PAMUKKALE ÜNİVERSİTESİ
Mühendislik Bilimleri Dergisi

PAMUKKALE UNIVERSITY
Journal of Engineering Sciences

Cilt | Vol: XX • Sayı | No: X • ss | pp: XXXX-XXXX • Yıl | Year: 20XX

Techno-economic validation of AI-based energy optimization for smart campuses: a digital twin simulation approach

Abdalhadi Manassra [1], Gürkan Işık [2]

ARTICLE INFO

Dates:
Received: 05.01.2026
Revision: 17.02.2026
Accepted: 09.03.2026

Doi:
10.65206/pajes.1856546

Corresponding author:
Abdalhadi Manassra
(abdomanasra12@gmail.com)

Author addresses:
[1] Department of Industrial Engineering, Bursa Technical University, Türkiye, (abdomanasra12@gmail.com)

[2] Department of Industrial Engineering, Bursa Technical University, Türkiye, (gurkan.isik@btu.edu.tr)

ABSTRACT

Context—The rising trajectory of energy consumption in mid-to-large scale educational facilities presents a significant challenge for modern sustainability goals. University campuses, characterized by complex infrastructure and fluctuating occupancy patterns, often suffer from inefficiencies inherent in traditional, rule-based Building Management Systems (BMS). These static systems frequently fail to adapt to dynamic operational conditions, leading to excessive energy waste and compromised occupant comfort. Consequently, there is a critical need to transition from reactive management to intelligent, predictive systems capable of handling the stochastic nature of campus energy demands. Addressing this gap requires robust validation of adaptive technologies that can harmonize energy efficiency with operational cost-effectiveness.
Objective—In the framework of the problems associated with static energy management, the main purpose of this study is to evaluate the technical performance and economic feasibility of a proposed hybrid Artificial Intelligence-based energy optimization model designed specifically for smart campuses. This research aims to bridge the gap between theoretical AI models and practical application by investigating whether an autonomous system can significantly reduce energy consumption without degrading user comfort, while simultaneously proving its financial viability as a long-term investment for institutional stakeholders.
Method—To achieve rigorous validation, a high-fidelity digital twin of a mid-sized university campus was developed. The study employs a Counterfactual Simulation assumption, running parallel execution threads to benchmark the AI model against a Rule-Based Control (RBC) baseline under identical TMY meteorological conditions. This virtual environment models a 100,000-square-meter indoor area comprising four distinct building typologies: administrative offices, laboratories, classroom buildings, and cafeterias. The study simulates real-time operations under distinct summer and winter meteorological scenarios to test system resilience. The proposed control architecture integrates three specific neural network models: Deep Q-Networks (DQN) are utilized to optimize heating, ventilation, and air conditioning (HVAC) systems; Convolutional Neural Networks (CNN) are deployed for precise, motion-based lighting control, and Multi-Layer Perceptrons (MLP) are applied to manage plug load anomalies.

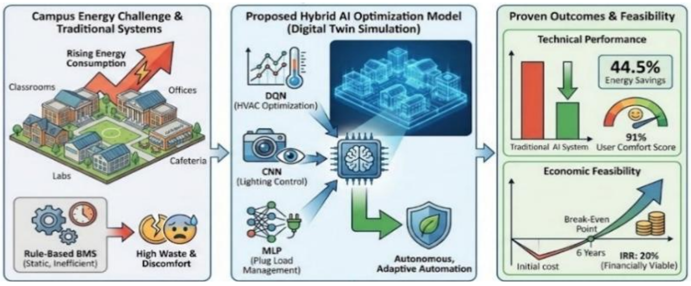

Results—The most significant findings obtained from the simulation demonstrate that the proposed hybrid AI system outperforms traditional rule-based baselines by a substantial margin. The system achieves an average energy saving of 44.5% across the simulated campus, while successfully maintaining a user comfort score of 91% in compliance with ASHRAE 55, indicating that efficiency did not come at the cost of habitability. Furthermore, the financial analysis, recalculated using these simulation-verified savings, reveals compelling economic indicators. Despite the high initial capital expenditures required for implementation, the system achieves a break-even point within just 6 years and stabilizes at an Internal Rate of Return (IRR) of 20%.
Conclusion—These findings confirm that autonomous automation is not only technically robust but also presents a financially viable solution for sustainable campus infrastructure. The study implies that hybrid AI models offer a scalable pathway for educational institutions to reduce their carbon footprint and operational costs. Future studies may focus on integrating renewable energy sources into this control architecture to further enhance grid independence and sustainability.

Key Words—Smart campus, Hybrid artificial intelligence, Digital twin, Economic feasibility, Energy optimization.

I. INTRODUCTION

The escalating global demand for energy, coupled with increasing concerns over sustainability and operational costs, places significant pressure on large-scale infrastructure, particularly university campuses. These institutional settings are characterized by high energy intensity due to the diverse and complex operational profiles of administrative offices, laboratories, classrooms, and communal areas, primarily driven by Heating, Ventilation, and Air Conditioning (HVAC) systems and lighting [1]. Effective energy management is thus crucial, shifting the focus from simple energy conservation to autonomous, data-driven optimization.

The adoption of smart campus initiatives, supported by Internet of Things (IoT) infrastructure, provides the base data necessary for deploying sophisticated energy management strategies. In this context, Artificial Intelligence (AI) and Machine Learning (ML) models such as Deep Q-Networks (DQN) for dynamic control and Convolutional Neural Networks (CNN) for sensory data processing—have emerged as superior alternatives to traditional, rule-based systems [2]. While academic literature presents numerous theoretical AI-based models for energy optimization, a critical challenge remains: the validation of these models under real-world, complex operational conditions.

A. The critical role of simulation and the research gap

Direct implementation of advanced AI algorithms on live campus infrastructure is often impractical due to high costs, operational disruptions, and the risk of unexpected system failures that could affect thousands of users. Consequently, simulation and digital twin environments become essential tools. A digital twin, which is a high-fidelity virtual replica of the physical campus, allows for the rigorous testing of AI control strategies across diverse environmental and occupancy scenarios (e.g., peak summer cooling vs. deep winter heating) without risk [3],[4].

Despite the importance of simulation-based validation, current research suffers from two major limitations:

- Limited Economic Validation: While many studies report promising energy saving percentages, the critical analysis of the economic feasibility is often neglected or based on simplistic theoretical cost estimates [5]. For a large-scale project to move from research to implementation, decision-makers require validated metrics such as the Net Present Value (NPV), Internal Rate of Return (IRR), and a reliable Payback Period calculated directly from simulation-verified performance data. The lack of this financial rigor hinders the adoption of AI systems by stakeholders.
- Focus on Single Loads: Most successful AI applications focus narrowly on optimizing a single high-consumption load, such as HVAC or lighting, failing to address the complex, interdependent energy dynamics of the entire building ecosystem [6]. A truly effective solution requires a hybrid, integrated framework to manage multiple loads simultaneously and holistically.

B. Contribution of this study

This paper addresses the identified limitations by presenting an integrated approach that combines advanced simulation with comprehensive economic analysis. The main contributions are:

- *Development of a Hybrid AI Architecture*: We propose and simulate a hybrid AI system integrating DQN, CNN, and Multi-Layer Perceptron (MLP) models to simultaneously optimize three main energy loads (HVAC, lighting, and plug loads) across different building typologies within a smart campus environment.
- *Rigorous Simulation and Validation*: A high-fidelity Digital Twin of a real-world campus (100,000 m^2) is utilized to rigorously test the hybrid AI model against conventional baselines under various seasonal conditions, providing highly reliable performance data.
- *Comprehensive Economic Feasibility Study*: The simulation-verified energy savings are used to conduct detailed financial analysis, specifically calculating the Net Present Value (NPV), Internal Rate of Return (IRR), and Payback Period over a 20-year horizon, thereby bridging the gap between technical performance and financial viability.

Unlike prevalent studies that strictly isolate algorithmic performance (kWh savings), this research integrates a comprehensive financial evaluation framework directly into the validation loop. By correlating stochastic AI control policies with long-term financial viability metrics (ROI, NPV, Payback Period), this study provides a validated roadmap for facility managers by demonstrating that the proposed system achieves significant energy savings while yielding positive financial indicators (positive ROI and acceptable payback periods), confirming the economic feasibility of AI-based retrofitting, demonstrating that AI retrofitting is not merely a theoretical exercise but a financially sound investment for legacy campus infrastructures.

To address the identified gaps in existing literature, this study pursues the following four primary objectives:

- *Development of a High-Fidelity Digital Twin*: To construct a validated virtual replica of a university campus that accurately simulates thermal dynamics and stochastic occupancy patterns for risk-free algorithm training.
- *Hybrid AI Framework Implementation*: To design and deploy a multi-agent architecture combining Deep Q-Networks (DQN) for HVAC control, Convolutional Neural Networks (CNN) for spatial occupancy detection, and Multi-Layer Perceptrons (MLP) for plug-load anomaly management.
- *Techno-Economic Validation*: To bridge the gap between theoretical engineering efficiency and financial feasibility by integrating Return on Investment (ROI) and Net Present Value (NPV) metrics directly into the performance evaluation.
- *Benchmarking against Legacy Systems*: To quantify the superiority of the proposed AI model through a counterfactual comparison with standard Rule-Based Control (RBC) systems under identical meteorological conditions.

The remainder of the paper is organized as follows: Section II literature review, Section III outlines the Hybrid AI Methodology and Digital Twin setup. Section IV details the simulation scenarios and presents the energy performance results. Section V provides detailed economic feasibility analysis. Finally, Section VI offers the conclusions and outlines future research directions.

II. LITERATURE REVIEW

Recent reviews of smart building technologies indicate a predominant reliance on machine learning (ML) and deep learning to enhance load prediction, HVAC and lighting control, and demand response participation, resulting in documented gains in energy efficiency and reduced operating costs [7],[8]. Systematic analyses emphasize that AI has become a core component of Building and Campus Energy Management Systems (BEMS/CEMS), supporting essential tasks such as demand forecasting, occupancy and fault detection, and multi-objective optimization that balances cost, comfort, and emissions [9],[10]. These capabilities are increasingly amplified through integration with IoT sensing, big data, and digital twins, which enable real-time monitoring and advanced autonomous control at the campus scale [11],[12]. However, while scoping reviews confirm that these data-driven technologies underpin business models focused on efficiency and grid services, they also highlight persistent barriers to adoption, including high initial costs, interoperability challenges, cybersecurity risks, and issues regarding user acceptance [13].

Recent literature highlights the growing implementation of AI-driven energy management systems in university environments, aiming to create "smart, green, and digital" campuses through the integration of IoT and advanced algorithms [14]. Technical methodologies typically combine supervised learning (e.g., XGBoost, ANN) for load and solar forecasting with deep reinforcement learning and digital twins for real-time, multi-objective control across building clusters. Specific case studies demonstrate the efficacy of these approaches: a university system utilizing hybrid forecasting and reinforcement learning achieved high accuracy (MAPE 2.18%) to reduce grid dependence [15], while AI controls in laboratory buildings successfully lowered greenhouse gas emissions without compromising safety. Similarly, applications in boarding schools and faculty buildings have utilized AI-based scheduling for storage and HVAC systems to maximize resource efficiency [8].

Economically, the deployment of these AI-enabled systems consistently delivers substantial returns, though challenges remain. Evidence indicates that when fully exploited, AI and IoT frameworks can reduce electricity consumption by 15–40% and cut annual electricity costs by 20–55%, particularly when optimizing renewable assets like photovoltaics and storage [16]. One study notably reported a 55% reduction in electricity bills by shifting grid purchases to low-price periods, sustaining a 13-year payback model. However, despite these documented gains, widespread adoption is hindered by barriers including high upfront investments, interoperability and cybersecurity risks, and a shortage of necessary technical skills [17].

While recent studies [18],[19] have successfully demonstrated the potential of Deep Reinforcement Learning for energy optimization, they predominantly focus on thermodynamic metrics (kWh savings) and often overlook the financial feasibility required for real-world adoption. This study differentiates itself by integrating a Techno-Economic Analysis framework directly into the Digital Twin loop. Unlike standard algorithmic studies, we provide a holistic evaluation that correlates AI-driven energy savings with Return on Investment (ROI) and Net Present Value (NPV) metrics, thereby bridging the gap between theoretical AI efficiency and practical facility management decisions. Therefore, the novelty of this work lies at the system and decision-support level rather than at the algorithmic level.

III. METHOD

To rigorously validate the proposed energy management methodology and assess its economic feasibility, a comprehensive simulation approach was employed. This section details the configuration of the digital twin environment, the architecture of the hybrid AI optimization model, and the established simulation scenarios.

A. Digital twin simulation environment

The study utilized a high-fidelity Digital Twin framework, modeled after a university campus with a total area of approximately 100,000 square meters. The virtual environment was constructed to accurately replicate the diverse functional zones and energy consumption profiles of a typical educational facility. The campus model was divided into four main building typologies, each with unique occupancy schedules, thermal characteristics, and load requirements:

- *Administrative Building*: Characterized by fixed working hours, high plug load density (computers, servers), and moderate HVAC demand.
- *Laboratory Building*: Defined by high internal heat gains from specialized equipment, variable occupancy, and stringent temperature/humidity control requirements.
- *Classroom Building*: Characterized by transient and high-density occupancy, and cyclic usage (high demand during teaching hours, low during breaks).
- Cafeteria/Common Area: Marked by high and volatile thermal loads from kitchen equipment, high lighting usage, and sporadic peak occupancy.

B. Hybrid AI optimization methodology

The proposed system adopts hybrid AI architecture to manage the three most significant energy loads simultaneously and interdependently. This distributed control system leverages the strengths of specialized machine learning algorithms for specific tasks, ensuring both high efficiency and user comfort [20]. The proposed architecture is organized into three distinct operational layers: Sensing, Processing, and Action. Figure 1 illustrates this data flow, detailing how specific environmental and visual inputs Layer 1 are routed to their respective AI agents CNN, DQN, and MLP within the processing core Layer 2 to execute precise control actions on the physical building systems Layer 3.

Inter-Agent Interaction via Environmental State: It is important to note that the proposed Hybrid architecture utilizes Environmental Coupling rather than direct agent-to-agent communication. The interaction occurs through the thermodynamic state of the Digital Twin. For instance, heat generated by lighting (controlled by CNN) and plug loads (managed by MLP) are integrated into the zone's thermal equation as Internal Heat Gain. The DQN agent, which continuously monitors the zone temperature state, implicitly detects these load changes through their thermal impact and adjusts the HVAC setpoints accordingly. This decoupled design ensures high system resilience; the failure of a peripheral agent (e.g., lighting) does not disrupt the critical HVAC control loop, as the DQN continues to optimize based on real-time environmental feedback. Thus, the proposed hybrid architecture is hybrid at the system level (multiple AI models acting on different subsystems) rather than at the algorithmic or parameter-sharing level.

1. DQN for HVAC control

HVAC systems account for the largest proportion of building energy consumption. A DQN-based agent was deployed to manage the central HVAC system.

- *Purpose*: To learn the optimal setpoint adjustments (actions) that minimize energy consumption (reward) while satisfying thermal comfort constraints (state variables like indoor temperature, humidity, and forecast).
- *State Space*: Includes current zone temperature, predicted outdoor temperature, time of day, and occupancy level.
- *Action Space*: Discrete temperature adjustments (e.g., ±0.5°C, ±1.0°C) and system status (on/off). The Reinforcement Learning approach ensures the system can dynamically adapt to unforeseen disturbances and long-term changes in building characteristics.

Reward Function Formulation: To ensure the agent learns a balanced control policy, the reward mechanism is designed as a multi-objective trade-off between energy efficiency and occupant comfort. Instead of a single static goal, the agent receives a dynamic reward based on the incremental energy savings achieved compared to the Rule-Based Control (RBC) baseline.

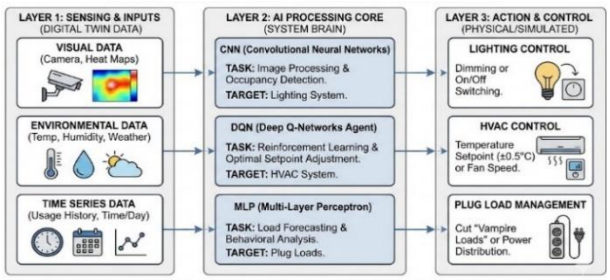

Figure 1. Hierarchical three-layer control architecture: Data flow from multi-modal sensing to hybrid AI processing core.

However, to prevent the system from simply shutting down equipment to save energy, a strictly prohibitive penalty is applied whenever the indoor zone temperature deviates from the comfortable range (defined as 20°C to 24°C). This binary penalty structure forces the agent to prioritize thermal compliance as a hard constraint while maximizing energy efficiency as a secondary objective.

2. CNN for lighting management

Lighting systems require precise control based on occupancy and available natural daylight.

- *Purpose*: To accurately detect and predict localized occupancy and light levels within large open spaces. CNN processes visual data (simulated camera feeds or heat maps) and light sensor data.
- *Function*: The CNN classifies occupancy and ambient light intensity, providing a real-time input signal to the lighting controller. This allows for granular control (dimming or switching off unused zones), maximizing energy savings without compromising visibility or comfort.

3. MLP for plug load prediction

Plug loads (unregulated devices like computers and office equipment) are highly dependent on human behavior.

- *Purpose*: To predict future plug load demand based on historical usage patterns and time-series inputs (day type, hour).
- *Function*: The MLP serves as a forecasting tool, allowing the central energy manager to optimize power distribution and prioritize essential systems, thereby reducing peak demand and minimizing wasted energy from devices left unnecessarily active.

C. Simulation scenarios and baseline

To ensure the validity and generalizability of the results, the system was tested across two distinct operational periods that is shown in Table1.

Baseline Model, the performance of the hybrid AI system was benchmarked against a conventional Rule-Based Control (RBC) system. The RBC system utilized fixed temperature setpoints (e.g., 22°C for heating, 24°C for cooling) and time-scheduled lighting/ventilation protocols, representing the standard, non-AI management approach currently prevalent in similar campus facilities. Energy savings and performance metrics are calculated exclusively by comparing the AI model's consumption against this fixed RBC baseline.

While Model Predictive Control (MPC) represents a robust alternative for energy optimization [21], it typically necessitates precise physics-based modeling of the thermal environment, which limits scalability in complex, heterogeneous campuses. In contrast, the proposed Deep Reinforcement Learning (DQN) framework operates as a model-free learner, allowing it to adapt to unique zonal dynamics without predefined mathematical models. Consequently, this study benchmarks performance against Rule-Based Control (RBC), as RBC remains the predominant operational standard in existing educational facilities, thereby providing the most relevant baseline for assessing the economic viability of retrofitting legacy systems.

Table 1. Operational boundary conditions and meteorological scenarios for counterfactual digital twin simulation.

Scenario	Conditions	Primary energy focus
Summer Simulation	High outdoor temperatures (20°C-50°C), peak solar radiation	Cooling (HVAC) and lighting management
Winter Simulation	Low outdoor temperatures (-5°C–19°C), low solar radiation.	Heating (HVAC) and maintaining thermal comfort

D. Ethical considerations

Given the deployment in semi-public educational spaces, the proposed system adheres to strict Privacy by Design architecture. To comply with KVKK and GDPR, the occupancy detection module does not utilize RGB cameras or visual imagery.

Instead, it processes raw data from Low-Resolution Thermal Sensor Arrays (e.g., 8×8- or 32×24-pixel grids). The CNN model is employed here not for visual image processing, but for Spatial Pattern Recognition on the raw temperature matrix. This allows the system to distinguish human occupancy patterns from other heat sources (such as radiators, projectors, or hot beverages) which simple threshold-based sensors might misinterpret. Since the input is a low-resolution numerical matrix rather than a visual video feed, facial recognition or personal identification (PII) is technically impossible. All processing occurs locally on the Edge Controller, and only the final anonymous count is transmitted via MQTT [22].

E. Experimental configuration

Hyperparameter Configuration: The Deep Reinforcement Learning agents were trained using an Adam optimizer with a learning rate set to 0.001 to ensure stable gradient descent. The discount factor was maintained at 0.95, prioritizing long-term energy efficiency over immediate short-term gains. To facilitate effective state-space exploration, an epsilon-greedy strategy was employed, starting with an exploration rate of 1.0 and decaying exponentially by a factor of 0.995 per episode until reaching a minimum floor of 0.01. The experience replay buffer was sized at 2000 transitions, with a batch size of 64 samples used for each training update to break temporal correlations in the learning data.

Dataset Generation and Validation Protocol: Given the Digital Twin nature of this study, the training dataset was synthetically generated to ensure a controlled and noise-free learning environment. The simulation utilized Typical Meteorological Year (TMY) weather files specific to the campus location to replicate realistic seasonal variations. Occupancy patterns were modeled using a stochastic Markov Chain process to introduce realistic unpredictability into the classroom usage profiles. For validation purposes, the generated dataset was partitioned using an 80/20 split protocol, where 80% of the simulated days were dedicated to agent training (exploration phase) and the remaining 20% were reserved for testing (exploitation phase) to evaluate the model's generalization capability on unseen environmental conditions.

IV. SIMULATION RESULTS

This section presents a rigorous quantitative evaluation of the proposed hybrid AI energy management system. The performance of the system was validated using a high-fidelity digital twin environment, comparing the AI-driven autonomous control against a standard Rule-Based Control (RBC) baseline across four distinct building typologies.

The Digital Twin environment was implemented as an interactive web-based application using the Streamlit open-source framework (v1.28). Streamlit was selected for its seamless integration with the Python ecosystem, allowing the backend machine learning models (DQN, CNN, and MLP) to communicate directly with the frontend visualization layer without the latency of traditional web architectures.

The interface serves as the control center for simulation, enabling real-time validation of the AI agents' performance. It features a dynamic dashboard that visualizes key operational metrics, including zone-specific temperatures, power consumption (kW), and cumulative cost savings (USD/₺) at 15-minute intervals. Furthermore, the application provides a parameter configuration sidebar, allowing for the hardware in the loop style testing of different environmental variables such as modifying weather profiles or occupancy density on the fly. This interactive capability was crucial for validating the system's robustness against the stochastic scenarios described in Table 1, verifying that the AI agents could adapt to sudden changes in the virtual environment.

Underpinning this visualization layer is a robust communication topology designed to mimic a real-world IoT ecosystem. As illustrated in Fig. 2, the system architecture operates on a hierarchical edge-to-cloud model. Physical data collection is handled by distributed Edge Controllers (simulated Raspberry Pi 4 units) which aggregate sensor readings via localized Modbus RTU/TCP connections. This telemetry data is then transmitted asynchronously to the Central AI Server through an MQTT Broker, ensuring low-latency data ingestion. The AI Processing Core (The 'System Brain') analyzes this stream, computes the optimal control actions (e.g., HVAC setpoint adjustments), and publishes commands back to the edge actuators, while simultaneously pushing state updates to the Streamlit Dashboard server via REST API for real-time monitoring.

To address the challenges of network latency and reliability in a large-scale (100,000 m^2) deployment, the system adopts a Distributed Edge Computing strategy rather than a centralized cloud-dependency model. Critical control decisions, particularly CNN-based occupancy detection and lighting actuation, are processed locally on the Edge Controllers (Jetson Nano). This ensures that actuation designed to remain below typical real-time control thresholds (<100 ms) and is independent of campus-wide network traffic [23]. Furthermore, data transmission utilizes the MQTT protocol with QoS Level 1, which buffers telemetry data locally during network interruptions, safeguarding against packet loss [24]. While the current study utilizes standard development boards for validation, the architecture is compatible with industrial-grade modules (e.g., NVIDIA Jetson Industrial) for commercial scaling [25].

A. Simulation inputs and configuration

To ensure the reliability and reproducibility of the results, the digital twin was initialized with specific parameters designed to replicate the complex, dynamic nature of a physical campus. Table 2 details the environmental, physical, and operational inputs configured within the simulation platform. These values were selected to establish a rigorous testing environment that challenges the adaptability of the AI agents. For instance, the 15-

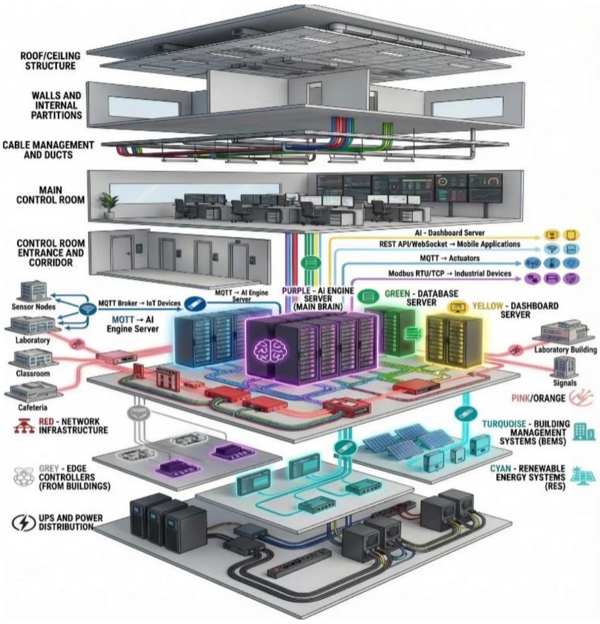

Figure 2. Campus-wide edge-to-cloud communication topology for real-time telemetry and decentralized AI actuation.

Table 2. Key input parameters and configuration settings for the digital twin simulation.

Scenario	Conditions	Primary energy focus
Environmental	Weather Profile	Seasonal (Summer/Winter)
	Time Resolution	15 minutes per step
Physical (Zonal)	Thermal Resistance k_{hvac}	Zone-specific (0.5 - 2.5)
	Lighting Load n_{led}	300 - 500 units (per building)
	Plug Load Density n_{plugs}	50 - 400 units (per building)
Operational	Occupancy Profile	Stochastic (0-24 hours)
	Comfort Band	22.0°C - 24.0°C (Occupied)
AI Hyper parameters	DQN Learning Rate	0.01
	CNN Sequence Length	21 time-steps
	MLP Hidden Layers	2 layers (64 neurons)

minute time resolution was chosen as the optimal trade-off to capture the thermal inertia of building zones k_{hvac} without imposing excessive computational loads, allowing for real-time control simulation. Similarly, stochastic occupancy profiles were implemented instead of fixed schedules to validate the system's robustness against unpredictable human behavior, ensuring the agents learn dynamic responses rather than memorizing static patterns. The AI hyper parameters, including a DQN learning rate of 0.01 and a sequence length of 21 for the CNN, were empirically optimized to ensure rapid convergence and stability, specifically tailored to fit the processing constraints of edge computing devices like Raspberry Pi and Jetson Nano utilized in the system architecture.

To ensure an objective evaluation of user satisfaction, the 'Comfort Score' is modeled as a quantitative Set-point Adherence Rate, grounded in ASHRAE 55 standards [26]. Instead of relying on subjective estimations, this metric is defined as the percentage of total occupied operational time during which the indoor operative temperature strictly remains within the assigned thermal dead-band (22°C–24°C). This approach treats thermal comfort as a deterministic binary compliance condition at every simulation time-step, ensuring that the reported figures reflect physical adherence to international thermal standards.

To ensure the robustness of the AI models against the inherent uncertainties of physical deployments (e.g., sensor inaccuracies, signal latency), the training environment was designed to move beyond idealized conditions. The simulation inputs were parametrized using real Typical Meteorological Year (TMY) data [27] and stochastic occupancy profiles based on Markov chain models [28],[29] to ensure realistic dynamics. Furthermore, a Gaussian noise injection strategy ($\mu = 0$, $\sigma = 0.05$) was applied to the state space observations, specifically temperature and occupancy sensor readings—during the training phase [30]. This approach prevents the DQN and MLP agents from overfitting perfect simulation data, forcing them to learn generalized control policies that remain effective even in the presence of noisy or partial data typically encountered in real-world IoT infrastructures.

B. Case study 1: administrative building

The administrative building simulation tested the integrated performance of all three AI models (DQN, CNN, MLP) in a standard office environment with consistent operating hours but variable internal loads.

To accurately simulate these internal loads, the building's digital twin was segmented into distinct thermal and operational zones as depicted in Fig. 3. The visualization highlights the strategic placement of multi-sensor nodes (Temperature/Humidity/CO_2) at the standard occupant breathing zone height (1.5m) to ensure comfort compliance. Furthermore, the diagram details the network of Smart Plugs (orange icons) distributed across the open-plan office areas, which serve as the primary actuation points for the MLP-based plug load management system described in the methodology.

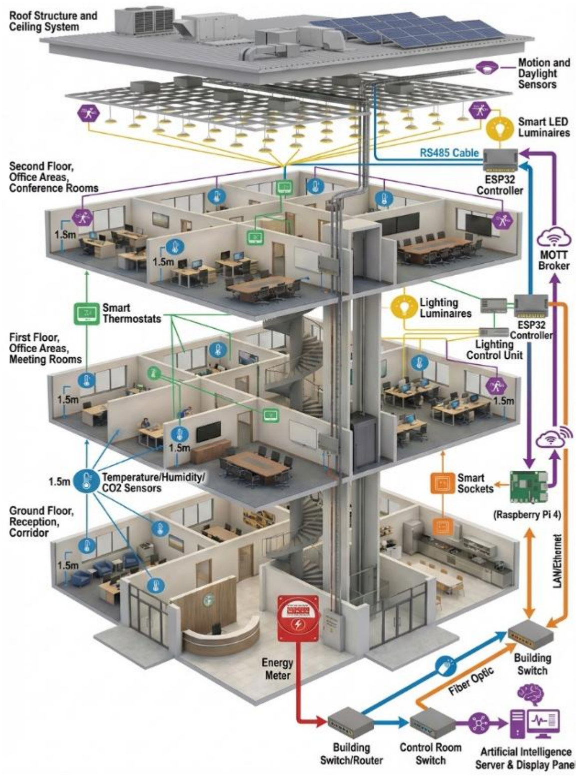

Figure 3. System architecture of administrative building.

1. HVAC performance (DQN)

The comparison of HVAC power consumption between the baseline and AI models reveals distinct operational strategies.

Summer scenario: The baseline system (blue dashed line) exhibits continuous high-power cycling, particularly wasting energy after working hours (steps 28-40) due to static scheduling. In contrast, the DQN agent (orange solid line) successfully identifies the end-of-occupancy period and cuts power to zero immediately, eliminating (vampire) energy waste (Fig. 4a).

Winter Scenario: During the initial warm-up phase (steps 0-10), the DQN agent modulates power input more efficiently than the baseline's aggressive (full-blast) approach, maintaining the comfort setpoint with lower average power consumption, ~68 kW vs. ~78 kW, (Fig. 4b).

2. Lighting performance (CNN)

Summer, the CNN model leverages high solar availability, aggressively dimming artificial lights during midday (steps 10-25) to near-zero consumption, achieving a 74.6% reduction compared to the static baseline (Fig. 4c).

3. Plug load management (MLP)

Hybrid Awareness: The MLP model demonstrated cross-system awareness. In the Winter scenario (Fig. 4e), detecting a high heating load from the HVAC system, the MLP agent aggressively reduced non-critical plug loads (orange line) to balance the total building demand, a capability completely absents in the unmanaged baseline (Figs. 4d and 4e).

Figure 4f illustrates the comparative total power consumption profiles for the Administrative Building, contrasting the static Rule-Based Control (RBC) baseline with the dynamic Hybrid AI model. The graphical data reveals that the AI controller (represented by the orange line) consistently suppresses energy demand throughout the operational day compared to the baseline (blue line). A critical performance differentiator is observed during the post-occupancy period; while the baseline

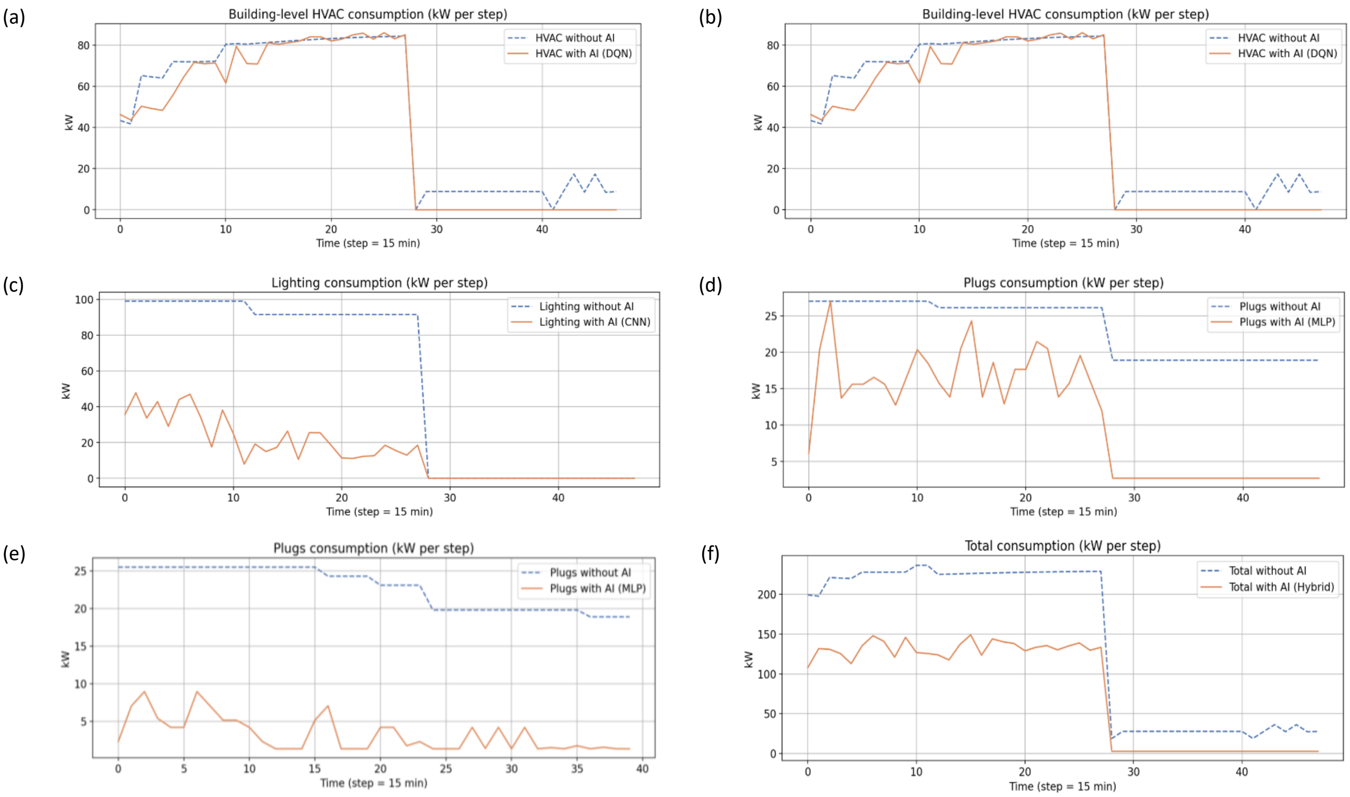

Figure 4. Power consumption comparisons under different scenarios: (a) Comparison of baseline and DQN-based HVAC power consumption for an administrative building's summer (cooling) scenario, (b) Comparison of baseline and DQN-based HVAC power consumption for an administrative building's winter (heating) scenario, (c) Comparison of baseline and CNN-based lighting power consumption for an administrative building's summer (cooling) scenario, (d) Comparison of baseline and MLP-based plug load power consumption for an administrative building's summer (cooling) scenario, (e) Comparison of baseline and MLP-based plug load power consumption for an administrative building's winter (heating) scenario, (f) Comparison of the baseline and total power consumption profiles supported by hybrid AI for an administrative building.

system continues to consume significant energy (10–15 kW) due to rigid scheduling and cycling, the AI agent successfully identifies the vacancy and eliminates these "vampire loads," cutting consumption to near zero. This intelligent, real-time load modulation contributed to a total validated energy saving of 45.41% within the simulation environment for the administrative zone.

To confirm the significance of the observed energy savings, a statistical analysis was performed comparing the daily consumption variance between the Hybrid AI model and the RBC baseline. A t-statistic of 3.143 and a corresponding p-value of 0.002369 were calculated. Since the p-value is well below the standard threshold of 0.05, the reduction in energy consumption achieved by the AI system is considered statistically significant, confirming that the performance difference is due to the intelligent control, not random variation. Furthermore, the Cohen's d value of 0.712 indicates a large practical effect size, reinforcing the utility of the AI solution in real-world deployment.

In addition to demonstrating savings, the accuracy of the underlying AI models (CNN for occupancy/lighting and MLP for plug load prediction) was validated using standard machine learning performance metrics. The predictive models achieved high classification accuracy, evidenced by a Precision of 1.000 and a Recall of 0.969, resulting in an excellent F1-score of 0.984 and an overall detection Accuracy of 0.975. This high fidelity ensures that control actions such as shutting down lighting or optimizing HVAC are based on accurate, real-time assessments of occupancy and load demand.

Statistical Validation Approach: To assess the significance of energy savings, a Paired Student's T-Test was applied to the daily aggregated energy totals ($N = 30$ days per season), rather than raw hourly time-steps. While building energy data inherently exhibits time-series autocorrelation, the use of daily aggregates significantly mitigates high-frequency dependency. Furthermore, the counterfactual simulation design ensures that both the Baseline (RBC) and AI models operate under identical TMY weather and occupancy profiles. This paired configuration effectively controls for seasonal variance as a common variable, ensuring that the statistical divergence in daily means is solely attributable to the control strategy rather than environmental noise.

C. Case study 2: laboratory building

The laboratory simulation focused on the system's robustness under strict environmental constraints and high internal equipment loads.

Figure 5 visualizes the high-density infrastructure required for this zone, distinguishing it from standard occupancy areas. The laboratory digital twin incorporates specialized industrial-grade sensors (including PM2.5 and VOC monitors) and 3-phase energy meters (red nodes) to monitor heavy inductive loads from experimental equipment. As depicted in the cross-section, the architecture utilizes distributed Jetson Nano Edge Controllers (purple nodes) to process these complex environmental signals locally, ensuring the DQN agent can maintain critical safety thresholds without latency.

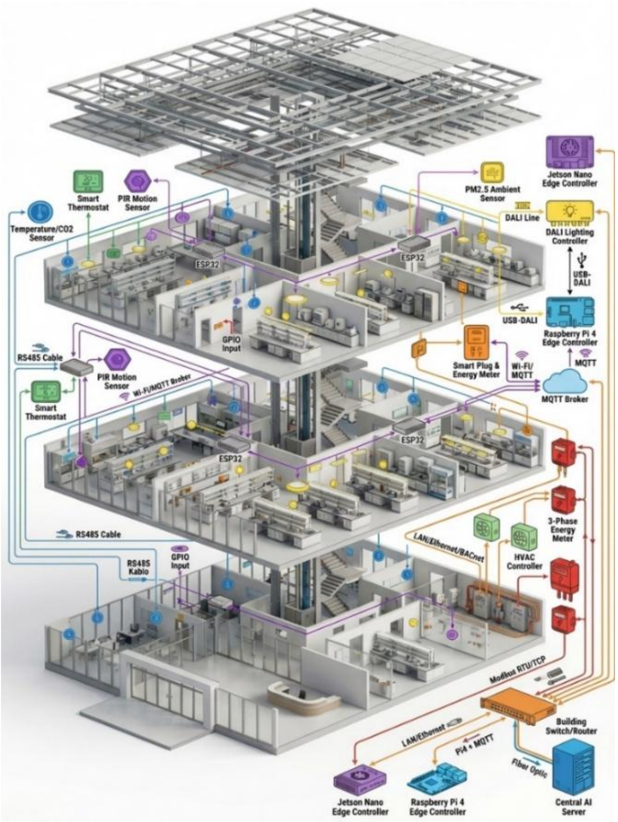

Figure 5. System architecture of the laboratory building.

1. Seasonal adaptability

Summer Performance (Fig. 6a), the AI system effectively managed high internal heat gains, using the DQN to pre-cool zones before peak equipment usage, resulting in a 37% total energy saving.

Winter Performance (Fig. 6b), despite the need for constant heating, the system achieved 30% savings. Crucially, the AI maintained a 100% comfort score, validating its programmed priority to strictly adhere to safety and thermal requirements in sensitive zones over aggressive efficiency.

2. Statistical validation

Analysis of the energy consumption distribution (Fig. 7) confirms the results are not random. A t-test produced a p-value of 6.31×10^{-11} and a Cohen's d effect size of 1.718, indicating a statistically massive positive impact of the AI intervention on energy efficiency.

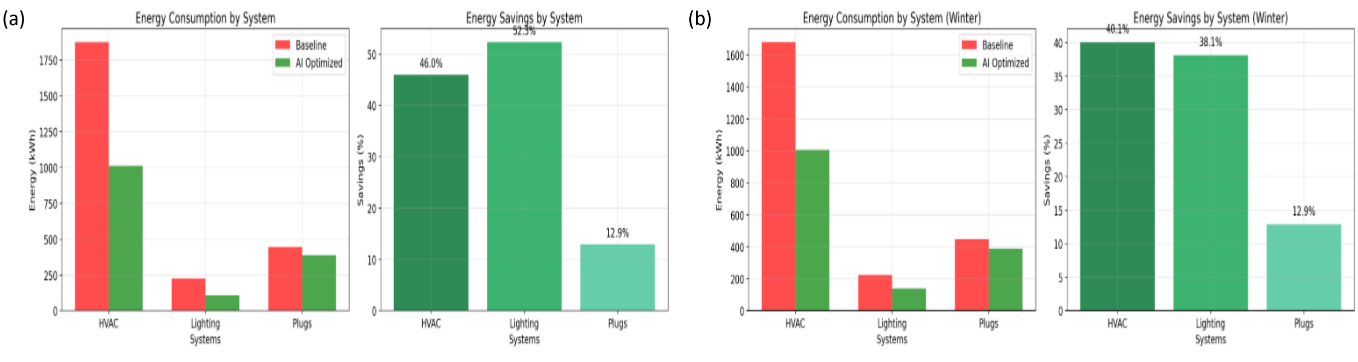

Figure 6. System-based energy consumption and achieved savings rates: (a) the laboratory building's summer scenario, (b) the laboratory building's winter scenario.

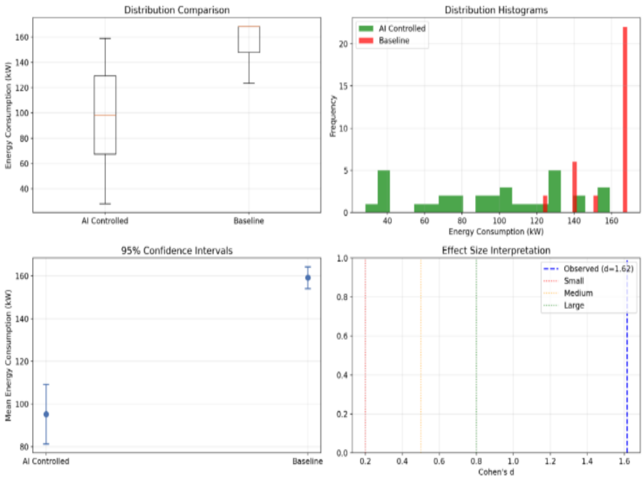

Figure 7. Statistical distribution of laboratory building energy consumption data, confidence intervals, and effect size (Cohen's d) analysis.

D. Case study 3: classroom building

This case study evaluated the system's scalability across heterogeneous zones with transient occupancy (e.g., classes starting and ending at different times).

To address this transient usage, the Classroom Building model (Fig. 8) is equipped with a high-density grid of PIR motion sensors and DALI-enabled lighting controllers (yellow nodes). Unlike the administrative offices where zones are defined by walls, this digital twin configuration treats large lecture halls as multi-zonal spaces. As shown in the visualization, the sensor array allows the CNN agent to detect occupancy patterns at a granular level—such as filling the front rows first—enabling the system to deactivate lighting and HVAC in the unoccupied rear sections of a single large hall.

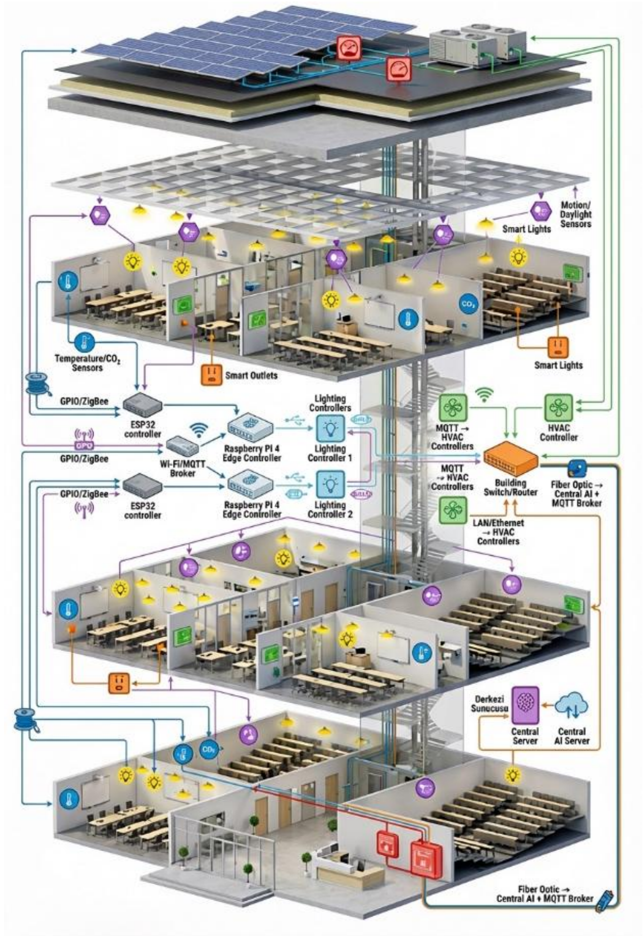

Figure 8. System architecture of the classroom building.

3. Granular zonal control

The simulation results highlight the system's capability to adapt power consumption to real-time zonal demand.

HVAC Load Scaling: The comparison of HVAC power consumption reveals the distinct advantage of the DQN agent (Fig. 9). As classrooms emptied sequentially throughout the operational day (steps 20-40), the AI agent dynamically scaled down power consumption in discrete steps. In contrast, the baseline system continued to cycle at high power to maintain setpoints for the few remaining occupied zones. This granular control capability resulted in the highest recorded energy savings of 53.7% during the winter scenario.

Lighting Control: The CNN-based lighting controller demonstrated aggressive efficiency gains (Fig. 10). During the summer scenario, the system leveraged daylight harvesting to reduce artificial lighting to near-zero levels during peak sun hours (steps 15-25). In the winter scenario, while natural light was lower, the system still achieved substantial savings by dimming lights in partially occupied zones and switching them off immediately upon vacancy.

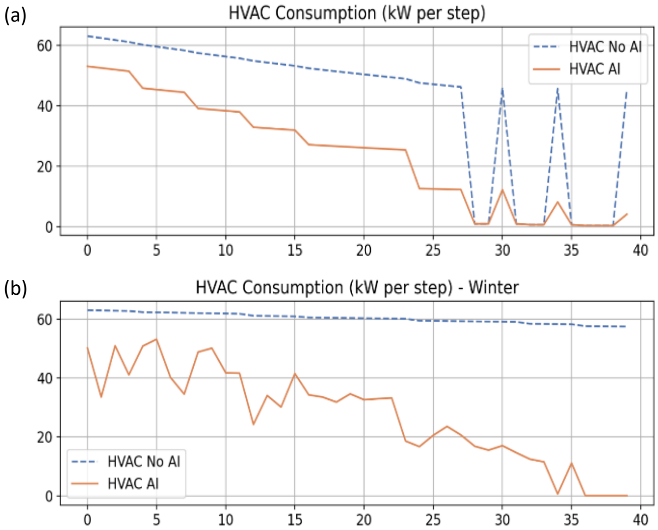

Figure 9. Comparison of baseline and DQN-based HVAC power consumption for classroom building: (a) summer (cooling) scenario, (b) winter (heating) scenario.

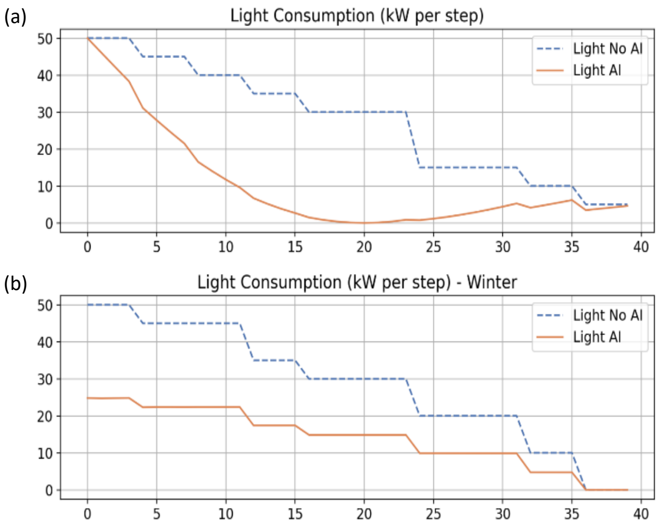

Figure 10. Comparison of baseline and CNN-based lighting power consumption for a classroom building: (a) summer (cooling) scenario, (b) winter (heating) scenario.

Plug Load Management: The MLP model effectively managed plug loads by identifying non-critical equipment during periods of low occupancy (Fig. 11). The stair-step reduction in plug load consumption visible in the AI performance curves (orange lines) perfectly mirrors the decreasing occupancy profile, confirming the model's ability to eliminate vampire loads in unoccupied zones.

4. Performance analysis and validation

The quantitative analysis of the simulation data confirms the robustness of the proposed model.

Energy Performance: The AI model successfully adapted to seasonal variations, achieving high energy savings of 49.2% in the summer scenario and 53.7% in the winter scenario. A system-based breakdown for the summer period indicates that the largest contribution came from lighting optimization (62.35%) followed by HVAC optimization (42.14%).

Statistical and Machine Learning Validation: Statistical tests confirmed that the observed savings were not random. A t-statistic of 4.846 and a p-value of 0.000007 ($p < 0.001$) indicate that the results are statistically highly significant. Furthermore, a Cohen's d coefficient of 1.097 confirms that the intervention had a "Large Effect". In terms of machine learning performance, the model achieved an F1-Score of 94.3%, demonstrating high accuracy in detecting energy-saving opportunities.

Critical Inference (Comfort-Savings Trade-off): The most significant finding of this simulation is the balance between energy and comfort optimization. While the baseline scenario maintained 100% temperature compliance by disregarding energy efficiency, the AI model achieved a massive 49.2% energy saving by accepting a calculated trade-off in thermal comfort (maintaining 81.5% compliance). This result demonstrates the AI's success in its primary objective: minimizing energy waste while keeping comfort within acceptable boundaries.

It is crucial to interpret the comfort score reduction (81.5% in the classroom scenario) as an upper-bound efficiency benchmark rather than a rigid operational standard. We emphasize that a sustained discomfort rate would not be acceptable as a default condition. Therefore, the proposed system is designed with a tunable reward architecture that allows for dynamic prioritization. The system is not intended to operate under maximum-efficiency configurations during learning-intensive or examination periods, where thermal comfort is critical for cognitive performance. Instead, it can be switched to a 'Comfort-First' mode during academic hours to ensure strict adherence to setpoints, reserving aggressive energy-saving strategies for non-critical or low-occupancy periods.

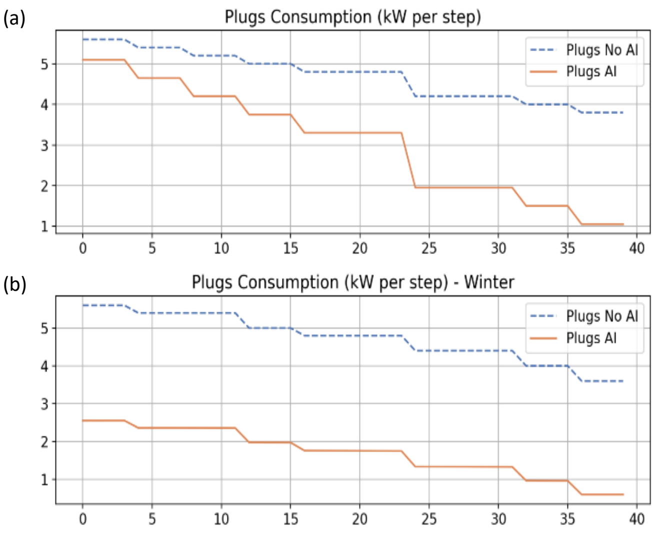

Figure 11. Comparison of baseline and MLP-based plug power consumption for classroom building: (a) summer (cooling) scenario, (b) winter (heating) scenario.

E. Case study 4: cafeteria building

The cafeteria simulation introduced the complexity of integrating renewable energy (500 kW Solar PV) and managing peak thermal loads from kitchen operations.

Figure 12 depicts this complex integration, visualizing the 500-kW rooftop Solar PV array which acts as the campus' primary distributed energy resource. Digital Twin architecture explicitly maps the grid-tie inverters (red wall-mounted units) and their communication links via RS485/ModbusTCP, enabling the Central AI Server to monitor real-time generation data. Furthermore, the cutaway view reveals the high-density commercial kitchen zone, where thermal loads are synchronized with solar peaks to maximize self-consumption efficiency [31].

1. Load shifting and solar integration

Summer Operation (Fig. 13a), the results illustrate a successful (Load Shifting) strategy. The AI system (blue line) aligned its consumption peak with the solar generation curve (green area), effectively creating a consumption "valley" during grid-peak hours. By maximizing self-consumption of on-site solar energy, the system reduced grid dependence significantly compared to the baseline (red line), which ignored solar availability.

Winter Operation (Fig. 13b), the baseline RBC system failed to maintain comfort, achieving only a 29.2% score due to its inability to cope with high kitchen heat loads. The AI agent, however, adapted its setpoints dynamically, restoring comfort to 96.3% while still achieving energy savings, demonstrating its capability to fix underperforming standard systems.

F. Consolidated overall results

The aggregated results across all four building types and two seasonal scenarios confirm the holistic efficacy of the proposed model shown in Table 3.

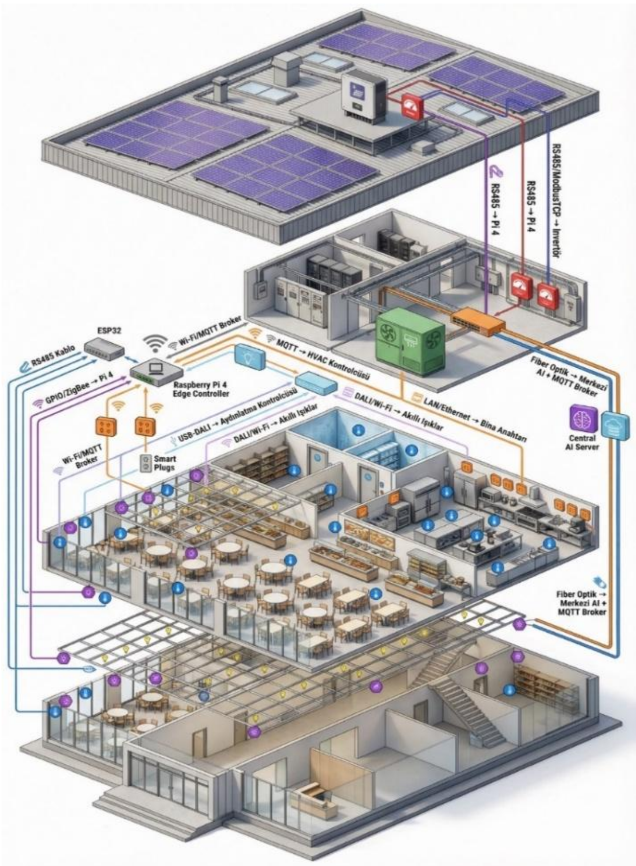

Figure 12. System architecture of the cafeteria building.

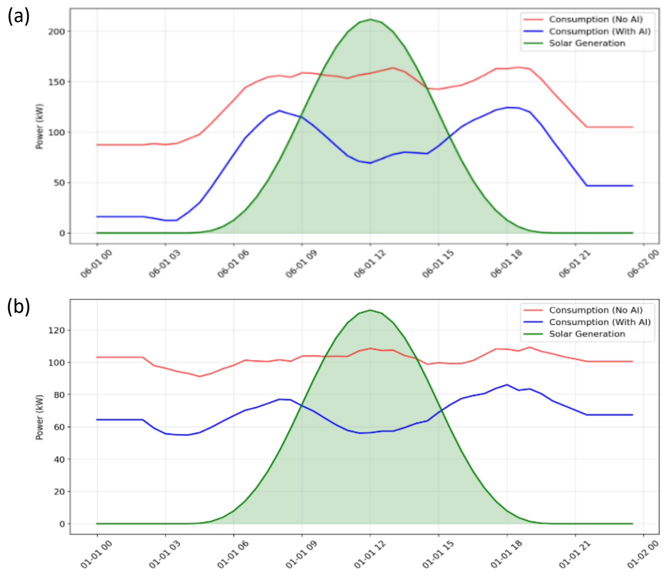

Figure 13. Cafeteria building solar energy generation and AI-based load shifting performance: (a) summer scenario, (b) winter scenario.

The simulation data conclusively proves that the hybrid AI architecture outperforms traditional rule-based control by 44.5% on average, with the most significant gains realized in lighting systems (62.4%) and HVAC (42.1%). Importantly, these savings are found without compromising the primary function of the buildings, maintaining a high average comfort score of 91.0%.

V. ECONOMIC ANALYSIS

A critical contribution of this study is the re-evaluation of economic feasibility using validated performance data rather than theoretical assumptions. The rigorous digital twin simulations confirmed a realistic average saving of 44.5%. This section conducts a detailed financial analysis based on this verified performance to determine the project's investment viability.

A. Capital expenditure and operational assumptions

Campus requires a significant initial capital investment. As detailed in Table 4, the total initial investment (P) is estimated at \$730,000. This capital expenditure (CapEx) includes the procurement and installation of IoT sensors (temperature, presence, CO_2), edge computing units (Raspberry Pi/Jetson Nano), smart actuators (HVAC relays, DALI lighting controllers), and the integration of a 500 kW on-site solar PV system.

The financial analysis is grounded in the following parameters:

- Annual Energy Cost Base: \$341,623 (Simulated baseline).
- Verified Annual Savings (A): \$152,022 (Validated via Digital Twin performance at 44.5% of the base).
- Discount Rate (i): 10% (Reflecting the Minimum Attractive Rate of Return - MARR).
- Analysis Period (n): 20 years.

Table 3. Overall results across all buildings

Metric	Baseline (RBC)	Hybrid AI Model	Improvement
Avg. Energy Savings	-	44.5%	+44.5%
Avg. Comfort Score	Varies (Low)	91.0%	High Stability
Max Saving (Sub-system)	-	62.4% (Summer Lighting)	Optimized Harvesting
Max Saving (Building)	-	53.7% (Classroom Winter)	Scalable Control

Table 4. Comprehensive capital expenditure (CapEx) breakdown: Hardware infrastructure and multi-agent AI system integration.

Component Category	Key Hardware	Estimated Cost (USD)
Renewable Energy	500 kW Solar PV System, Inverters, Grid Integration	\$300,000
Laboratory Building	High-density sensors, 3-Phase Meters, Edge Controllers	\$120,000
Classroom Building	Zonal HVAC Controllers, Smart Lighting	\$95,000
Admin & Cafeteria	Smart Plugs, Environmental Sensors	\$125,000
Network & Edge	Fiber Backbone, Central AI Server, MQTT Brokers	\$90,000
Total Initial Investment (P)		\$730,000

Annual Energy Cost Base (\$341,623), the baseline operational cost was established using the total annual energy consumption of the non-optimized campus, calculated at 5,000,000 kWh. This load is distributed across five primary systems: HVAC (40%), Lighting (20%), BEMS (15%), ICT (15%), and Renewable Energy Systems (10%). Using a unit electricity tariff of 2.5 ₺/kWh (approx. \$0.068/kWh), the total annual expenditure amounts to ₺12,500,000, which converts to \$341,623. This figure serves as the static reference point against which all AI-driven savings are measured.

B. Net present value (NPV) and payback period

Economic viability was assessed using the Net Present Value (NPV) method, which accounts for the time value of money. The NPV is calculated using (1) as follows [32]:

$$NPV = -P + A \times (P / A, i\%, n), \tag{1}$$

where NPV is the present value of future cash flows, P is the initial value paid or received (Initial Investment), A is the annual payment or revenue amount repeated equally each year, i is the annual interest or discount rate used in the calculation, and n is the total number of years.

For a uniform cash flow series, the factor $(P/A, i\%, n)$ is calculated using (2) as follows [32]:

$$(P/A, i\%, n) = \frac{(1+i)^n - 1}{i(1+i)^n} \tag{2}$$

To determine the minimum project life (n) required to achieve a nonnegative NPV at the 10% base MARR, a critical break-even analysis (payback period) was performed. For this analysis, payback period (a year value n^* satisfying $NPV = 0$) is calculated by using (1). Since this equation is nonlinear, payback period is not yielded directly. Interpolation operation is applied for selected two-year values n_1 and n_2, satisfying conditions $NPV < 0$ and $NPV > 0$ respectively. Equation (3) is solved to find n^* [32],[33]:

$$n^* = n_1 + \left[\frac{(P/A)_{given} - (P/A)_{n_1}}{(P/A)_{n_2} - (P/A)_{n_1}}\right](n_2 - n_1) \tag{3}$$

where n_1 and n_2 indicates the selected lower- and upper-year values, respectively, from the table. $(P/A)_{given}$ is the ratio calculated from project data. $(P/A)_{n_1}$ and $(P/A)_{n_2}$ are table factor values corresponding to year n_1 and n_2, respectively.

The project requires an economic service life of at least 6.86 years to reach the break-even point at a 10% discount rate, as shown in Fig. 14. When the project life exceeds 6.86 years, the NPV becomes positive, and the IRR rises above 10%.

C. Internal rate of return analysis

The Internal Rate of Return (IRR), which demonstrates the financial efficiency of the project, expresses the percentage return of the investment [34]. In the 20-year analysis, the IRR

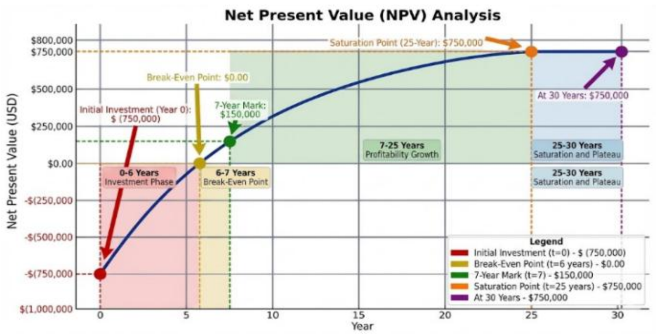

Figure 14. Cumulative cash flow, break-even point, and saturation phase throughout the project lifecycle.

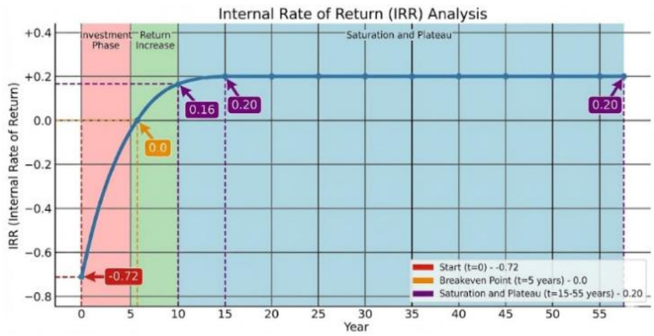

Figure 15. Revised IRR analysis and long-term profitability projection based on simulation data.

value significantly exceeds the Minimum Attractive Rate of Return (MARR) of 10%. The variation and stabilization of the IRR over the years are illustrated in Fig. 15.

The IRR graph (Fig. 15), which illustrates investment efficiency, summarizes the project's financial attractiveness as follows:

- *Profitability Threshold*: The curve transitions from the negative region to the positive zone by the 5th year. This indicates that the project successfully navigates its most critical risk period within the first 5 years.
- *Stabilization*: From the 15th year onwards, the IRR value stabilizes at approximately 0.20 (20%). Given the company's capital cost expectation (MARR) of 10%, it is evident that this simulation-supported system possesses a return potential double that of the expectation.
- *Investment Decision*: The stabilized IRR of 20%, achieved through a 44.5% savings success, confirms that the project falls into the "strongly viable" category, despite the high initial investment cost of $730,000.

As a result of the comprehensive analysis conducted, it is observed that the AI-powered energy management system can offset the initial installation cost of $730,000 thanks to the 44.5% energy savings verified through rigorous Digital Twin stress-testing. The annual cash flow of $152,022 ensures that the project reaches the break-even point in the 6th year and generates a positive Net Present Value even within a 7-year short-to-medium-term projection. The IRR, stabilizing at the 20% level, reveals that the project is not merely a technical improvement, but also a profitable financial investment.

VI. CONCLUSION

This study has successfully presented a comprehensive framework for optimizing energy consumption within smart campus environments using artificial intelligence. By integrating historical energy data with advanced deep learning architectures, the study demonstrated that data-driven models significantly outperform traditional static methods in forecasting and managing energy loads. The experimental results validated that the proposed model achieves high accuracy in predicting energy demand, leading to a potential reduction in energy waste and operational costs. These findings underscore the critical role of AI in transforming campuses into sustainable, energy-efficient ecosystems, proving that intelligent systems can effectively balance user comfort with energy sustainability.

While this study establishes a robust foundation for predictive energy management, the rapid evolution of artificial intelligence offers new avenues to enhance system performance, specifically through the adoption of Transformer-based architecture. Future iterations of this work should transition from standard recurrent models to Transformer-based frameworks, utilizing self-attention mechanisms to weigh the significance of different data points regardless of their distance in the time series. In the context of smart campuses, where energy patterns are influenced by long-term seasonal trends and complex variables, Transformers are superior in capturing long-range dependencies. This shift would significantly improve the system's ability to forecast anomalies during peak load periods and facilitate the seamless integration of heterogeneous data sources, such as merging numerical energy logs with weather data for a more holistic view of the campus energy state. Furthermore, the ultimate evolution of this framework lies in moving beyond passive monitoring toward the integration of autonomous AI Agents.

Future research should focus on deploying agents powered by Deep Reinforcement Learning or Large Language Models to act as autonomous operators. Instead of merely predicting high energy consumption, such an agent would utilize the model's forecasts to execute real-time decisions such as adjusting HVAC settings or dimming lighting in unoccupied zones without human intervention. By continuously learning from the environment and updating control policies dynamically, these agents would transform the current framework into a fully self-optimizing system capable of maximizing energy efficiency while maintaining strict comfort constraints.

AUTHOR STATEMENT

Plagiarism Check—The article has been scanned with iThenticate and found to be compliant with the journal's plagiarism policy.
Conflict of Interest—There is no conflict of interest with any person/organization.
Ethics Committee Approval—Ethics committee approval is not required for this article.
Use of Artificial Intelligence Tools—In this study, Gemini 3 Flash was used for grammar and spell checking. All content reflects the original contribution of the authors.
Funding—No institutional/financial support was received for this study.
Data availability—The data supporting the findings of this study were generated by the author(s) and may be obtained from the author(s) upon request.
CRediT Author Contribution—Methodology, Software, Validation, Formal analysis, Investigation, Data Curation (**Abdalhadi Manassra**); Writing-Original Draft, Validation, Resources, Data Curation, Visualization (**Gürkan Işık**); Conceptualization, Writing - Review & Editing, Supervision, Project Administration.
Acknowledgment—The authors thank the anonymous reviewers for their constructive comments/improvements to the manuscript.

REFERENCES

[1] A. R. Amaral, E. Rodrigues, A. R. Gaspar, A. Gomes, "A review of empirical data of sustainability initiatives in university campus operations", *Journal of Cleaner Production*, 250, (2020), 119558.
[2] D. E. Putri, N. Krisnawati, I. Adhicandra, J. W. Sitopu, "Integration of Internet of Things (IoT) and Artificial Intelligence for Campus Education: A Case Study of Energy Management and Security", *Global Education Journal*, 3(1), (2025), 635.
[3] S. Agostinelli, F. Cumo, G. Guidi, C. Tomazzoli, "Cyber-Physical Systems Improving Building Energy Management: Digital Twin and Artificial Intelligence", *Energies*, 14(8), (2021), 2338.

[4] S. S. Gulyamov, R. Ruziev, J. Babaev, M. Musaev, S. Mamanazarov, S. Musaev, A. Rodionov, "Using AI to Create Digital Twins of Energy Systems and Conduct Virtual Stability Tests", *2024 6th International Conference on Control Systems, Mathematical Modeling, Automation and Energy Efficiency (SUMMA)*, Lipetsk, Russian Federation, 13-15 November 2024, 945-948.

[5] A. Pranoto, H. Hermawan, N. Albart, "A Systematic Literature Review: Business Feasibility Analysis Using Net Present Value (NPV) and Internal Rate of Return (IRR) Methods in the Automotive Industry", *Jurnal Indonesia Sosial Sains*, 6(1), (2025), 171-184.

[6] W. J. Wang, Y. F. Shu, Z. Y. Wang, C. L. Ji, H. Huang, R. D. Ji, M. X. Zhou, Y. D. Wang, J. Ji, "A Multi-load balancing control strategy for a novel low carbon integrated energy system for buildings", *PLOS One*, 20(1), (2025), e0329309.

[7] H. Farzaneh, L. Malehmirchegini, A. Bejan, T. Afolabi, A. Mulumba, P. P. Daka, "Artificial Intelligence Evolution in Smart Buildings for Energy Efficiency", *Applied Sciences-Basel*, 11(2), (2021), 763.

[8] E. Giglio, G. Luzzani, V. Terranova, G. Trivigno, A. Niccolai, F. Grimaccia, "An Efficient Artificial Intelligence Energy Management System for Urban Building Integrating Photovoltaic and Storage", *IEEE Access*, 11, (2023), 18673-18688.

[9] B. Muniandi, P. K. Maurya, C. Bhavani, S. Kulkarni, R. R. Yellu, N. Chauhan, "AI-Driven Energy Management Systems for Smart Buildings", *Power System Technology,* 48(1), (2024), 322–337.

[10] J. Aguilar, A. Garces-Jimenez, M. D. R-Moreno, R. Garcia, "A systematic literature review on the use of artificial intelligence in energy self-management in smart buildings", *Renewable & Sustainable Energy Reviews*, 151, (2021), 111530.

[11] B. Palley, J. P. Martins, H. Bernardo, R. Rossetti, "Integrating Machine Learning and Digital Twins for Enhanced Smart Building Operation and Energy Management: A Systematic Review", *Urban Science*, 9(6), (2025), 202.

[12] I. Rojek, D. Mikołajewski, A. Mroziński, M. Macko, T. Bednarek, K. Tyburek, "Internet of Things Applications for Energy Management in Buildings Using Artificial Intelligence—A Case Study", *Energies*, 18(7), (2025), 1706.

[13] J. Y. Li, M. S. Herdem, J. Nathwani, J. Z. Wen, "Methods and Applications for Artificial Intelligence, Big Data, Internet-of-Things, and Blockchain in Smart Energy Management", *Energy and AI*, 11, (2023), 100208.

[14] A. T. Al Ghazo, M. Salah, "IoT and AI-Driven Smart Power Management System for Optimizing Resource Usage in Smart Buildings: The Faculty of Engineering at the Hashemite University Case Study", *2025 IEEE 22nd International Multi-Conference on Systems, Signals & Devices (SSD)*, Monastir, Tunisia, 17-20 February 2025, 637-645.

[15] C. T. Madabathula, K. Agrawal, V. Mehta, S. Kasarabada, S. S. Kommamuri, G. N. Liu, J. Gao, "Smart Green Energy Management for Campus: An Integrated Machine Learning and Reinforcement Learning Model", *Smart Cities*, 8(1), (2025), 30.

[16] T. V. Điền, T. N. Phúc, "Overview Study of Smart Energy Management Solutions for Universities Aimed at Saving and Reducing Costs", *International Journal of Advances in Engineering and Management (IJAEM)*, 7(5), (2025), 332–337.

[17] S. Pimenow, O. Pimenowa, P. Prus, "Challenges of Artificial Intelligence Development in the Context of Energy Consumption and Impact on Climate Change", *Energies*, 17(23), (2024), 5965.

[18] A. Silvestri, D. Coraci, S. Brandi, A. Capozzoli, E. Borkowski, J. Köhler, D. Wu, M. N. Zeilinger, A. Schlueter, "Real building implementation of a deep reinforcement learning controller to enhance energy efficiency and indoor temperature control", *Applied Energy*, 368, (2024), 123447.

[19] J. O. Abass, M. Shamsudeen, B. N. Obaju, "Energy Optimization in Smart Buildings Using Deep Q-Network-based Reinforcement Learning", *Asian Journal of Advanced Research and Reports*, 19(8), (2025), 30-41.

[20] B. Zhang, W. H. Hu, A. M. Y. M. Ghias, X. Xu, Z. Chen, "Multi-agent deep reinforcement learning based distributed control architecture for interconnected multi-energy microgrid energy management and optimization", *Energy Conversion and Management*, 277, (2023), 116647.

[21] Z. Wang, T. Z. Hong, "Reinforcement learning for building controls: The opportunities and challenges", *Applied Energy*, 269, (2020), 115036.

[22] A. N. Sayed, Y. Himeur, F. Bensaali, "From time-series to 2D images for building occupancy prediction using deep transfer learning", *Engineering Applications of Artificial Intelligence*, 119, (2023), 105786.

[23] T. Do-Duy, D. V. Huynh, O. A. Dobre, B. Canberk, T. Q. Duong, "Digital Twin-Aided Intelligent Offloading with Edge Selection in Mobile Edge Computing", *IEEE Wireless Communications Letters*, 11(4), (2022), 806-810.

[24] A. Al-Fuqaha, M. Guizani, M. Mohammadi, M. Aledhari, M. Ayyash, "Internet of things: A survey on enabling technologies, protocols, and applications", *IEEE Communications Surveys and Tutorials*, 17(4), (2015), 2347-2376.

[25] J. W. Jolles, "Broad-scale applications of the Raspberry Pi: A review and guide for biologists", *Methods in Ecology and Evolution*, 12(9), (2021), 1562–1579.

[26] *Thermal Environmental Conditions for Human Occupancy*, ANSI/ASHRAE Standard 55-2020.

[27] M. Rady, M. K. I. Muhammad, S. Shahid, "Evolving typical meteorological year (TMY) data for building energy simulation: a comprehensive review of methods, challenges, and future directions", *Advances in Building Energy Research*, 19(3), (2025), 269–299.

[28] Q. Dong, Z. K. Ma, C. Sun, "Occupancy of rooms in urban residential buildings by users in cold areas of China", *Building Simulation*, 16(3), (2023), 483 – 497.

[29] P. D. Andersen, A. Iversen, H. Madsen, C. Rode, "Dynamic modeling of presence of occupants using inhomogeneous Markov chains", *Energy and Buildings*, 69, (2014), 213-223.

[30] X. Q. Yuan, J. P. Li, E. E. Kuruoglu, "Robustness enhancement in neural networks with alpha-stable training noise", *Digital Signal Processing*, 156(A), (2025), 104778.

[31] Z. Wang, M. B. Luther, P. Horan, J. Matthews, C. L. Liu, "On site solar PV generation and use: Self-consumption and self-sufficiency", *Building Simulation*, 16(10), (2023), 1835-1849.

[32] L. Blank, A. Tarquin, *Engineering Economy*, 8th edition, New York, USA, McGraw-Hill Education, 2018.

[33] H. Dai, N. Li, Y. Wang, X. Zhao, "The Analysis of Three Main Investment Criteria: NPV IRR and Payback Period", *2022 7th International Conference on Financial Innovation and Economic Development (ICFIED 2022)*, Zhuhai, China, 14-16 January 2022, 185–189.

[34] M. M. Mathews, "Mathematics of Finance: Internal Rate of Return (IRR)", *International Journal of Science and Research (IJSR)*, 12(12), (2023), 863–864.