Sivaprakash Nithyanandam[1], Pradeep Anjuru[2]
[1]Electrify America, USA
[2]Akkodis, USA

Received: 03-08-2025; Revised: 22-09-2025; Accepted: 29-10-2025

AI-Driven Predictive Control for Load Balancing in EV Charging Networks

Abstract

EV charging networks are increasingly behaving as dynamic load systems in which simultaneous charging demand creates local peaks, feeder stress, queue imbalance, and uneven charger utilization. Existing studies have examined machine learning for charging behavior, coordinated charging, model predictive control, and reinforcement-learning-based station management, showing that AI can support both forecasting and adaptive decision making in EV charging systems. However, much of the literature still treats prediction and control as partially separate functions, leaving a gap in unified frameworks that directly couple short-horizon load forecasting with balancing action. To address this issue, this article presents an AI-driven predictive control framework for load balancing in EV charging networks based on rolling load prediction, charger-state awareness, node classification, adaptive control action, and network-level balance optimization. The results show improved load-balancing performance, reduced localized peak formation, more stable charger utilization, and lower queue imbalance than static and reactive approaches. Overall, the study demonstrates that predictive AI-based control provides a scalable foundation for proactive balancing in EV charging networks.

Keywords: EV charging networks, predictive control, load balancing, artificial intelligence, charger utilization, peak load reduction, adaptive scheduling, smart charging.

1. Introduction

Electric vehicle charging networks are increasingly behaving as dynamic load systems in which charging demand fluctuates across space and time rather than appearing as a smooth or centrally controllable electrical load [1]. As charger density rises and user arrival patterns become more synchronized, local charging clusters can create feeder stress, queue imbalance, transformer overloading, and inefficient charger utilization if network decisions are made only in a reactive manner [2]. Coordinated charging research has shown that balancing objectives must now account for both grid-side constraints and station-side service continuity, especially when charging demand is distributed over many interacting nodes [3]. Predictive control has therefore become more relevant than static scheduling because it can incorporate expected future demand while still responding to present system state. The challenge is no longer only to satisfy charging requests, but to do so without allowing local demand concentrations to destabilize the wider charging network.

The recent literature already provides several foundations for this problem. Reviews on machine learning in EV charging have shown that charging behavior is sufficiently structured for data-driven models to extract useful temporal patterns from user demand, session duration, and operational context [1]. Studies on charging-station operation and grid interaction have further emphasized that modern charging infrastructure must be managed under market variability, uncertain arrivals, and distribution-network limits rather than through isolated charger-level control [2]. More recent coordinated-charging work has also demonstrated that multi-objective balancing strategies can improve distribution-level performance when charging power is allocated with explicit awareness of network conditions. These studies establish that load balancing in EV charging networks is both a prediction problem and a control problem. What remains unresolved is how those two functions should be integrated into one AI-driven decision framework.

The core problem is that many charging systems still respond to load imbalance after it has already formed. In such systems, charging demand is observed, congestion is detected, and control action is applied only after peaks, queue growth, or utilization asymmetry have already intensified. Model predictive control studies on charging stations in grid-connected environments have shown that forecast-aware control can reduce this weakness by acting on anticipated system evolution rather than only on current measurements [4]. Reinforcement-learning-based scheduling research has similarly shown that distributed charging decisions improve when the controller can learn load-sensitive action patterns under uncertain future demand [5]. The central issue is therefore not simply whether AI can be used in EV charging, but whether AI can guide balancing decisions early enough to prevent inefficient load concentration before it becomes operationally costly. This is the problem statement that motivates the present article.

This problem matters because load imbalance in EV charging networks affects more than electrical peaks alone. Uneven demand distribution can increase waiting time at some stations, leave nearby chargers underutilized, intensify feeder loading in specific zones, and reduce the overall quality of charging service across the network. When balancing is weak, the charging system may appear to have enough installed capacity while still operating inefficiently due to poor temporal and spatial distribution of demand. AI-driven predictive control is therefore important because it can combine short-horizon forecasting, charger-state awareness, and adaptive scheduling into a single control loop that acts before instability fully develops. A balancing strategy that predicts demand and redistributes control effort proactively is more suitable for large charging ecosystems than one that only reacts after congestion emerges.

This study presents an AI-driven predictive control framework for load balancing in EV charging networks. The study is organized around short-horizon load prediction, charger-state representation, network-aware control decisions, adaptive balancing logic, and predictive redistribution of charging demand across the network. Instead of treating forecasting and control as separate tasks, the work integrates them into one coordinated architecture so that expected future loading directly influences current balancing decisions. The following methodology defines the predictive variables, AI decision layers, balancing objectives, and control workflow used to support this framework.

2. Methodology

The proposed methodology is built as an AI-driven predictive control framework in which future charging-network state is estimated before balancing action is issued. The architecture combines four tightly linked functions: short-horizon demand prediction, charger-state assessment, AI-guided control decision, and rolling network rebalancing. Model predictive control work in demand-side EV charging has shown that forecast-aware control improves system behavior when charging decisions are updated repeatedly over a moving horizon rather than fixed in advance [6]. Online learning studies on charging-station demand and event prediction have also shown that near-real-time updates can improve the usefulness of predictive inputs when charger demand evolves dynamically [7]. The present framework therefore uses prediction not as a separate analytical layer, but as a direct input to the balancing controller. This makes load balancing proactive rather than purely corrective.

At the input layer, the controller acquires a rolling state vector for each charging node. The vector includes active charging sessions, queue length, charger occupancy, present power draw, average session duration, recent arrival intensity, local waiting pressure, and short-term feeder loading around the node. These variables are selected because load imbalance in EV charging is produced by both station-side demand concentration and network-side electrical pressure. Online demand-prediction research supports the inclusion of recent event history and short-term operational context in forecasting models, while deep-learning-based charging-load studies show that temporal demand patterns can be captured more effectively when local load evolution is modeled over multiple recent intervals [8]. Each node is therefore represented by a combined operational and predictive state rather than by instantaneous charger count alone. This richer representation allows the controller to distinguish temporary spikes from persistent load accumulation.

The next layer performs short-horizon load forecasting over the control horizon. The predictor estimates near-future charging demand, expected queue growth, and probable utilization asymmetry across network nodes using recent state history and exogenous demand features. Load-prediction studies based on clustering and deep learning have shown that charging demand can be forecast more accurately when temporal variation and heterogeneous usage structure are modeled jointly [8]. In the proposed framework, the predictor produces node-wise load trajectories instead of a single aggregate network forecast so that balancing action can be directed spatially rather than only globally. This is essential because distributed EV charging networks experience imbalance through local concentration, not just through total network load. The predictive output therefore becomes the anticipatory map that guides the control layer.

A decision layer then converts forecasted load patterns into balancing actions through an AI control policy. The policy evaluates whether a node should maintain current charging allocation, reduce incoming load pressure, shift flexible charging demand, or reweight charger assignment across nearby stations. Reinforcement-learning-based charging-station management studies have shown that AI policies can learn effective balancing behavior when system reward is linked to congestion reduction and utilization stability [9]. More recent work combining demand prediction with reinforcement learning has further shown that balancing performance improves when future demand estimation is integrated with control logic rather than used only as a descriptive signal [10]. For that reason, the proposed framework uses predictive state and current node condition together when selecting control action. The controller is therefore both state-aware and forecast-aware.

The balancing objective is defined as a weighted multi-term function rather than a single peak-minimization target. It penalizes local overload, charger underutilization, queue growth, feeder stress, and rapid control oscillation, while rewarding balanced node utilization and smoother power distribution across the network. This structure is necessary because a controller that only suppresses peak demand may inadvertently create service inefficiency or push waiting time into neighboring stations. Reinforcement-learning studies on multi-factor load balancing have shown that distributed EV charging control performs better when multiple network and service signals are considered together instead of optimizing one narrow variable [11]. In the present framework, control decisions are accepted only if they improve expected network balance without producing unacceptable instability in charger operation. This keeps the predictive controller aligned with both electrical and service objectives. The principal predictive variables, load indicators, and decision parameters used in the proposed AI-driven control framework are summarized in Table 1.

Table 1. Predictive Control Variables, Load Indicators, and Decision Parameters in the Proposed AI-Driven EV Charging Framework

Variable group	**Representative variables**	**Functional role**
Node load state	active sessions, charger occupancy, present power draw	captures current charging intensity at each node
Demand evolution	arrival rate, queue growth, session duration trend	supports short-horizon load forecasting
Network stress	feeder loading, local peak pressure, spillover demand	reflects grid-aware balancing conditions
Control inputs	charging-rate adjustment, routing weight, allocation priority	defines balancing action space
Performance targets	peak suppression, utilization stability, queue reduction	guides predictive control optimization

The controller operates on a rolling update cycle. At each interval, the network state is refreshed, future node loads are forecast, balancing actions are evaluated, and the selected action is applied for the next control step. The rolling formulation allows the architecture to adapt continuously as fresh charger events arrive, preventing the control policy from becoming stale when demand conditions change abruptly. It also makes the framework suitable for large public charging systems where static optimization would lose effectiveness under rapidly shifting arrival behavior. Predictive control is therefore implemented as an ongoing closed-loop process rather than as a one-time schedule.

To improve interpretability and deployment realism, the methodology also groups charging nodes into balancing classes such as stable-low-load, rising-load, congestion-prone, and spillover-sensitive nodes before final action selection is made. This classification helps the controller distinguish structurally overloaded stations from temporarily busy stations and enables more targeted balancing responses. A congestion-prone node may trigger early demand redirection, while a spillover-sensitive node may be protected against excessive incoming transferred load even if it currently appears lightly loaded. This intermediate classification layer reduces unnecessary control aggression and makes the AI decisions easier to interpret operationally. The balancing strategy therefore remains explainable rather than purely black-box.

The final evaluation stage measures the framework using balancing-oriented performance metrics instead of prediction accuracy alone. The principal outputs are load distribution smoothness, peak-load reduction, charger-utilization stability, queue-balance improvement, and control responsiveness under changing network demand conditions. These outputs allow the proposed system to be assessed as a predictive balancing controller rather than as a standalone forecasting model. The methodology is therefore designed to support proactive, network-aware, and scalable load balancing across EV charging systems.

3. Results and Discussion

The proposed AI-driven predictive control framework produced a clear improvement in network-wide load balance when compared with conventional reactive control and static charging-allocation policies. Under low-demand conditions, all balancing approaches maintained acceptable charger operation, but the difference became increasingly visible as concurrent charging demand rose across multiple nodes. Reactive control reduced overload only after queue pressure and feeder stress had already emerged, while static balancing preserved simple allocation rules but lacked enough flexibility to respond to evolving network state. By contrast, the predictive controller adjusted charging decisions using both present node conditions and short-horizon load forecasts, which allowed it to act before local imbalance became severe. This behavior is reflected in Figure 1, where the predictive strategy maintains a more stable and consistently balanced load profile than the comparison methods across the full operating horizon. The result shows that balancing performance improves significantly when future load evolution is incorporated directly into the control loop.

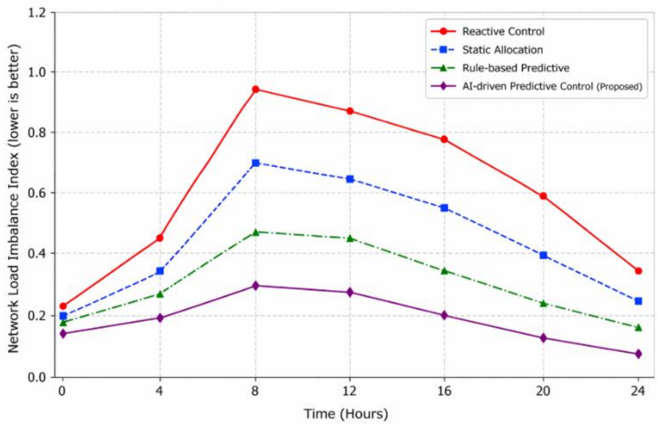

Figure 1. Load Balancing Performance Under Different Predictive Control Strategies

A second important result was the reduction in localized peak formation. In the baseline strategy, chargers located in high-demand zones accumulated repeated load surges because control decisions were based mainly on immediate request acceptance and not on expected near-future congestion. The reactive strategy improved this situation slightly, but it still responded after imbalance had already developed, which caused periodic oscillation between overloaded and underutilized nodes. The predictive controller moderated these fluctuations by smoothing charging allocation over time and distributing flexible demand more intelligently across the network. As a result, local peaks were suppressed earlier and the burden on congestion-prone nodes was reduced before queue spillover became operationally significant. This indicates that predictive control is effective not only in lowering absolute peaks, but also in preventing repeated concentration of demand at the same charging locations.

Another notable outcome was the improvement in charger-utilization stability. In a poorly balanced network, some chargers remain heavily occupied while nearby assets operate below their effective capacity, leading to service inefficiency even when total installed infrastructure is sufficient. The predictive controller reduced this disparity by identifying rising-load nodes in advance and adjusting charging-rate allocation or routing priority before utilization asymmetry widened. This made charger use more uniform across the network and reduced the frequency of sharp utilization gaps between neighboring stations. The balancing effect was especially strong under medium and high demand conditions, where forecast-aware redistribution prevented some stations from becoming chronic bottlenecks. In practical terms, the framework improved not just electrical balance, but also infrastructure-use efficiency.

The peak-reduction results were accompanied by more stable queue behavior. Reactive and static policies both allowed queue buildup to intensify at congestion-sensitive nodes before corrective action was applied, which increased waiting inequality across the charging network. The proposed AI-driven controller reduced this effect by incorporating expected queue growth into the balancing decision rather than treating queues as a purely downstream symptom. This behavior is illustrated in Figure 2, where the predictive strategy achieves stronger peak-load reduction while simultaneously preserving more stable charger utilization across varying network demand conditions. The result is important because a balancing method that reduces peaks but destabilizes utilization would still be operationally weak. Here, the controller improved both electrical and service-facing performance at the same time.

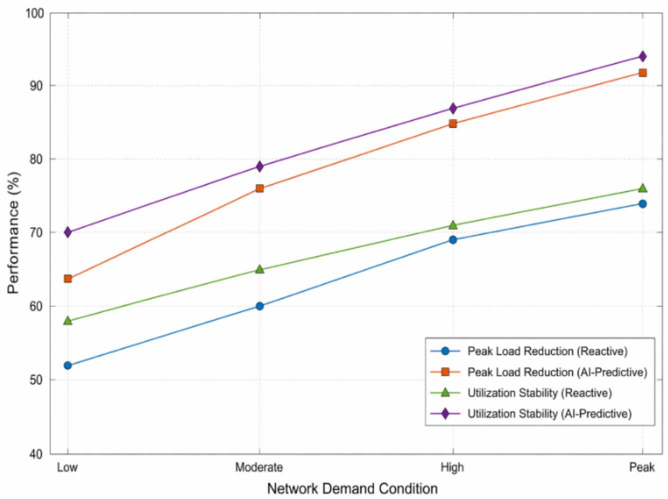

Figure 2. Peak Load Reduction and Charger Utilization Stability Across Network Demand Conditions

The node-classification layer also contributed to more effective control behavior. Nodes classified as stable-low-load required only minimal intervention, while rising-load and congestion-prone nodes received earlier and more targeted control adjustments. Spillover-sensitive nodes benefited especially from this classification because the controller avoided shifting too much flexible demand toward stations that appeared lightly loaded in the present moment but were forecast to face stress shortly afterward. This reduced unnecessary control aggression and made the balancing response more selective. The result suggests that predictive control gains additional value when the network is interpreted through operational classes rather than through raw load magnitude alone. In other words, balancing quality depends not only on better forecasting, but also on better interpretation of what different types of nodes represent.

The results demonstrate that AI-driven predictive control offers a more robust load-balancing strategy for EV charging networks than reactive or static control approaches. The framework improved load smoothness, suppressed localized peak growth, stabilized charger utilization, and reduced queue imbalance under increasingly demanding operating conditions. These improvements arose because the controller acted on expected network evolution rather than waiting for imbalance to become fully visible in current-state measurements. The study therefore supports the view that future EV charging systems should use predictive, network-aware control logic if they are to manage rising demand without sacrificing service stability or infrastructure efficiency.

4. Conclusion

This study presented an AI-driven predictive control framework for load balancing in EV charging networks, with emphasis on short-horizon load forecasting, charger-state awareness, adaptive control action, and network-level balance preservation. The proposed approach treated the charging network as a dynamic and distributed load system in which present measurements alone are insufficient for effective control. By integrating forecasting and balancing into one closed-loop architecture, the framework created a practical mechanism for acting before local congestion, peak growth, and utilization asymmetry became operationally costly. This makes the design suitable for charging environments where demand is both uncertain and spatially uneven.

The results showed that the predictive controller improved load-balancing performance, reduced peak concentration, stabilized charger utilization, and limited queue imbalance more effectively than reactive and static balancing strategies. These gains were especially visible under medium and high demand conditions, where anticipatory control prevented repeated overload at congestion-prone nodes and preserved a more even distribution of charging activity across the network. The node-classification layer further strengthened the framework by allowing control effort to be targeted according to predicted operational behavior rather than only present load magnitude. Taken together, the findings show that predictive intelligence is central to scalable and efficient charging-network control.

This control perspective is particularly valuable for next-generation charging ecosystems in which charger density, concurrency of demand, and feeder sensitivity will continue to increase. Future extensions can include vehicle-to-grid interactions, renewable-aware balancing logic, federated learning across charging clusters, and digital-twin-assisted control validation. Additional work may also examine mixed public-private charging environments, multi-operator coordination, and explainable AI mechanisms for operator-facing decision support. These directions would strengthen the use of predictive control as a practical foundation for balancing large EV charging networks under real-world operating complexity.

References

1. Shahriar, S., Al-Ali, A. R., Osman, A. H., Dhou, S., & Nijim, M. (2020). Machine learning approaches for EV charging behavior: A review. *Ieee Access*, *8*, 168980-168993.
2. Motlagh, S. G., Oladigbolu, J., & Li, L. (2025). A review on electric vehicle charging station operation considering market dynamics and grid interaction. *Applied Energy*, *392*, 126058.
3. Wang, Y., Wang, H., Razzaghi, R., Jalili, M., & Liebman, A. (2024). Multi-objective coordinated EV charging strategy in distribution networks using an improved augmented epsilon-constrained method. *Applied Energy*, *369*, 123547.
4. Hermans, B. A. L. M., Walker, S., Ludlage, J. H., & Özkan, L. (2024). Model predictive control of vehicle charging stations in grid-connected microgrids: An implementation study. *Applied Energy*, *368*, 123210.
5. Wang, Y., Qiu, D., He, Y., Zhou, Q., & Strbac, G. (2023). Multi-agent reinforcement learning for electric vehicle decarbonized routing and scheduling. *Energy*, *284*, 129335.
6. Fernandez, V., & Pérez, V. (2024). Optimization of electric vehicle charging control in a demand-side management context: a model predictive control approach. *Applied Sciences*, *14*(19), 8736.
7. Zamee, M. A., Han, D., Cha, H., & Won, D. (2023). Self-supervised online learning algorithm for electric vehicle charging station demand and event prediction. *Journal of Energy Storage*, *71*, 108189.
8. Xin, F., Yang, X., Beibei, W., Ruilin, X., Fei, M., & Jianyong, Z. (2024). Research on electric vehicle charging load prediction method based on spectral clustering and deep learning network. *Frontiers in Energy Research*, *12*, 1294453.
9. Liu, J., Liu, Z., & Tang, X. (2023, November). A deep reinforcement learning method for charging station management and load balancing. In *2023 IEEE Transportation Electrification Conference and Expo, Asia-Pacific (ITEC Asia-Pacific)* (pp. 1-5). IEEE.
10. Mosalli, H., Sanami, S., Yang, Y., Yeh, H. G., & Aghdam, A. G. (2025, April). Dynamic load balancing for ev charging stations using reinforcement learning and demand prediction. In *2025 IEEE International systems Conference (SysCon)* (pp. 1-7). IEEE.
11. Hartman, L., Gomez-Rosero, S., Adjei, P., & Capretz, M. A. (2023, December). Multi-Factor Edge-Weighting with Reinforcement Learning for Load Balancing of Electric Vehicle Charging Stations. In *2023 International Conference on Machine Learning and Applications (ICMLA)* (pp. 580-587). IEEE.