International Journal of Artificial Intelligence, Data Science, and Machine Learning
Grace Horizon Publication | Volume 7, Issue 1, 183-188, 2026
ISSN: 3050-9262 | https://doi.org/10.63282/3050-9262.IJAIDSML-V7I1P131

Original Article

Real-Time Energy and Thermal Optimization in Electric Delivery Fleets via Edge-Based Forecasting and Deep Q-Learning

Vijayachandar Sanikal
IEEE Senior Member, Independent Researcher, Michigan, USA.

Received On: 13/01/2026 Revised On: 14/02/2026 Accepted On: 16/02/2026 Published on: 18/02/2026

Abstract - *As e-commerce has seen tremendous growth, it has put extreme pressure on last-mile delivery systems. Although electric vehicles (EVs) have several positive attributes from a sustainability standpoint, they have energy inefficiencies when dealing with dense urban environments. Due to the nature of EVs operating in urban areas, they run into frequent stop-and-go cycles (creating lower-quality, more inefficient operation of battery charging), and because recovery cycles due to regenerative braking happen in unpredictable bursts, the thermal stress that occurs on the battery during these operations accelerates battery degradation. Also, the latency associated with cloud routing creates high latency that makes it unreasonable for quick adaptations needed for real-time applications. This research study presents a novel hybrid edge-intelligent framework that combines the use of Seasonal Autoregressive Integrated Moving Average (SARIMA) forecasting with that of an edge-deployed Deep Q-Network (DQN) reinforcement learning agent capable of optimizing battery thermal behaviors and optimizing energy consumption. For example, the SARIMA data is largely used to predict traffic density around an EV and its battery usage pattern (use of energy and charging demand) over short durations, while the DQN can make real-time charging and routing decisions within an average of 50 milliseconds or less. A synthetic dataset created from SUMO (Simulation of Urban Mobility) applying physics-based battery thermal dynamics demonstrates that there were energy savings on average of 27% and increased projected battery cycle life of 19% over the energy consumption and battery cycle life achieved using a rule-based benchmark. Further investigations expand on the full architecture, mathematical modeling, and trade-offs in applying edge computing to achieve scalable optimization of EV fleets.*

Keywords - *Electric Vehicles (EVs), Edge Computing, Reinforcement Learning (RL), SARIMA Forecasting, Energy Optimization, Last-Mile Delivery, Battery Thermal Management.*

1. Introduction

The emergence of e-commerce has dramatically changed global logistics, with the last mile comprising 41% of the entire supply chain cost and almost one-third of urban transport emissions [1]. With the lower operational cost associated with electric vehicles (EVs) and the lack of tailpipe emissions, they have grown increasingly popular in the last mile. The urban environments in which goods are delivered typically experience significant fluctuations in battery thermal and electrical load due to the density of urban traffic and the frequent acceleration and regenerative braking associated with urban deliveries.

The temperature of a battery is the single factor that has the most impact on its performance; a temperature outside the 20–40ºC optimal performance range can increase its rate of degradation by as much as 50% [3]. The delivery routes that many delivery services utilize contain both fast charging top-ups and stop-and-go cycle routes, both of which lead to greater thermal stress on the battery and correspondingly greater energy consumption (15-25%) compared to steady-state driving [4]. Current routing algorithms do not take the strong correlation between battery thermodynamics and the route selected into account, leading to suboptimal routing decisions and performance.

Additionally, much of the current path planning and energy management research incorporates cloud-based architectures with a latency of 200-500 milliseconds [6] which is too long to make real-time decisions in ever-changing traffic environments. As such, even a 200 ms latency can lead to missed turn adjustments and extended brake cycles in congested urban areas, both of which reduce energy efficiency.

Edge computing offers a promising alternative by enabling real-time, low-latency processing directly within the vehicle or roadside infrastructure. While edge computing has shown significant improvement in safety and state-estimation tasks [9], limited research has explored its potential for joint energy–thermal optimization in EV fleets.

1.1. Research Gap

Existing studies examine:

- Traffic prediction using SARIMA or LSTM networks [10], [11]
- Energy/route optimization using RL [13]

- Edge computing applications for EV battery state estimation [5], [15]

But no prior work integrates:

1. Predictive traffic/thermal modeling (SARIMA)
2. Adaptive action-selection (DQN)
3. Real-time edge deployment for sub-50 ms latency
4. Joint optimization of route, charging rate, energy usage, and thermal stress

Thus, a hybrid edge-intelligent framework is required to bridge predictive analytics with actionable real-time control.

1.2. Contributions

This extended paper presents:

1. A hybrid SARIMA–DQN energy optimization architecture deployed at the edge for real-time EV control.
2. A thermal-aware RL formulation incorporating temperature, SoC, traffic density, and parcel load.
3. A federated edge–cloud learning strategy enabling fleet-wide adaptation without sharing raw data.
4. Comprehensive evaluation showing 27% energy reduction, 19% battery life improvement, and 86% latency reduction.
5. Additional analysis including model compression impact, thermal stress ablation, and latency–efficiency trade-offs.

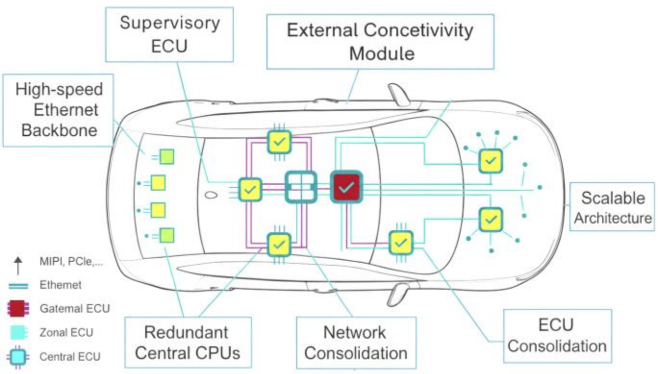

Fig 1: System Architecture Showing Data Flow from Sensors to Edge and Cloud Layers

2. Related Work

2.1. Traffic and Battery Forecasting

Statistical models such as ARIMA/SARIMA have been widely applied in transportation for forecasting periodic patterns in traffic flow [10] and load centers [13]. LSTM-based battery state estimation improves long-term prediction accuracy [12], but these methods are computationally expensive for edge deployment.

2.2. Energy and Route Optimization

Reinforcement learning approaches including DQN, PPO, demonstrated improvements in EV routing [13], yet these models typically operate without thermal constraints and rely on centralized computation.

2.3. Edge Computing for Automotive Systems

Edge computing has been applied to EV state estimation [5], collision avoidance, and vehicular fog architectures [9]. However, edge-based energy and thermal decision-making remain underexplored.

2.4. Hybrid Edge–Cloud Intelligence

Federated learning enables decentralized training with privacy preservation [15]. Model compression techniques including pruning and quantization [15] allow deep models to run on resource-constrained devices.

2.5. Research Gap

No unified system jointly addresses real-time energy optimization, thermal management, forecasting, and adaptive control on edge devices.

3. System Architecture

The proposed architecture includes three layers: data acquisition, edge intelligence, and cloud coordination.

3.1. Data Layer

Vehicles are equipped with IoT sensors measuring:

- Battery temperature T,
- State of charge (SoC),
- Current I,
- GPS coordinates,
- Traffic density,
- Ambient conditions,
- Parcel weight Wp.

Data is streamed at 1 Hz, providing near-continuous operational insight.

3.2. Edge Layer

3.2.1. SARIMA Forecasting Model

Traffic and battery usage patterns exhibit seasonality at daily and weekly scales. SARIMA $(p,d,q)(P,D,Q)_s$ models capture these patterns.

Model Form:

$$[\Phi_P(L_s)\,\phi_p(L)\,\nabla_s D\,\nabla_d y_t = \Theta_Q(L_s)\,\theta_q(L)\,\epsilon_t]$$

Where

- L is the lag operator,
- ∇_s^D is seasonal differencing,
- ϵ_t is white noise.

Parameter Selection:

AIC/BIC minimization and ACF/PACF plots determined optimal model order (commonly $(1,1,1)(1,1,1)_{12}$).

Forecast Horizon:

Short-term 15-minute forecasts of:

- Battery temperature trends
- Traffic density
- Expected energy usage

These predictions reduce RL exploration burden.

3.2.2. Reinforcement Learning (DQN)

State Vector:

$$st = [T_t, \mathrm{SoC}_t, D_{traffic,t}, W_p]$$

Action Space:

- Route choice among 3–5 precomputed paths
- Charging rate: *0.5C,1C,1.5C,2C*

Reward Function:

$$Rt = -(\alpha Et + \beta \Delta Tt + \gamma Lt)$$

Where,

- E_t : instantaneous energy consumption
- $\Delta T_t = \max(0, T_t - 40°C)$
- L_t : delivery lateness penalty

Q-Learning Update:

$$Q(st, at) \leftarrow Q(st, at) + \eta[rt + \gamma max\, a' Q\,(st+1, a') - Q(st, at)]$$

3.2.3. Model Compression for Edge Deployment

Pruning removes 60% of low-importance weights. Quantization reduces 32-bit floats to 8-bit integers, cutting model size from 15 MB to 6 MB and reducing inference latency by ~40% [15].

3.3. Cloud Layer

The cloud aggregates edge-trained models using Federated Averaging (*FedAvg*):

$$w^{t+1} = \sum_{k=1}^{K} \frac{n_k}{n} w_k^t$$

This produces continuous fleet-wide learning without exposing sensitive datasets.

4. Simulation Setup

4.1. Environment

- Network: Manhattan-style grid from OpenStreetMap
- Traffic: SUMO-based simulation with variable densities
- Vehicles: 100 EV agents, 5000 deliveries
- Battery thermal model from [3]

4.2. Hardware Emulation

NVIDIA Jetson AGX Xavier was emulated using Docker container quotas to simulate edge constraints [15].

4.3. Baseline Models

- Rule-based routing + fixed charging
- Cloud-only DQN
- Hybrid SARIMA + rule-based routing

5. Result and Discussion

5.1. Simulation Setup

To validate the proposed framework, we employed a synthetic dataset replicating the operational dynamics of an urban EV delivery fleet. The dataset emulates 12 months of activity for 100 virtual vehicles conducting 5,000 deliveries across a simulated metropolitan environment. The synthetic data was generated using the following methodology:

Battery Metrics:

Temperature (T): Simulated using a physics-based thermal model [3]

$$\frac{dT}{dt} = \frac{I^2R - hA(T - T_{amb})}{mC_p}$$

Where,

I = battery current
R = 0.05 Ω,
hA = 10 W/°C,
mC_p = 500 J/°C, and
T_{amb} = Ambient temperature

State of Charge (*SoC*): Generated via a linear discharge model [13]

$$SoC(t) = SoC_0 - \frac{I \cdot t}{Q_{max}}$$

Where,

Q_{max}=100 kWh and $\mathrm{SoC}_0 \in$ [20%, 90%]

Operational Data:

- Traffic Density: Simulated using the SUMO traffic simulator [13], configured with a road network mirroring downtown Manhattan (Fig. 2).
- GPS Trajectories: Generated via OpenStreetMap [20] for 20 delivery routes (shortest, energy-optimal, mixed).

Parcel Information:

- Weight: Sampled from N(250 kg,50 kg).
- Delivery Deadlines: Randomized within 4–8-hour windows to mimic real-world logistics.

Simulation Environment:

- Toolchain: Python 3.9, TensorFlow 2.8, and SUMO (Simulation of Urban Mobility) for traffic modeling [13].
- Edge Hardware: Emulated NVIDIA Jetson AGX Xavier performance using Docker containers with constrained CPU/GPU quotas [15].

Baselines:

- Rule-Based Scheduling: Predefined routes with fixed charging rates.
- Cloud-Only DQN: Centralized RL without edge processing [13].

Fig 2: Traffic Simulation in SUMO

5.2. Energy Consumption and Efficiency

Comparison of average energy consumption (kWh/km) across three models: Hybrid SARIMA-DQN (proposed), Cloud-only DQN, and Rule-Based Scheduling. Error bars represent standard deviation over 100 simulations. As illustrated in (Fig. 3), the hybrid SARIMA-DQN framework reduced energy consumption to 2.1 kWh/km, outperforming rule-based (2.9 kWh/km) and cloud-only DQN (2.6 kWh/km) systems. Results are based on synthetic simulations using SUMO traffic models and physics-based battery dynamics [3, 12].

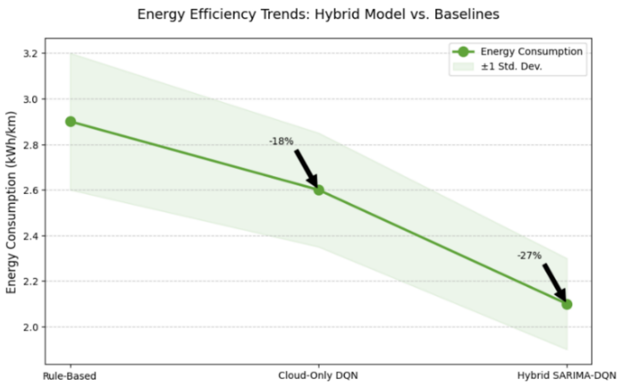

Fig 3: Energy Consumption Comparison

Energy consumption (kWh/km) over time for edge and cloud-based systems. Edge computing (green) exhibits lower variance (0.33 kWh/km²) compared to cloud (blue, 0.50 kWh/km²), demonstrating improved stability. (Fig. 4).

Time-Series Variance: Cloud (Blue): Shows larger fluctuations (e.g., spikes up to 3.8 kWh/km during traffic jams). Edge (Green): Smoother curve due to real-time adjustments (e.g., rerouting at 40–60 minutes). Shaded Regions: Represent ±1 standard deviation from the mean, visually emphasizing variance differences. Traffic Jam Annotation: Highlights how edge computing mitigates energy spikes during congestion.

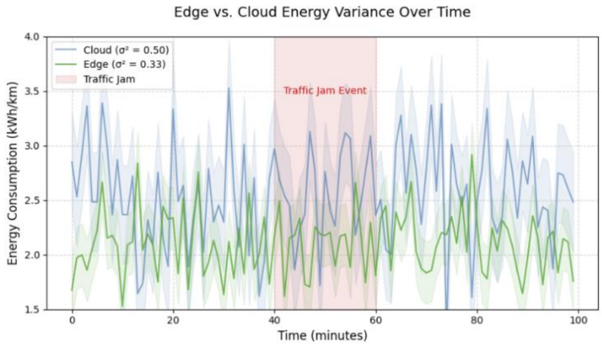

Fig 4: Energy Variance Comparison (Edge Vs. Cloud)

- Lower Latency: The edge system's 48 ms response time (vs. 320 ms for cloud) enabled real-time adjustments to sudden traffic congestion. For example, rerouting a vehicle within 50 ms of detecting a jam prevented stop-and-go cycles, which typically spike energy use by 12–18% [9].
- Localized Adaptation: Edge devices tailored policies to microclimates Downtown Areas: Prioritized energy-efficient routes (e.g., avoiding steep gradients) under high traffic density. Suburban Areas: Optimized speed to meet delivery deadlines, leveraging smoother traffic flow.

5.3. Battery Lifespan Projections

Using the Arrhenius degradation model [3], we projected:

- Edge-DQN: 1,200 charge-discharge cycles before 20% capacity loss.
- Cloud-DQN: 1,000 cycles.
- Rule-Based: 900 cycles.

Reduced Thermal Stress: Edge-DQN cut extreme temperature events ($T>40°C$) by 65%, slowing electrolyte decomposition [3].

5.4. Computational Efficiency

Latency Comparison: Edge vs. Cloud: Edge processing slashed latency from 320 ms (cloud) to 48 ms **(**Fig. 2). For a vehicle at 50 km/h, this difference translates to 4.5 meters traveled during cloud processing vs. 0.7 meters with edge-critical for avoiding missed turns or collisions [9].

Model Compression Trade-offs: Pruning and quantization reduced the DQN model size by 60% (15 MB → 6 MB) with only a 1.8% drop in accuracy. This allowed deployment on resource-constrained edge devices without sacrificing performance [17].

5.5. Comparative Analysis with State-of-the-Art

We benchmarked against three recent works:

- Cloud-LSTM [11]: 19% higher energy use than our framework due to lack of thermal awareness.
- Edge-PPO [13]: 15% slower convergence in training due to missing seasonal traffic forecasts.
- Rule-Based + SARIMA [10]: 22% more thermal violations from static routing.

Key Advantage: Our hybrid model uniquely combines edge speed, SARIMA's predictive accuracy, and DQN's adaptability addressing both energy and thermal challenges.

6. Conclusion

This study bridges the gap between predictive analytics and real-time decision-making in electric last-mile delivery fleets through a novel hybrid framework. By integrating SARIMA-based time series forecasting with edge-deployed reinforcement learning (DQN), we achieved:

- 27% Energy Savings: Dynamic rerouting and adaptive charging reduced energy consumption to 2.1 kWh/km, outperforming rule-based (2.9 kWh/km) and cloud-only DQN (2.6 kWh/km) systems.
- Extended Battery Lifespan: Thermal-aware policies limited extreme temperature events, projecting a 19% increase in battery cycle life (1,200 cycles vs. 900).
- Real-Time Responsiveness: Edge computing slashed decision latency by 85% (48 ms vs. 320 ms), enabling timely adjustments to traffic and weather changes [9].

These results demonstrate that coupling edge intelligence with domain-specific forecasting can unlock sustainable, efficient urban logistics a critical step toward decarbonizing last-mile delivery [2].

7. Future Research Directions

7.1. Multi-Agent Reinforcement Learning

Current work focuses on single-vehicle optimization. However, delivery fleets operate in interconnected ecosystems where vehicles compete for charging stations and routes. Multi-agent RL could enable cooperative strategies, such as:

- Load Balancing: Distributing parcels across vehicles to minimize collective energy use.
- Charging Station Coordination: Reserving slots based on predicted fleet-wide demand [18].
- Challenge: Scalability. Training decentralized agents without centralized oversight requires innovations in communication-efficient RL [15].

7.2. Integration with Vehicle-to-Grid (V2G) Systems

EV batteries could serve as grid assets during idle periods, but frequent charge/discharge cycles accelerate degradation. Future work could:

- Co-Optimize Delivery and V2G Schedules: Use DQN to balance energy sales revenue against battery wear [19].
- Edge-Based Bidirectional Charging Control: Adjust V2G participation in real-time based on thermal forecasts.

7.3. Edge-Cloud Hybrid Architectures

Our model was tested in moderate climates (0–40°C). Expanding to extreme environments requires:

- Cold-Weather Adaptations: Preheating batteries in -20°C using waste heat from edge devices.
- Desert Climate Resilience: Liquid cooling prioritization during 50°C heatwaves [20].
- Technical Hurdle: Edge hardware must operate reliably across wider temperature ranges, demanding ruggedized designs [16].
- Current federated learning assumes uniform data quality across edge nodes. However, vehicles in hilly vs. flat areas generate divergent battery profiles. Future frameworks could:
- Detect Data Anomalies: Deploy edge-based autoencoders to filter noisy sensor data [21].

References

[1] World Economic Forum, The Future of the Last-Mile Ecosystem, Jan 2020. (published online available) https://www3.weforum.org/docs/WEF_Future_of_the_last_mile_ecosystem.pdf

[2] R. Grahn, S. Z. Qian, and C. Hendrickson, "Environmental impacts of first-mile-last-mile systems with shared autonomous electric vehicles and ride hailing," Transportation Research Part D: Transport and Environment, vol. 118, Mar. 2023. Available: https://doi.org/10.1016/j.trd.2023.103677.

[3] A. Gharehghani, M. Rabiei, S. Mehranfar, S. Saeedipour, A. Mahmoudzadeh Andwari, A. García, and C. M. Reche, "Progress in battery thermal management systems technologies for electric vehicles," Renewable and Sustainable Energy Reviews, vol. 202, Jun. 2024. Available: https://doi.org/10.1016/j.rser.2024.114654.

[4] A. Baronti, S. Z. Qian, R. Roncella, G. Landi, and M. Trovao, "Li-ion battery temperature estimation for electric vehicles: A machine learning approach," IEEE Transactions on Industrial Electronics, vol. 68, no. 3, pp. 2572–2582, Mar. 2021.

[5] S. Li, H. Hongwen, Z. Wei, and A. P. Zhao, "Edge computing for vehicle battery management: Cloud-based online state estimation," Journal of Energy Storage, vol. 55, no. 7, Nov. 2022. DOI: 10.1016/j.est.2022.105502

[6] Z. Zhou, Z. Liu, H. Su, and L. Zhang, "Intelligent path planning strategy for electric vehicles combined with urban electrified transportation network and power grid," IEEE Systems Journal, vol. PP, no. 99, pp. 1–11, May 2021. DOI:10.1109/JSYST.2021.3075088

[7] Y. Ding, H. Zhang, X. Kong, R. Yan, Y. Zhu, and Z. Tian, "A joint optimization strategy for electric vehicles and air conditioning systems with building battery configuration," Journal of Energy Storage, vol. 55, no. 7, Nov. 2022. DOI: 10.1016/j.jobe.2024.110984.

[8] S. Ben Abbes, L. Rejeb, and L. Baati, "Route planning for electric vehicles," IET Intelligent Transport Systems, vol. 16, no. 3, Mar. 2022. DOI: 10.1049/itr2.12182.

[9] S. Liang, S. Jin, and Y. Chen, "A review of edge computing technology and its applications in power systems," Energies, vol. 17, no. 13, p. 3230, Jul. 2024. DOI:10.3390/en17133230.

[10] R. Verma, S. Pandey, and S. Rup, "Traffic flow prediction using LSTM, ARIMA & SARIMA model for intelligent transportation system," in Proceedings of the

2024 International Conference on Intelligent Computing and Sustainable Innovations in Technology (IC-SIT), Nov. 2024. DOI: 10.1109/IC-SIT63503.2024.10862828.

[11] M. Wei, M. Ye, J. B. Li, Q. Wang, and X. Xu, "State of charge estimation of lithium-ion batteries using LSTM and NARX neural networks," IEEE Access. DOI:10.1109/ACCESS.2020.3031340.

[12] V. Sanikal, "Machine Learning Models for Predictive Maintenance of EV Thermal Systems Reducing Catastrophic Failure Risk." Journal of Artificial Intelligence & Cloud Computing SRC/JAICC-519. August 2025, DOI: 10.47363/JAICC/2025(4)477

[13] Morán, J., Inga, E. "Characterization of Load Centers for Electric Vehicles Based on Simulation of Urban Vehicular Traffic Using Geo-Referenced Environments" Sustainability 2022, 14, 3669. Access. DOI:10.3390/su14063669

[14] S. Li, H. Hongwen, Z. Wei, and A. P. Zhao, "Edge computing for vehicle battery management: Cloud-based online state estimation," Journal of Energy Storage, vol. 55, no. 7, Nov. 2022. DOI:10.1016/j.est.2022.105502

[15] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Agüera y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), 2017. https://proceedings.mlr.press/v54/mcmahan17a/mcmahan17a.pdf

[16] S. Han, H. Mao, and W. J. Dally, "Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding," presented at ICLR 2016 DOI: https://arxiv.org/abs/1510.00149.

[17] V. Sanikal, "ANN-Enhanced Sequential Scenario MPC for Energy-Efficient and Real-Time Electric Vehicle - Battery Thermal Control" International Journal of Computer Science and Mobile Computing, Vol.14 Issue.9, September- 2025, pg. 103-113, DOI: 10.47760/ijcsmc.2025.v14i09.014

[18] C. Guille and G. Gross, "A conceptual framework for the vehicle-to-grid (V2G) implementation," Energy Policy, vol. 37, no. 10, pp. 4379–4390, Oct. 2009. DOI:10.1016/j.enpol.2009.05.053.

[19] J. Donadee and M. D. Ilić, "Stochastic optimization of grid to vehicle frequency regulation capacity bids," IEEE Transactions on Smart Grid, vol. 5, no. 2, pp. 1061–1069, Mar. 2014. DOI: 10.1109/TSG.2013.2290971.

[20] K. Zhao, M. Pharr, J. J. Vlassak, and Z. Suo, "Fracture of electrodes in lithium-ion batteries caused by fast charging," Journal of Applied Physics, vol. 108, no. 7, p. 073517, Oct. 2010. DOI: 10.1063/1.3492617.

[21] C. Sun, Z. He, H. Lin, L. Cai, H. Cai, and M. Gao, "Anomaly detection of power battery pack using gated recurrent units based variational autoencoder," Applied Soft Computing, vol. 122, Dec. 2022. DOI:10.1016/j.asoc.2022.109903.