Energy Optimization in Smart Building Controls: A Predictive Study of Reinforcement Learning Policies

Xia Xia*

Department of Computer Science and Technology, Tsinghua University, School of Aerospace Engineering, Tsinghua University, Beijing, China

Corresponding author: xia.xia@tsinghua.edu.cn

Abstract

The rapid urbanization and modernization of infrastructure have escalated the energy consumption of commercial and residential buildings globally, necessitating advanced control methodologies to balance occupant comfort with energy efficiency. Traditional control mechanisms often fall short in dynamically adapting to complex thermodynamic interactions within multi-zone structures. Reinforcement learning has emerged as a promising alternative, capable of discovering optimal control strategies through continuous interaction with the building environment. However, evaluating and predicting the efficacy of these policies prior to exhaustive simulation or physical deployment remains a significant computational challenge. This paper introduces a novel framework that leverages graph analysis to predict the energy optimization potential of reinforcement learning policies in smart building controls. By representing the physical architecture and thermal dynamics of a building as a complex graph, where nodes denote distinct spatial zones and edges represent thermodynamic relationships, we extract sophisticated structural features from the control policies. These graph-based representations are subsequently utilized to predict energy savings and operational efficiency without the need for prolonged environmental interaction. The analysis demonstrates that incorporating spatial and thermodynamic dependencies significantly enhances the predictive accuracy of policy evaluation models. The findings provide a critical stepping stone for the rapid prototyping and safe deployment of advanced machine learning algorithms in critical building infrastructure, ultimately contributing to broader sustainability goals.

Keywords: Reinforcement Learning; Graph Analysis; Smart Buildings; Energy Optimization; Predictive Modeling

1. Introduction

1.1 Background on Smart Building Energy Management

The continuous expansion of the built environment has established buildings as one of the primary consumers of global energy and a significant contributor to greenhouse gas emissions. Managing this energy consumption while ensuring the safety, productivity, and thermal comfort of occupants is the central objective of heating, ventilation, and air conditioning systems. Historically, these systems have relied on conventional proportional-integral-derivative controllers, which operate based on simple error-correction mechanisms. While robust and easy to implement, these conventional methods treat different zones of a building in isolation, ignoring the complex thermodynamic interactions that occur across walls, ceilings, and open corridors. This isolated approach inevitably leads to suboptimal energy usage, simultaneous heating and cooling of adjacent spaces, and a delayed response to sudden changes in occupancy or external weather conditions. Consequently, researchers have increasingly investigated more advanced, proactive control methodologies [1]. Model predictive control emerged as a highly effective solution, utilizing mathematical models of building physics to forecast future states and optimize control actions over a receding horizon. However, the requirement for highly accurate, bespoke physical models for every individual building renders model predictive control difficult to scale and expensive to maintain in the face of structural or operational changes over time.

1.2 The Role of Reinforcement Learning in Energy Optimization

To address the limitations of model-dependent approaches, the focus of the building control community has shifted toward data-driven methodologies, predominantly reinforcement learning. In a reinforcement learning paradigm, an artificial agent learns to make sequential decisions by interacting with an environment, receiving feedback in the form of rewards or penalties based on the energy consumed and the thermal comfort maintained [2]. Over successive iterations, the agent develops a control policy that maps observed environmental states, such as room temperatures, humidity levels, and external weather forecasts, to optimal control actions, such as adjusting damper positions, chilled water valve states, or variable air volume flow rates. A significant advantage of reinforcement learning is its ability to adapt dynamically to the unique characteristics of a building without requiring explicit thermodynamic equations. Furthermore, advanced algorithms, notably deep reinforcement learning, leverage multi-layer neural networks to handle high-dimensional state spaces characteristic of modern smart buildings equipped with thousands of Internet of Things sensors. Despite these advantages, the application of reinforcement learning in real-world infrastructure is severely hindered by the black-box nature of the learned policies and the massive computational expense required to evaluate their long-term efficacy through exhaustive simulations [3].

1.3 Problem Statement and Research Objectives

The core challenge addressed in this paper is the computationally prohibitive process of evaluating how well a specific reinforcement learning policy will optimize energy consumption before it is fully deployed or exhaustively simulated. Traditional policy evaluation methods require running the agent through hundreds of episodic simulations across diverse weather profiles, which consumes immense computational resources and delays the iterative design process. This research proposes that the structural and relational data inherent in a building's physical layout contains critical predictive signals regarding a policy's effectiveness. The primary objective is to develop and validate a methodology that uses graph analysis to evaluate these reinforcement learning policies. By modeling the smart building as a graph network, we aim to extract localized and global control features from the policy itself, effectively predicting its energy optimization potential through rapid, forward-pass graph analysis rather than full-scale thermodynamic simulation.

2. Literature Review

2.1 Evolution of Building Control Systems

The trajectory of building control systems reflects a broader industrial trend toward automation, data integration, and intelligent decision-making. Early control strategies were purely mechanical or relied on rudimentary relay logic, gradually evolving into the microprocessor-based direct digital control systems that dominate today. Direct digital control allowed for the widespread implementation of rule-based logic and proportional-integral-derivative loops, which significantly improved the stability of indoor environments [4]. However, the inherent limitations of these localized loops became apparent as buildings grew larger and more complex, featuring intertwined mechanical subsystems like chillers, boilers, cooling towers, and air handling units. The advent of building energy management systems provided a centralized platform for monitoring and supervisory control, enabling the implementation of scheduled setbacks and demand response strategies [5]. Despite these advancements, the overarching control logic remained largely reactive. The transition to predictive and optimization-based control systems, such as model predictive control, marked a significant paradigm shift. Researchers demonstrated that integrating weather forecasts, occupancy schedules, and real-time electricity pricing into a cohesive mathematical optimization framework could yield substantial energy savings. Nevertheless, the high cost of formulating, linearizing, and updating the underlying physics-based models for individual buildings prevented widespread commercial adoption, thereby paving the way for model-free, machine learning approaches.

2.2 Reinforcement Learning Applications in Heating and Ventilation

The application of reinforcement learning in building climate control has experienced exponential growth, driven by breakthroughs in algorithm design and computational hardware. Early implementations utilized tabular Q-learning to manage single-zone environments, demonstrating that an agent could independently learn to balance temperature constraints with energy minimization objectives [6]. As researchers tackled larger multi-zone buildings, the state and action spaces exploded in dimensionality, necessitating the integration of deep learning. Deep Q-networks and policy gradient methods, such as Proximal Policy Optimization and Soft Actor-Critic, have since become the standard tools for continuous control tasks in building energy management [7]. These deep reinforcement learning agents are typically trained in simulated environments constructed using software like EnergyPlus or TRNSYS, which provide highly accurate thermodynamic feedback. A critical area of ongoing research involves multi-agent reinforcement learning, where distributed agents control localized zones while communicating with neighboring agents to achieve global building efficiency. While multi-agent systems alleviate the scalability issues of centralized controllers, they introduce complexities regarding convergence, non-stationarity, and communication overhead.

2.3 Graph Analysis in Interconnected Physical Systems

Graph theory and graph neural networks have recently revolutionized the analysis of data structured in non-Euclidean domains. Unlike traditional convolutional neural networks that process grid-like data such as images, graph neural networks operate on nodes and edges, making them exceptionally well-suited for systems characterized by complex topological relationships. In the context of physical infrastructure, graph analysis has been successfully applied to traffic network forecasting, power grid stability assessment, and water distribution management. For smart buildings, the spatial arrangement of rooms and the layout of the ventilation ductwork naturally form a graph structure. Recent studies have begun exploring the use of graph convolutional networks to model the thermal dynamics of buildings, demonstrating that aggregating temperature information from adjacent nodes significantly improves the accuracy of indoor temperature forecasting. However, the utilization of graph analysis to interpret and predict the efficacy of reinforcement learning control policies remains a largely unexplored frontier. This paper bridges that gap by proposing that the inherent quality of a distributed or centralized control policy can be quantified by analyzing how effectively it addresses the structural and thermodynamic dependencies represented within the building graph.

3. Theoretical Framework and System Architecture

3.1 Conceptualizing Smart Buildings as Graphs

To apply graph analysis to policy evaluation, the physical and mechanical structure of the smart building must first be formalized as a mathematical graph. Let the building be defined as a directed, weighted graph where the set of nodes represents the distinct operational entities within the building, such as individual thermal zones, common areas, or specific air handling units [8]. The set of edges represents the physical and thermodynamic interactions between these nodes. Edges can be categorized into multiple types. Spatial edges exist between nodes that share a physical boundary, representing conductive heat transfer through walls and convective heat transfer through open doors or unsealed partitions. Mechanical edges represent the layout of the ductwork and piping, capturing the flow of conditioned air or chilled water from centralized mechanical equipment to the terminal units in each zone [9]. The weight of an edge corresponds to the magnitude of the thermal coupling; for example, a poorly insulated internal wall would possess a higher spatial edge weight than a highly insulated exterior envelope. By structuring the building in this manner, we capture the reality that a cooling action taken in one room inherently impacts the thermal state and energy requirements of adjacent rooms.

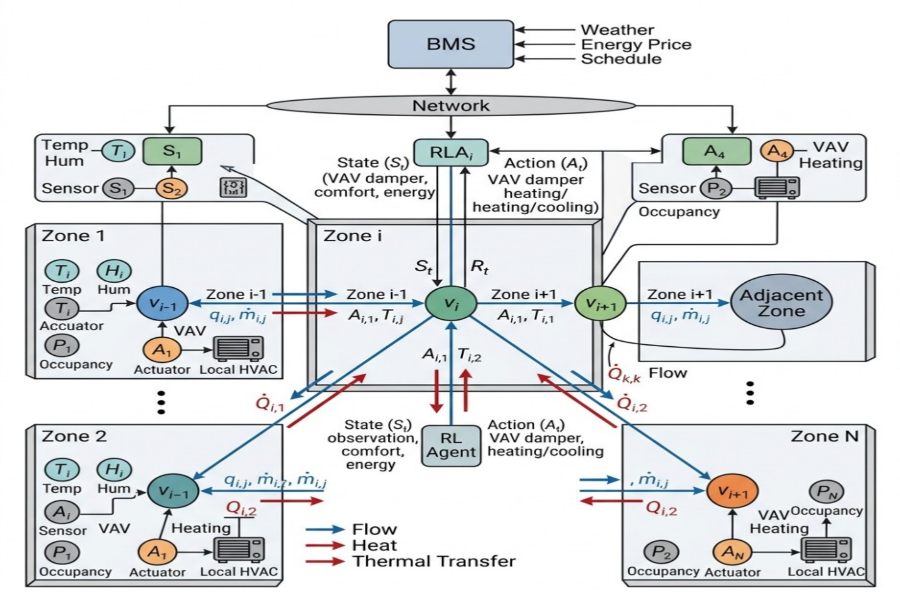

Figure 1: Comprehensive Schematic of Graph Representation in Smart Building Architecture

3.2 Policy Representation in Reinforcement Learning

In a standard reinforcement learning setup, the policy is a function that maps the current state vector of the building to an action vector. For predictive evaluation via graph analysis, we must extract meaningful representations from this policy without running long-term simulations. If the policy is parameterized by a deep neural network, the raw weights and biases are generally too abstract and opaque to be analyzed directly. Instead, we generate a behavioral representation of the policy by sampling a diverse set of representative building states, encompassing various weather conditions and occupancy scenarios. We then process these states through the policy network to produce an action distribution. By associating the state-action pairs with their corresponding nodes in the building graph, we create a graph-structured feature set. For instance, the node features for a specific thermal zone would include the current local temperature, the occupancy status, and the specific control action dictated by the reinforcement learning policy, such as the intended airflow rate. This transforms the abstract policy into a concrete, topologically grounded dataset that reflects how the agent intends to operate the building under a spectrum of conditions.

3.3 Graph Neural Networks for Policy Evaluation

With the building represented as a graph and the policy's behavioral footprint encoded as node features, we employ advanced graph neural network architectures to predict the policy's overall energy optimization performance. The core mechanism of a graph neural network involves message passing, wherein each node updates its internal representation by aggregating feature information from its connected neighbors. In our framework, this message-passing process conceptually mirrors the thermodynamic flow within the building. As the graph neural network processes the policy features, it learns to identify instances where the policy successfully coordinates actions across adjacent zones to minimize energy waste, as well as instances where conflicting actions might lead to inefficiencies. The final representations from all nodes are pooled into a single graph-level embedding, which is then passed through fully connected layers to output a continuous prediction of the expected energy savings and thermal comfort violations. This forward pass takes only fractions of a second, offering a massive computational advantage over iterative thermodynamic simulation.

4. Methodology for Predicting Energy Optimization

4.1 Data Collection and Graph Construction

The foundation of the predictive methodology relies on the rigorous construction of the building graph and the compilation of a comprehensive dataset of reinforcement learning policies. The physical properties of the target building, including the geometry of the thermal zones, the thermal resistance of the construction materials, and the topology of the variable air volume network, are parsed from architectural building information models. These properties are programmatically converted into the adjacency matrices that define our structural and mechanical graphs [10]. To generate the training data for the predictive model, we must first train a large ensemble of reinforcement learning agents. We deploy multiple algorithms, including deterministic and stochastic policy gradient methods, initializing them with varying hyperparameters, reward function weights, and neural network architectures. As these agents progress through their training regimens in a high-fidelity building simulator, we periodically save their policy checkpoints. Each checkpoint represents a unique control strategy with a distinct level of competence, ranging from highly erratic, energy-wasting behaviors early in training to sophisticated, highly optimized strategies in the later stages.

4.2 Extracting Reinforcement Learning Policy Features

Once the repository of policy checkpoints is established, the next phase involves extracting the localized features that will populate the nodes of our graph [11]. We establish a standardized evaluation dataset consisting of numerous static building states that cover the complete operational envelope of the building. This includes extreme summer cooling scenarios, mild transition seasons, and deep winter heating demands, coupled with varying internal heat gains from occupants and equipment. Each policy checkpoint is evaluated against this standardized dataset. For every state, the policy dictates a set of control actions. The average magnitude, variance, and entropy of these actions for each specific thermal zone are calculated and assigned as features to the corresponding node in the graph. Additionally, we extract the gradient of the policy network with respect to the state inputs, which provides insight into how sensitive the agent's decisions are to localized temperature fluctuations. This comprehensive feature extraction results in a rich, multi-dimensional dataset where each policy is represented not merely as a monolithic neural network, but as a spatially distributed behavioral profile mapped directly onto the physical topology of the building.

4.3 Predictive Modeling using Graph Analysis

The predictive model is constructed using a spatial graph convolutional network architecture designed specifically to process the multi-relational nature of the building graph. The network utilizes parallel convolutional layers to handle the different edge types independently. One layer processes information flow along the spatial adjacency matrix, capturing conductive and convective heat transfer interactions, while another layer operates on the mechanical adjacency matrix, capturing the operational dependencies of the air distribution system. The outputs of these parallel layers are concatenated and passed through subsequent non-linear activation functions and pooling layers. The training objective of this graph network is to minimize the mean squared error between its predicted energy efficiency score and the actual ground-truth energy efficiency score, which is obtained by running the corresponding policy through a comprehensive, year-long simulation in the building simulator. Through backpropagation, the graph neural network learns to correlate specific spatial patterns of control actions with overall building energy performance. For example, it learns that policies exhibiting high variance in action outputs across highly coupled spatial nodes generally result in poor energy efficiency due to simultaneous heating and cooling conflicts.

5. Experimental Setup and Simulation Environment

5.1 Building Characteristics and Heating Configuration

To validate the proposed graph-based prediction framework, a rigorous experimental environment was established using an industry-standard building energy simulation platform. The chosen testbed is a highly detailed, calibrated model of a medium-sized commercial office building. This structure encompasses a diverse array of functional spaces, including open-plan work areas, enclosed private offices, central conference rooms, and peripheral corridors, resulting in a complex multi-zone thermodynamic environment [12]. The building features a centralized variable air volume air handling unit equipped with heating and cooling coils, dynamically supplying conditioned air to terminal boxes in each distinct thermal zone. The graph representation of this building comprises numerous nodes corresponding to these zones and the main mechanical equipment. The edge weights were determined based on the physical surface areas shared between adjacent zones and the designed maximum volumetric flow rates of the duct branches. The simulation utilizes historical meteorological data to provide realistic external temperature, solar irradiance, and humidity conditions over an entire meteorological year.

Table 1: Building Simulation Parameters and Thermal Characteristics

Parameter	Value Range	Distribution Strategy
Total Thermal Zones	35 to 40 individual zones	Spatial mapping across multiple floors
Exterior Wall Insulation	R-15 to R-20 continuous	Applied to perimeter structural envelope
Internal Heat Gain	10 to 15 Watts per square meter	Scheduled based on occupancy profiles
Maximum Supply Airflow	0.5 to 2.0 cubic meters per second	Modulated via variable frequency drives

5.2 Training Regimen for Baseline Reinforcement Learning Policies

The generation of the diverse policy dataset required a massive, parallelized training effort [13]. Over one hundred independent deep reinforcement learning agents were initialized to learn the building control task. The agents were tasked with modulating the supply air temperature setpoint of the central air handling unit and the individual damper positions of the variable air volume terminal units. The reward function was designed as a weighted sum of two conflicting objectives: the minimization of total electrical and thermal energy consumption, and the minimization of temperature deviations from predefined occupant comfort bounds. By systematically varying the penalty weight assigned to comfort violations across different training runs, we induced the agents to learn a wide spectrum of behavioral strategies, ranging from highly aggressive energy-saving policies that frequently violate comfort constraints, to highly conservative policies that maintain strict temperature control at the expense of elevated energy use. Policy checkpoints were saved at regular intervals throughout the millions of simulated timesteps, ultimately yielding a dataset of several thousand distinct policies, each with an empirically derived ground-truth performance metric obtained via exhaustive simulation.

6. Results and Comparative Analysis

6.1 Predictive Accuracy of Graph-Based Models

The primary evaluation metric for the proposed methodology is the accuracy with which the graph neural network can predict the actual simulated energy performance of a given policy. To establish a rigorous benchmark, the performance of the proposed multi-relational graph convolutional network was compared against several baseline predictive models. These baselines included a standard multi-layer perceptron that processes the flattened policy feature vectors without any topological awareness, and a basic convolutional neural network that treats the building zones as an artificial grid. The evaluation was conducted on a held-out test set of policy checkpoints that the predictive models had not encountered during their training phase. The results indicate a substantial and statistically significant advantage for the graph-based methodology [14]. By intrinsically incorporating the thermodynamic and mechanical connectivity of the building, the graph network was able to accurately forecast energy savings with remarkably low error margins. The non-graph baseline models frequently struggled to accurately predict the performance of policies that exhibited complex, coordinated multi-zone actions, underscoring the critical necessity of spatial context in evaluating holistic building operations.

Table 2: Predictive Accuracy Metrics of Graph-Based Policy Evaluation Models

Evaluation Model	Mean Absolute Error	Root Mean Square Error
Standard Multi-Layer Perceptron	14.5 Percent	18.2 Percent
Grid-Based Convolutional Network	11.2 Percent	14.7 Percent
Multi-Relational Graph Network	4.8 Percent	6.1 Percent

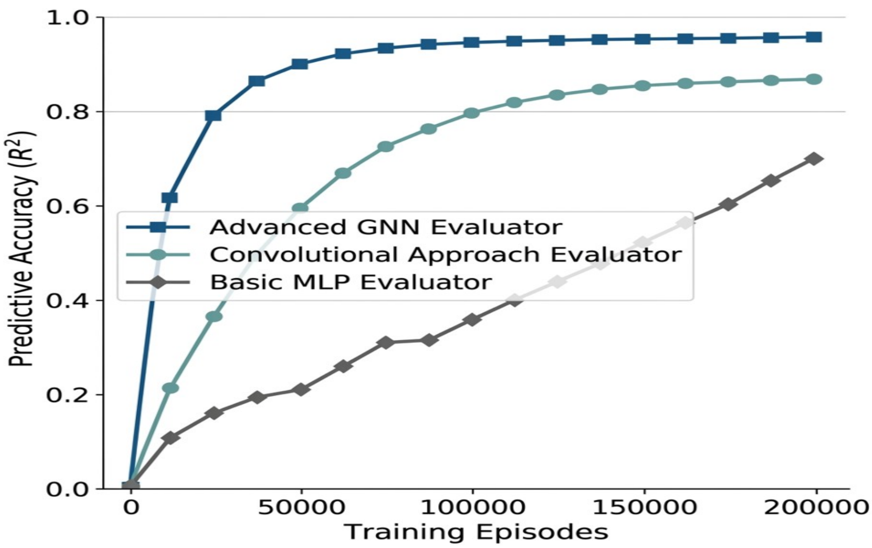

Figure 2: Predictive Accuracy Trajectories of Graph

6.2 Impact of Node Centrality on Energy Efficiency

Further analysis of the trained graph neural network revealed profound insights into the underlying mechanisms of effective building control [15]. By examining the attention weights and feature gradients within the graph network, we could identify which specific nodes and edges were most influential in determining the overall energy efficiency score of a policy. The analysis demonstrated that nodes with high eigenvector centrality, typically corresponding to large, centrally located thermal zones with multiple shared partitions, exerted a disproportionate impact on the building's global energy profile. Reinforcement learning policies that achieved the highest energy savings were found to exhibit highly coordinated, synchronized control actions specifically tailored to these high-centrality nodes. Conversely, policies that performed poorly often exhibited erratic or independent control actions in these central hubs, leading to cascading thermal imbalances that forced peripheral zones to consume excessive energy to compensate. This finding empirically validates the hypothesis that the topological location of a space within a building fundamentally alters its optimal operational strategy, a nuance that is effectively captured and evaluated by our graph-based predictive framework.

7. Discussion and Operational Implications

7.1 Interpretability of Reinforcement Learning Policies via Graph Metrics

One of the most persistent barriers to the commercial adoption of reinforcement learning in critical infrastructure is the lack of interpretability. Facility managers and control engineers require transparent assurances that an autonomous agent will behave rationally and safely under all foreseeable conditions [16]. The proposed graph analysis framework provides a novel mechanism for demystifying deep reinforcement learning policies. By translating abstract neural network weights into spatially contextualized behavioral features, engineers can visually map and inspect the intended control strategies before deployment. If the graph analysis predicts poor energy performance or highlights severe action conflicts between strongly coupled adjacent zones, the engineering team can safely discard the policy or adjust the reinforcement learning reward structure without risking occupant discomfort or equipment damage in the real world. Furthermore, this predictive capability drastically accelerates the reinforcement learning development cycle. Instead of relying exclusively on time-consuming thermodynamic simulations to evaluate every iteration of an agent, developers can utilize the rapidly executing graph model to screen thousands of candidate policies, reserving rigorous simulation only for the most promising candidates.

7.2 Limitations of the Current Approach

While the results demonstrate a significant advancement in policy evaluation, several limitations must be acknowledged. The accuracy of the graph-based predictive model is inherently constrained by the fidelity of the graph representation itself. Accurately quantifying the edge weights for complex convective airflow pathways, such as multi-story atriums or open stairwells, remains a challenging task that often requires approximation. Additionally, the current methodology relies on static structural graphs. In reality, the thermal coupling between zones can be highly dynamic, influenced by occupant behavior such as the opening and closing of internal doors or exterior windows. The predictive model in its current iteration does not fully account for these dynamic topological shifts during its rapid evaluation pass. Furthermore, while the model accurately predicts long-term energy averages, it may obscure short-term, transient violations of thermal comfort bounds that occur during extreme weather anomalies not adequately represented in the standardized feature extraction dataset.

8. Conclusion

8.1 Summary of Findings

The integration of advanced machine learning into the operation of the built environment holds immense potential for reducing global energy consumption and mitigating the impacts of climate change. This research has successfully demonstrated that the complex, often opaque policies generated by deep reinforcement learning algorithms can be effectively evaluated and predicted using advanced graph analysis techniques. By formalizing the structural and mechanical systems of a smart building as a multi-relational graph, we transformed abstract control policies into topologically grounded data representations. The application of graph neural networks to these representations yielded highly accurate predictions of building energy performance, vastly outperforming traditional evaluation metrics that lack spatial awareness. The findings confirm that the key to optimizing multi-zone environments lies in understanding and exploiting the thermodynamic dependencies that exist between interconnected spaces.

8.2 Future Research Directions

The methodology established in this paper lays the groundwork for numerous avenues of future investigation. Subsequent research should focus on incorporating dynamic graph structures that can adapt to real-time changes in building topology, such as adaptive facades or variable internal partitioning [17]. Additionally, exploring the integration of temporal graph networks could enable the prediction of a policy's performance over varying time horizons, capturing the delayed thermal inertia of heavy construction materials. Finally, translating this predictive framework into an active mechanism for policy shaping presents a compelling opportunity. Rather than simply evaluating fully trained policies post-hoc, the graph-based performance gradients could theoretically be fed back into the reinforcement learning training loop, directly guiding the agent toward spatially coordinated, highly efficient control strategies from the very onset of the learning process. Such advancements will be instrumental in realizing the vision of fully autonomous, inherently efficient smart buildings.

References

1. Namoune, A.; Chaker, A.; Saouane, I. A Dual-Approach Machine Learning and Deep Learning Framework for Enhanced Fault Detection in Photovoltaic Systems: Incorporating SDM Parameter Analysis and Thermal Imaging. J. Renew. Sustain. Energy 2025, 17, 033502.

2. Agarwal, A.; de Jesus Colwell, F.; Bello Rodriguez, J.; Sommer, S.; Barman, M.; Correa Galvis, V.A.; Hill, T.R.; Boonham, N.; Prashar, A. Synergistic 3D, multispectral, and thermal image analysis via supervised machine learning for improved detection of root rot symptoms in hydroponically grown flat-leaf parsley. Smart Agric. Technol. 2025, 12, 101364.

3. Kaewdornhan, N.; Srithapon, C.; Liemthong, R.; Chatthaworn, R. Real-Time Multi-Home Energy Management with EV Charging Scheduling Using Multi-Agent Deep Reinforcement Learning Optimization. Energies 2023, 16, 2357.

4. Rahman, M.K.; Shamsuddoha, M. AI contribution for developing green visibility and integration toward sustainability performance in supply chain. J. Enterp. Inf. Manag. 2025, 1–39.

5. Putri, R.I.; Rifa'i, M.; Riskitasari, S. Telemonitoring for Photovoltaic Systems Using Internet of Things. In Proceedings of the 2023 International Conference on Advanced Mechatronics, Intelligent Manufacture and Industrial Automation (ICAMIMIA), Lombok, Indonesia, 14–15 November 2023; pp. 1–5.

6. Amini Toosi, H.; Del Pero, C.; Leonforte, F.; Lavagna, M.; Aste, N. Machine Learning for Performance Prediction in Smart Buildings: Photovoltaic Self-Consumption and Life Cycle Cost Optimization. Appl. Energy 2023, 334, 120648.

7. Qin, Y.; Zhao, Q.; Zhou, D.; Shi, Y.; Shou, H.; Li, M.; Zhang, W.; Jiang, C. Application of flash GC e-nose and FT-NIR combined with deep learning algorithm in preventing age fraud and quality evaluation of pericarpium citri reticulatae. Food Chem. X 2024, 21, 101220.

8. Lee, N.-Y.; Na, I.-S.; Lee, K.-W.; Lee, D.-H.; Kim, J.-W.; Kook, M.-C.; Hong, S.-J.; Son, J.-Y.; Lee, A.Y.; Om, A.-S.; et al. Detection of physical hazards from fruit processed products using hyperspectral imaging and prediction based on PLS-DA and logistic regression machine learning models. Appl. Food Res. 2024, 4, 100506.

9. Xie, Y.; Chang, Y.; Li, M.; Qin, A.K.; Zhang, X. AutoSGRL: Automated framework construction for self-supervised graph representation learning. Neural Netw. 2026, 194, 108119.

10. Sutton, R.S.; Barto, A.G. Reinforcement Learning: An Introduction, 2nd ed.; MIT Press: Cambridge, MA, USA, 2018.

11. Sun, H.; Hua, Z.; Yin, C.; Li, F.; Shi, Y. Geographical traceability of soybean: An electronic nose coupled with an effective deep learning method. Food Chem. 2024, 440, 138207.

12. Sarkar, S. AI-Enabled Digital Twin Framework for Predictive Maintenance and Energy Optimization in Industrial Systems. ASRC Procedia Glob. Perspect. Sci. Scholarsh. 2025, 1, 1359–1389.

13. Shi, T.; Xu, C.; Dong, W.; Zhou, H.; Bokhari, A.; Klemeš, J.J.; Han, N. Research on Energy Management of Hydrogen Electric Coupling System Based on Deep Reinforcement Learning. Energy 2023, 282, 128174.

14. Jiang, X.; Liu, Q.; Yan, L.; Cao, X.; Chen, Y.; Wei, Y.; Wang, F.; Xing, H. Hyperspectral imaging combined with spectral-imagery feature fusion convolutional neural network to discriminate different geographical origins of wolfberries. J. Food Compos. Anal. 2024, 132, 106259.

15. Yagli, G.M.; Yang, D.; Srinivasan, D. Ensemble Solar Forecasting and Post-Processing Using Dropout Neural Network and Information from Neighboring Satellite Pixels. Renew. Sustain. Energy Rev. 2022, 155, 111909.

16. Li, Y.; Zhou, W.; Wang, Y.; Miao, S.; Yao, W.; Gao, W. Interpretable Deep Learning Framework for Hourly Solar Radiation Forecasting Based on Decomposing Multi-Scale Variations. Appl. Energy 2025, 377, 124409.

17. Chakraborty, D.; Mondal, J.; Barua, H.B.; Bhattacharjee, A. Computational Solar Energy—Ensemble Learning Methods for Prediction of Solar Power Generation Based on Meteorological Parameters in Eastern India. Renew. Energy Focus 2023, 44, 277–294.