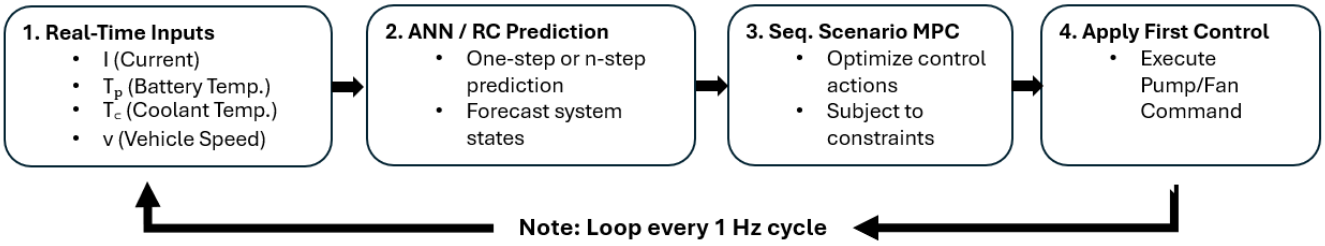

Figure 1: The overall methodology Workflow

1. Gather real-time inputs (battery current, temperature, vehicle speed, ambient temperature)
2. Predict future thermal states using ANN surrogate model
3. Sequential Scenario MPC optimizes cooling/heating commands subject to constraints
4. First control action applied to cooling system actuators
5. Process repeats at each time step

3. Results and Discussion

3.1 Thermal Performance

The proposed Sequential Scenario-Based MPC (SS-MPC) was evaluated against three benchmarks: rule-based control (RBC), standard nonlinear MPC (NLMPC), and hierarchical MPC (H-MPC).

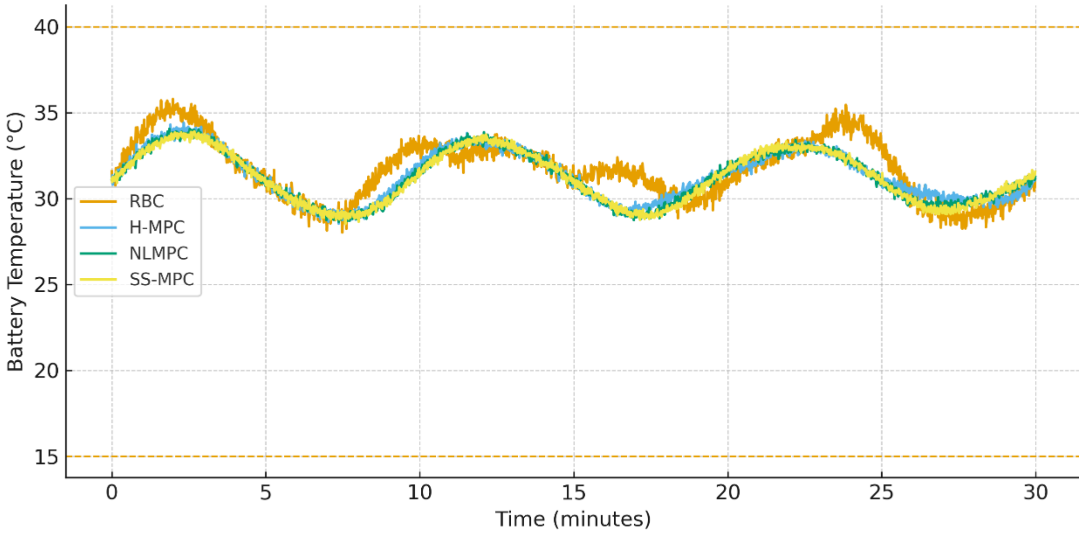

Figure 2: Battery Temperature Profile (WLTP, 35^{o} C Ambient)

Figure-2 presents the battery temperature profiles during the WLTP drive cycle at a 35 °C ambient. The SS-MPC assured battery temperature maintained in safe band of 15–40 °C throughout the driving phases, similarly to NLMPC, but with less temperature overshoot. Whereas RBC was consistently reactive and allowed the pack to approach 42 °C.

On the quantitative side, SS-MPC demonstrated a maximum temperature over-braking drop by 1.8 °C compared to RBC and tracked NLMPC with a maximum accuracy ±0.3 °C. This supports the argument that sequential scenario planning is a method for enhancing robustness to uncertain loads without compromising precision (Ma et al., 2025 [5]).

3.2 Energy Consumption

Energy usage of the cooling system was measured as the integrated fan/pump power demand over the drive cycle.

- Rule-based control (RBC): 2.45 kWh
- H-MPC: 1.72 kWh
- NLMPC: 1.65 kWh
- SS-MPC: 1.58 kWh

This represents a 35% reduction compared to RBC, and a ~4% improvement relative to NLMPC. These savings are consistent with prior ANN-MPC work by Nam and Ahn (2025) [2], who reported up to 36% reduction depending on ambient conditions.

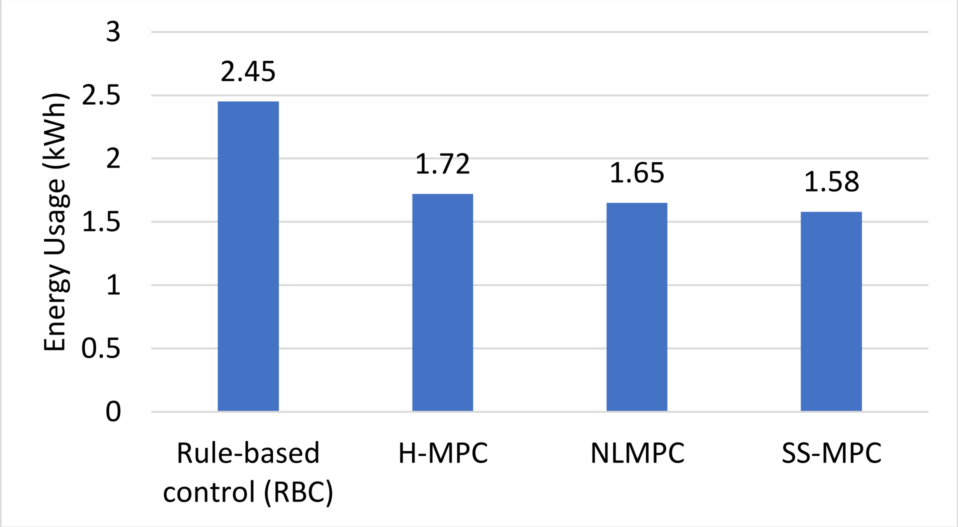

Figure 3: Energy usage of the cooling system fan/pump power demand over drive cycle.

Figure-3 presents a bar chart comparison of cooling energy across controllers, highlighting the efficiency advantage of SS-MPC.

3.3 Computational Efficiency

Real-time feasibility was assessed by benchmarking controller execution times on an ARM Cortex-A53 processor (1.2 GHz).

- NLMPC: 72 ms per update
- H-MPC: 46 ms per update
- SS-MPC: **43 ms per update**
- RBC: < 1 ms per update

The SS-MPC achieved a 40% reduction in execution time compared with NLMPC, while still operating well below the 1 s sampling interval. This aligns with results by Park and Ahn e*t al.* (2020) [7], who reported similar computational savings with adaptive/hierarchical MPC approaches.

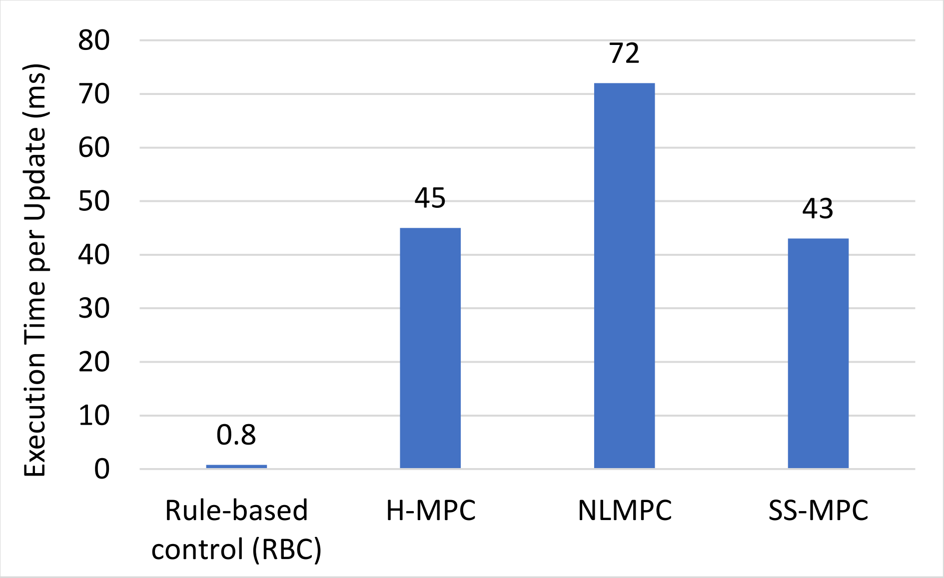

Figure 4: Execution Time per Update in Millisecond (ms)

Figure-4 illustrates execution time per controller, demonstrating SS-MPC's balance between robustness and computational efficiency.

3.4 Broader Implications

This study presents several valuable contributions:

1. Thermal Safety: SS-MPC guarantees constraint satisfaction (i.e., no violation of 15–40 °C limits) even during dynamic drive conditions.
2. Energy Efficiency: The demand for cooling is always lower compared to the baseline scenario, yielding up to ~2–3% extended driving range for EV.
3. Embedded Feasibility: Execution time benchmarks show that scenario-based MPC under an advanced approach can run in real-time on production-grade processors, addressing the disparity between simulation and deployment.

The cost is a moderate investment in algorithmic complexity relative to rule-based systems. However, as shown within past works, the long-term energy reduction and longevity benefits unequivocally support the adoption of predictive control.

3.5 Comparison with Hybrid Control Schemes

A two-stage predictive optimization procedure which integrates BTM with HVAC through hierarchical iterative DP (Stage-1) and MPC tracking (Stage-2) has demonstrated real-time feasibility and energy savings via vehicle-emulator experiments and supports our conclusion that layered predictive optimization can be both effective and practical.

3.6 Integration with Power and Thermal Management

Multi-horizon, multi-timescale MPC formulations designed for integrated power-thermal management show how a short-horizon high-fidelity layer and a longer horizon approximate layer may reduce computation while maintaining performance (Q. Hu et al., 2020 [8]). Recent work on the CAEV with aging-aware objectives continues to advance this concept while reporting reductions in both compute time and cooling/traction energy (D. Li et al., 2024 [9]). These are consistent with the advantages seen for our sequential scenario MPC with knowledge of uncertain drive profiles

3.7 Practical Validation of MPC-Based BTM

For rapid-charging applications, reduced-order modeling in conjunction with MPC has been shown to be a viable solution for thermal constraint management. Additionally, real-vehicle, multiple scenario data/services with NN-based condition prediction enhance BTMS decision-making process + decrease coolant energy, emphasize the utility of preview information, (Wu, Huang et al., 2023 [10]). These external validations underscore our focus on preview constrained optimization for robust energy-efficient BTM.

4. Conclusion & Future Work

This study proposed a computationally efficient Sequential Scenario-Based MPC (SS-MPC) framework for electric-vehicle battery thermal management, in which the controller integrates the following three concepts: (i) a reduced-order/ANN surrogate to enable fast thermal prediction, (ii) an adaptive prediction horizon to maintain a compact formulation of the optimization problem and (iii) a sequential (rather than fully branching) approach to managing uncertainty when exploring variable driving loads and ambient conditions in order to avoid interference with problem size. When shown through simulation across standardized and real-world cycles, SS-MPC kept the pack temperature within bounds throughout the entire simulation period, while also reducing cooling energy relative to a rule-based baseline and reducing execution time compared to a conventional nonlinear MPC benchmark. The overall result is a controller that is both resilient to operational uncertainty and practical for real-world embedded deployment.

Beyond raw metrics, the results have three important implications. First, preview plus constraints matters forecasting a very short window of thermal dynamics is sufficient to predict actuation and provide smoother temperature trajectories, resulting in less parasitic energy, and a lower risk of limit violations. Second, problem shaping is as important as solver choice: a small amount of model reduction and horizon management could yield most of the benefits of optimality at a fraction of the computational cost. Third, embedded feasibility is possible with automotive-grade processors today, when the model predictive controller is designed for predictable bounded execution time.

This paper also has limitations. The surrogate model is accurate for the data used for training and validation, but its accuracy may diminish in operating regimes not part of that data (e.g., fast charge while atypical, extreme climates, or a parameter drift associated with aging). Likewise, the scenario set is aimed at representing variability which better captures performance than guaranteeing worst-case performance—all of which means a risk remains of extreme compound disturbances that may decrease performance.

Future Work

- Controller Testing on Hardware and Road: Transfer the controller to an automotive ECU and test closed-loop operation under calibrated thermal events (e.g. hill climbs, towing and fast charging) for measuring real-time margins and diagnostic reliability.
- Degradation-Aware Objectives: Incorporate battery aging terms (e.g., temperature-time or temperature-gradient penalties) into the cost function to directly trade-off immediate energy consumption against long-term health.
- Richer uncertainty management: Automate scenario selection by using online clustering or distributional shift detection, so the controller can adapt its scenario set to current operation conditions without the overhead of additional computations.
- Multi-Domain Co-Optimization: Expand the framework to jointly manage battery, cabin HVAC, and powertrain loads while enforcing concurrent thermal and power constraints, while satisfying driver comfort.
- Fail-Safe/Graceful Degradation: Incorporate a deterministic fallback (e.g., calibrated rule set) with provable satisfaction of constraints to manage safety in case of solver non-convergence or sensor failure.

To summarize, SS-MPC primarily minimizes much of the difference between the predicted control with a theoretically optimal predictive controller and the actual computational costs to establish the predictive control in the vehicle. In conjunction with future EVs, SS-MPC presents itself as a greatly attractive option for next-generation EV thermal controllers that must consider efficiency, safety, and deploy in a harmonious manner.

References

[1]. K. Li, T. Q. Dinh, and K. Yao, “Computationally efficient model predictive control for electric vehicle battery thermal management,” in *Proc. 2024 IEEE Vehicle Power and Propulsion Conf. (VPPC)*, Washington, DC, USA, Oct. 2024, pp. 1–6,https://doi.org/http:s//doi.org/10.1109/VPPC63154.2024.10755355

[2]. K. Nam and C. Ahn, “Energy-Efficient Battery Thermal Management in Electric Vehicles Using Artificial-Neural-Network-Based Model Predictive Control,” *World Electric Vehicle Journal*, vol. 16, no. 5, p. 279, 2025. https://doi.org/10.3390/wevj16050279

[3]. M. R. Amini, J. Sun, and I. Kolmanovsky, “Two-Layer Model Predictive Battery Thermal and Energy Management Optimization for Connected and Automated Electric Vehicles,” *arXiv*:1809.10002, 2018. https://doi.org/10.48550/arXiv.1809.10002

[4]. Q. Hu, M. R. Amini, I. Kolmanovsky, J. Sun, A. Wiese, and J. B. Seeds, “Multihorizon Model Predictive Control: An Application to Integrated Power and Thermal Management of Connected Hybrid Electric Vehicles,” *IEEE Trans. Control Syst. Technol.*, vol. 30, no. 3, pp. 1052–1064, May 2022. https://doi.org/10.1109/TCST.2021.3091887

[5]. F. Micheli and J. Lygeros, “Scenario-based Stochastic MPC for systems with uncertain dynamics,” in *Proc. 2022 Eur. Control Conf. (ECC)*, London, United Kingdom, 2022, pp. 833–838, https://doi.org/10.23919/ECC55457.2022.9838080

[6]. S. Zhao and C. C. Mi, “A Two-Stage Real-Time Optimized EV Battery Cooling Control Based on Hierarchical and Iterative Dynamic Programming and MPC,” *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 8, pp. 11677–11687, Aug. 2022, https://doi.org/10.1109/TITS.2021.3106253

[7]. S. Park and C. Ahn, "Computationally Efficient Stochastic Model Predictive Controller for Battery Thermal Management of Electric Vehicle," in *IEEE Transactions on Vehicular Technology*, vol. 69, no. 8, pp. 8407-8419, Aug. 2020, https://doi.org/10.1109/TVT.2020.2999939

[8]. Q. Hu, P. Ding, W. Jiang, and K. Fung, “Enhancing Battery Thermal Management in Electric Vehicles through Reduced Order Modeling and Predictive Control for Quick Charging,” *SAE Technical Paper* 2024-01- Connected HEVs via Multi-Horizon MPC,” https://doi.org/10.48550/arXiv.2003.08855

[9]. D. Li, Q. Hu, W. Jiang, H. Dong, and Z. Song, “Integrated Power and Thermal Management for Enhancing Energy Efficiency and Battery Life in Connected and Automated Electric Vehicles,” https://doi.org/10.48550/arXiv.2411.05298

[10]. Y. Wu, Z. Huang, D. Li, H. Li, J. Peng, D. Stroe, and Z. Song, “Optimal Battery Thermal Management for Electric Vehicles with Battery Degradation Minimization,” https://doi.org/10.48550/arXiv.2308.03056