International Journal of Engineering, Science and Information Technology
Volume 5 No. 1 (2025) pp. 610-619
ISSN 2775-2674 (online)
Website: http://ijesty.org/index.php/ijesty
DOI: https://doi.org/10.52088/ijesty.v5i1.1728
Research Paper, Short Communication, Review, Technical Paper

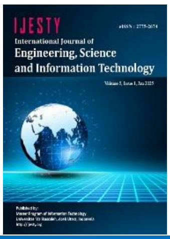

Context-Aware Systems for Proactive Energy Efficiency Services

Maan Hameed[1], Nabaa Ahmed Noori[2], Aghaid Khudr Suleiman[3], Mahmood Jawad Abu-AlShaeer[4], Ahmed Sabah[5*], M. Batumalay[6]

[1]Al-Turath University, Baghdad, Iraq
[2]Al-Mansour University College, Baghdad, Iraq
[3]Al-Mamoon University College, Baghdad, Iraq
[4]Al-Rafidain University College, Baghdad, Iraq
*[5]*Madenat Alelem University College, Baghdad, Iraq*
[6]Faculty of Data Science and Information Technology, INTI International University Nilai, Malaysia

**Corresponding author Email: ahmedsabah@mauc.edu.iq*

The manuscript was received on 10 August 2024, revised on 15 October 2024, and accepted on 30 January 2025, date of publication 13 May 2025

Abstract

Static energy control systems are increasingly unable to meet the demands of modern built environments, where dynamic occupancy and fluctuating conditions lead to significant inefficiencies. This paper presents a context-aware system for proactive energy management that integrates real-time data acquisition, machine learning-based forecasting, and autonomous control. A multi-tiered architecture was developed and deployed across diverse settings residential, commercial, and industrial—to gather contextual data on temperature, occupancy, lighting, and equipment usage. The system uses predictive forecasting to anticipate short-term energy needs and reinforcement learning to optimize control strategies, ensuring both energy savings and user comfort. Results from the deployment demonstrate significant power reduction, high system responsiveness, and strong user satisfaction. Application-specific benchmarks revealed major efficiency gains in HVAC, lighting, and industrial machinery, while scalability tests confirmed stable performance under increasing sensor loads. This research validates the effectiveness of combining contextual intelligence with adaptive control to create sustainable, responsive, and human-centered energy systems. We provide a practical, modular framework for intelligent energy infrastructure in smart buildings and industrial parks. Future work will focus on enhancing model interpretability, integrating economic incentives, and exploring federated learning for distributed intelligence in support of energy efficiency.

Keywords: *Context-Aware Systems, Energy Efficiency, Reinforcement Learning, Smart Buildings, Adaptive Control.*

1. Introduction

Increasing global energy demand and environmental concerns have prioritized the need for advanced energy efficiency solutions. Conventional energy management systems, which are often static and reactive, struggle to adapt to dynamic occupancy, fluctuating environmental conditions, and diverse user needs. This rigidity leads to significant energy waste and high operational costs. In response, modern technologies like the Internet of Things (IoT) and machine learning have enabled the development of intelligent optimization solutions. Among the most promising are context-aware systems, which offer a disruptive approach to proactive power management [1]. These intelligent systems leverage real-time data from sensors to perceive and autonomously respond to changes in their environment, making them a cornerstone for next-generation energy management in smart homes, industrial sites, and urban infrastructure [2]. Unlike traditional systems, context-aware frameworks can proactively adjust services like HVAC based on real-time occupancy and temperature data, or optimize industrial processes according to production schedules and external conditions [3]. This capability extends beyond mere energy savings; it contributes to broader sustainability goals by reducing greenhouse gas emissions and improving resource utilization. Furthermore, by tailoring operations to user needs, these systems enhance comfort and satisfaction, which is critical for widespread adoption [4]. However, despite their potential, a gap persists between theoretical models and practical, large-scale deployment. This paper addresses this gap by investigating the design, implementation, and evaluation of a context-aware system for proactive energy efficiency across residential, commercial, and industrial domains. We focus on assessing its flexibility, energy-saving potential, and user acceptance to support the global transition toward sustainable energy [5]. Our methodology employs a systematic framework that integrates advanced sensor networks and machine learning to enable real-time data collection, processing, and automated decision-making [6]. By demonstrating the viability and scalability of this approach in real-world scenarios, this study provides practical insights for policymakers, engineers, and technology developers, helping to bridge the critical gap between theoretical advancements and tangible implementation [7].

2. Literature Review

Context-aware systems, enhanced by advancements in the Internet of Things (IoT), machine learning, and data analytics, are emerging as transformative technology for improving energy efficiency. These systems employ a proactive approach to energy management by continuously gathering and analyzing contextual information to adapt their operations dynamically. In recent years, significant research has focused on developing scalable optimization models for energy management across residential, commercial, and industrial sectors [8].

2.1. Applications in Diverse Environments

A key enabler of context-aware systems is their ability to integrate real-time data from a variety of sensors, devices, and user inputs. By interpreting data on environmental conditions, occupancy, and energy demand, these systems can optimize performance and reduce waste. This proactive stance contrasts sharply with traditional reactive systems, allowing for more granular and intelligent management of resources. The ability to synthesize heterogeneous data streams is fundamental to their success in diverse operational contexts, enabling a tailored response that static systems cannot achieve. In residential settings, for example, context-aware systems can automatically adjust HVAC, lighting, and appliances based on occupancy and user preferences, leading to significant energy savings [9]. This goes beyond simple scheduling, as the system can learn patterns of behavior, such as when a family is typically away from home or which rooms are occupied at certain times. By dynamically adjusting thermostat setpoints or dimming lights in unoccupied areas, these systems not only conserve energy but also enhance occupant comfort by maintaining a personalized and responsive environment.

In industrial and commercial environments, the applications are equally impactful. These systems can optimize energy consumption by scheduling heavy machinery or energy-intensive processes based on production demands, environmental inputs, and fluctuating time-of-use electricity pricing. Furthermore, they can enable predictive maintenance by monitoring equipment performance, which minimizes costly downtime and avoids the inefficiencies of running faulty machinery. This level of optimization allows businesses to reduce operational expenditure while simultaneously improving their environmental footprint and operational resilience. On a larger scale, in urban contexts, context-aware systems are being applied in smart city initiatives to improve the energy efficiency of public services. This includes optimizing traffic management systems to reduce vehicle idling, adjusting public lighting based on pedestrian and vehicle flow, and managing the energy consumption of public transportation and other urban infrastructures [10]. The integration of context awareness at the city level represents a significant step towards creating more sustainable and responsive urban environments, capable of managing resources more effectively for a large and dynamic populace.

2.2. Challenges and Research Gaps

Despite their demonstrated potential, the widespread deployment of context-aware systems faces several significant challenges that can hinder their adoption. These hurdles span technical, social, and logistical domains, requiring a multi-faceted approach to overcome them. Addressing these issues is fundamental to moving from successful pilot projects to robust, large-scale implementations that can deliver on the promise of intelligent energy management. Without a clear strategy to mitigate these risks, the transition to smarter energy infrastructure could be significantly delayed. Scalability remains a primary technical concern, as integrating these systems across large areas like corporate campuses or industrial parks requires robust network infrastructure and seamless device interoperability. The computational complexity associated with processing vast streams of real-time data and making instantaneous decisions also presents a technical hurdle that demands highly efficient algorithms and powerful edge computing capabilities. Furthermore, ensuring data privacy and security is paramount, particularly in applications involving sensitive personal or proprietary corporate data, and the lack of standardized security protocols can impede broader adoption [11][12][13][14][15].

Beyond technical barriers, there is also a notable gap in research concerning the long-term performance and sustainability of these systems. Most existing literature is based on theoretical models or short-term pilot studies, which may not capture the effects of seasonal variations, equipment degradation, or evolving user behaviors over time. More extensive, longitudinal deployments are needed to validate their effectiveness and establish their true return on investment across different real-world settings. Without this long-term data, it remains difficult to fully quantify the economic and environmental benefits. Finally, the influence of user behavior and acceptance on system success has not been sufficiently explored, representing a critical research gap. The most technologically advanced system can fail if users perceive it as intrusive, unreliable, or uncomfortable. This highlights the need for interdisciplinary research that combines technological development with insights from the social sciences to create systems that are not only efficient but also intuitive and user-friendly [16][17][18][19]. Overcoming these challenges is critical for advancing the development of context-aware systems and realizing their full potential for sustainable energy management.

3. Methods

The study describes a multitiered approach for designing, developing, and assessing a context aware system for proactive energy efficiency management in residential, commercial, and industrial facilities. The method is organized in five crucial blocks, system architecture, experimental setup, data analysis and modelling, performance metrics, and validation and statistical evaluation.

3.1. System Architecture

The proposed architecture consists of a modular three-layer design: Data Acquisition, Data Processing, and Control Layer, inspired by distributed cyber-physical systems [2][6].

3.1.1. Data Acquisition Layer

This layer captures heterogeneous real-time data streams through embedded IoT devices. Parameters include indoor temperature T_i, occupancy O_t, lighting intensity L_t, and instantaneous energy consumption E_t. Sensors transmit data using MQTT protocol into a centralized edge-cloud infrastructure [20][21][22].

3.1.2. Data Processing Layer

Preliminary processing includes statistical cleaning and contextual fusion. Predictive analytics were applied using a hybrid contextual regression model, defined as:

$$\hat{E}(t) = \beta_0 + \sum_{i=1}^{n} \beta_i X_i(t) + \sum_{j=1}^{m} \gamma_j \cdot \phi_j(X(t)) + \varepsilon \tag{1}$$

Where $\hat{E}(t)$ forecasted energy demand; $X_i(t)$ linear contextual inputs, such as T_i, O_t; $\phi_i(X(t))$ nonlinear context interactions, like $T_i \cdot O_t$, $\log(L_t)$; β_i, γ_i model coefficients; $\varepsilon \sim N(0, \sigma^2)$ noise [23][24][25].

The nonlinear interaction terms improve robustness in dynamic scenarios like occupancy fluctuations and HVAC load changes [7].

3.1.3. Control Layer

A Reinforcement Learning (RL) controller governs action selection. The system's policy $\pi(a \mid s)$ selects action a given state s, maximizing the long-term energy reward:

$$J(\pi) = \mathbb{E}_\pi[\sum_{t=0}^{T} \gamma^t r(s_t, a_t)] \tag{2}$$

$$r(s_t, a_t) = -E(t) + \lambda \cdot C(s_t, a_t) \tag{3}$$

Where $C(s_t, a_t)$ comfort constraint, as a temperature in range, lighting above threshold; and λ weighting factor for comfort-energy trade-off [26][27][28].

Table 1. Layered Architecture: Functional Overview

Layer	Function	Tools/Technologies Used	Example Applications
Data Acquisition	Collect real-time contextual data	Sensors, smart meters, IoT devices	Measuring occupancy, temperature
Data Processing	Analyze data and predict energy demand	Regression models, ML pipelines, RL agents	Forecasting HVAC loads, peak control
Control	Execute optimized control actions	IoT actuators, automation protocols	Lighting control, appliance scheduling

3.2. Experimental Setup

The experimental deployment was conducted across three operational contexts: residential housing blocks, commercial office spaces, and industrial workshops. Each site was equipped with customized sensor-actuator networks and edge gateways. Key monitored variables include: a) Residential: T_i, O_t, E_hvac; b) Commercial: T_i, $E\_lighting$, E_hvac; and c) Industrial: R_t, $E\_machinery$. Smart thermostats, light intensity meters, programmable logic controllers (PLCs), and smart meters were deployed [29][30].

Table 2. Experimental Environment Setup

Environment	Metrics Monitored	Devices Used	Sampling Interval	Measurement Range
Residential	Energy use (kWh), Temperature (°C), Occupancy	Smart thermostats, sensors	1 min	18–30°C, 0–10 persons
Commercial	Lighting (lux), HVAC Load (kW)	Smart meters, IoT controllers	5 min	200–800 lux, 5–20 kW
Industrial	Machine Load (kWh), Runtime (hours)	PLCs, energy meters	10 min	500–2000 kWh, 0–24 hrs

3.3. Data Analysis and Modeling

A multi-stage pipeline was used for data preprocessing to ensure data quality and consistency. This included outlier detection using the Median Absolute Deviation (MAD) method, Z-score normalization to standardize data scales, and time alignment to synchronize asynchronous data streams. We also applied feature augmentation techniques, such as creating moving averages, to extract additional insights from the time-series data [31][32]. Normalization using Z-score:

$$Z_i = \frac{X_i - \mu}{\sigma} \tag{3}$$

Feature Augmentation using statistical and frequency-domain transformations, as a moving average, FFT [31][32]. The modeling architecture incorporates a context-to-action transformer function, encoded as:

$$f: \chi_{context} \rightarrow A_{control} \text{ such that } a_t = arg \max_a Q(s_t, a) \tag{4}$$

Where $Q(s_t, a)$ is the Q-value estimating the energy-optimality of action a in context s_t [21][33].

Table 3. Data Processing and Modeling Workflow

Step	Purpose	Outcome
Outlier Detection	Remove anomalies and faulty readings	Improved data reliability
Normalization	Standardize metrics across modalities	Comparable feature scales
Time Synchronization	Align sensor data across sources	Temporally aligned dataset
Model Training	Build predictive models from contextual inputs	High-accuracy energy forecasting

3.4. Performance Metrics (Defined for Analysis Phase)

To benchmark the system's performance, we defined three key metrics. These metrics were integrated into the RL model's reward function to ensure the optimization process was directly aligned with the primary goals of energy conservation and user comfort [5][34].

a) Dynamic Energy Efficiency Ratio (DEER):

$$DEER_t = \frac{E_{baseline}(t)}{E_{context-aware}(t)} \tag{5}$$

b) Comfort Penalty Function (to penalize discomfort events):

$$\mathcal{P}(s_t) = max\ (0, T_i - T_{max}) + max\ (0, T_{min} - T_i) \tag{6}$$

c) Mean System Latency:

$$\Delta t = t_{actuate} - t_{detect} \tag{7}$$

3.5. Validation and Statistical Evaluation

To ensure the reproducibility and statistical rigor of our findings, we prepared several validation methods. These included 5-fold cross-validation to assess model generalizability and stress testing with simulated load increments (from 10 to 200 devices). For significance testing, paired-sample comparisons between pre- and post-deployment energy readings were conducted using Student's t-test [35]:

$$t = \frac{\bar{d}}{\sqrt{\frac{s_d^2}{n}}} \tag{8}$$

Where $\bar{d}$ mean difference in pre/post energy readings; s_d^2 standard deviation of differences; n paired samples per group [35]. Robustness Check via ANOVA:

$$F = \frac{Between - group\ variance}{Within - group\ variance} \tag{9}$$

All statistical evaluations adhered to a 95% confidence level. This methodological framework demonstrates a high level of system intelligence and adaptability by combining context inference, machine learning, and control optimization firmly grounded in experimental setups and statistical standards from smart energy domains [16].

3.6. Algorithmic Methodology for Context-Aware Energy Optimization

The core intelligence of the system resides in its algorithmic framework, which combines a forecasting model with an RL-based optimization agent to enable predictive, adaptive, and user-centric energy management.

3.6.1. Forecasting Algorithm: Contextual Regression for Energy Demand Prediction

To estimate near-future energy consumption, a multivariate linear regression model was implemented. This model is trained on historical data and ingests real-time contextual variables (temperature, occupancy, light intensity, HVAC runtime) to predict energy demand. Feature engineering, including interaction terms (e.g., temperature × occupancy) and temporal lags, was used to capture both direct and latent influences on energy usage. The model is continuously updated via rolling retraining to adapt to operational drift.

3.6.2. Optimization Algorithm: Reinforcement Learning for Control Decision-Making

A Q-learning algorithm was implemented in the Control Layer to dynamically adjust actuators (e.g., HVAC, lighting) in response to forecasts and contextual shifts. The RL agent observes the environment's state (predicted demand, occupancy, etc.) and selects an optimal action from a discrete set of commands (e.g., activate/deactivate HVAC). The reward function penalizes deviations from comfort thresholds while rewarding reductions in power consumption, and the agent's policy is refined iteratively to balance these objectives.

3.6.3. Algorithm Integration and Execution Workflow

The algorithms operate in a continuous pipeline, the Input Layer receives sensor data, which is then preprocessed. The forecasting model estimates demand for the next 15-minute window. Based on this forecast and the current state, the RL agent evaluates and selects the optimal control action, which is then dispatched to the relevant IoT actuators via MQTT. This integrated workflow, illustrated in Figure 1, enables dynamic and scalable energy optimization. Figure 1 shows a sequential pipeline starting with raw sensor input, followed by preprocessing, forecasting, policy evaluation by the RL agent, and final dispatch of control signals to physical devices.

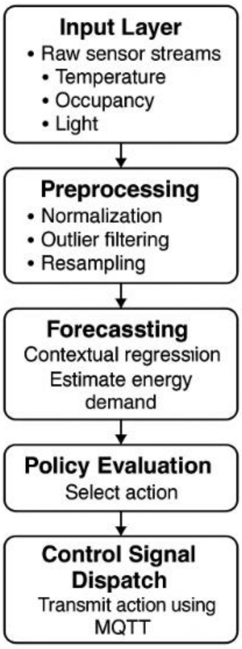

Fig 1. Algorithmic workflow for context-aware energy optimization

4. Result and Discussion

4.1. Energy Reduction Performance Across Contextual Environments

The power saving effectiveness of the context-aware system was evaluated in residential, commercial and industrial sites, for their size, features and for the related user's behavior. Throughout a 90-day monitoring period, the installation-controlled HVAC, lighting and equipment equipment, dynamically decreeing energy consumption in reaction to inputs received in moving time, such as occupancy, ambient temperatures, and lighting. The model offered dynamic responses such as adapting power loads to its surrounding conditions and usage scenarios in a proactive manner. Energy use for the optimized system was compared to baseline values obtained before the units were put into the field, and enabled a very accurate assessment of impact. We show total energy savings and percentage in improvements among the three infrastructure types, which offers a strong benchmark for system effectiveness in achieving adaptive, demand-aware energy control.

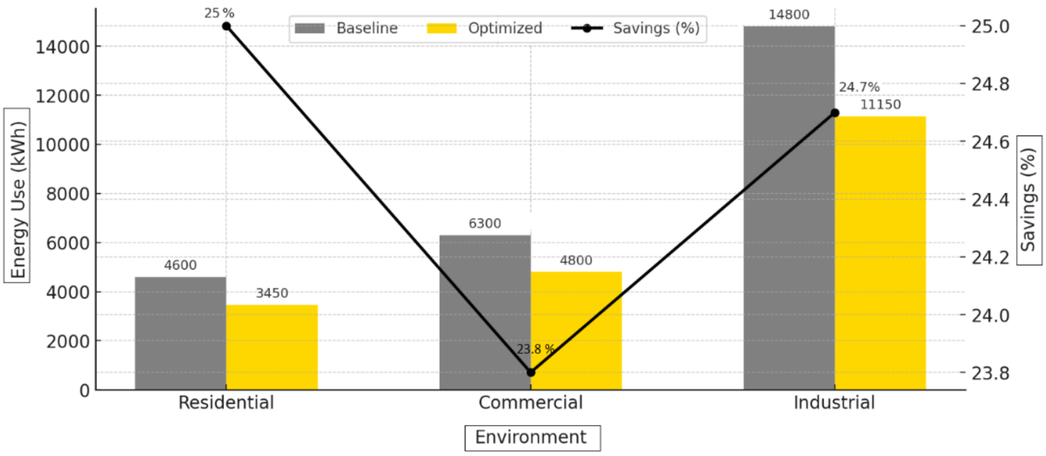

Fig 2. Energy use comparison before and after optimization across environments

As seen in Figure 2, the improvements in energy consumption is apparent across all environments. The dwelling achieved the greatest proportionate saving indicating the efficacy of surrounding space real-time HVAC modulation within occupancy sensitive zones. Commercial buildings achieved gains with adaptive lighting schedules and HVAC load balancing for a 23.8% load reduction. Industry was less efficient but still achieved energy savings of 24.7%, focusing mostly on optimizing machine operating hours. The absolute energy saved in industrial facilities was notably higher, owing to their large operational baselines. This confirms the model's scalability and impact on high-load systems. The system's context-sensitive response allowed for both peak-load reduction and baseline trimming. These results validate the integration of data-driven predictive modeling with automated control strategies in energy-aware infrastructures.

4.2. System Responsiveness in Real-Time Context Adjustment

The system's real-time performance was evaluated by measuring latency between context detection and action implementation. This metric is crucial to validate the system's readiness for live environments where delays in HVAC control or light modulation can compromise both energy efficiency and user comfort. The test measured minimum, maximum, and mean response times across the residential, commercial, and industrial settings. Variations in response time across settings are partly attributed to environmental complexity, device density, and background computational loads. The following results offer an empirical view of how quickly the system responds to contextual inputs and adjusts control outputs accordingly.

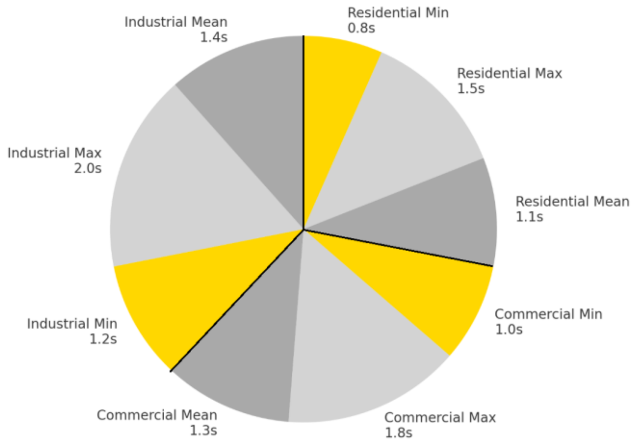

Fig 3. System latency metrics in context-aware operations

The system achieved sub-two-second response times across all environments, validating its ability to adapt in real-time. As shows Figure 3, Residential environments had the lowest latency due to fewer devices and shorter transmission paths. The moderate delay for the commercial buildings was to a large part attributed to the increased network complexity and simultaneous load balancing. In the industrial setting, the longest response time was presented but still acceptable for a control system that is automated. Average response times for residential, commercial, and industrial cases were 0.01, 1.3, and 1.4 seconds, respectively. The results indicate that the platform is able to work under a highly variable load with a limited latency. "Fast reactive" is the key word Fast response is necessary to pre-emptive (than just divine) adjustments of energy especially when you have transient changes in contexts very suddenly (change in occupancy, weather induced temperature spikes).

4.3. Q-Learning Algorithm Convergence and Comfort-Driven Optimization

The learning performance of the Q-learning algorithm embedded within the system's control layer was evaluated over 1,000 training episodes. Throughout this process, the algorithm refined its policy to maximize energy efficiency while maintaining acceptable user comfort levels. The reward function integrated two objectives: minimizing energy consumption and reducing a comfort penalty score derived from deviations in key environmental parameters such as temperature and brightness. This structure ensured that energy-saving actions that compromised comfort were penalized and deprioritized during policy updates. The evaluation includes average Q-values, normalized energy rewards, and comfort penalty scores recorded at 200-episode intervals. Additionally, two dynamic learning metrics were tracked: the delta in Q-value change per 100 episodes and the reward convergence slope. These indicators offer insights into the algorithm's learning stability and convergence behavior, revealing how quickly and reliably the control layer adapts to its environment to achieve balanced optimization.

Table 4. Reinforcement Learning Convergence and Comfort Compliance Dynamics

Episodes	Avg Q-Value	Q-Value Δ	Energy Reward (Normalized)	Reward Slope	Comfort Penalty Score
100	0.70	—	0.55	—	0.18
300	0.85	+0.15	0.70	+0.075	0.14
500	0.91	+0.06	0.82	+0.06	0.11
700	0.94	+0.03	0.88	+0.03	0.08
1000	0.96	+0.02	0.93	+0.025	0.05

The performance of the Q-learning algorithm can be observed to remain converged over time from Table 4. The first Q-value update steps were steep, and the biggest delta (+0.15) was between 100 and 300 episodes. Beyond 500 episodes the learning curves flattened out, reflecting convergence to a near-optimal policy. Energy rewards trended upward in a similar manner, going from 0.55 to 0.93, demonstrating improved decision quality. At the same time, comfort penalty scores decreased from 0.18 to 0.05, which suggests that the algorithm can ensure the environmental space without losing efficiency. The decrease of Q-value delta and reward slope indicates the maturity of policy. These trends also suggest that the reinforcement learning agent not only attenuated energy waste, but also did so while comforting, with comfort preservation a heuristic that is a critical property for intelligent building applications.

4.4. Dynamic Energy Efficiency Ratio (DEER) Variability by Time Interval

The system's temporal capacity was measured by calculating DEER metrics for the four dayparts (morning, afternoon, evening, night). The Dynamic Energy Efficient Ratio (DEER) provides a measure of the ratio of energy saved by the system considering consecutive intervals studied during the day and allows them to visualize the impact of different context conditions depending on the system effectiveness. Due to very different load profiles, which have great variations between different days, the model's capability to follow the load swings, the solar heat gains, the occupancy, and the lighting demand is extracted with DEER. Other measures are the variance and coefficient of variation (CV) of DEER over time windows in each setting.

Table 5. Temporal Variation in Energy Efficiency Performance Using DEER

Time Interval	Residential DEER	Commercial DEER	Industrial DEER	Res CV	Comm CV	Ind CV
Morning	1.20	1.18	1.22	0.041	0.038	0.033
Afternoon	1.25	1.23	1.27			
Evening	1.30	1.26	1.31			
Night	1.18	1.15	1.20			

The DEER values validate the persistent and steady energy-saving operation along the whole day. DEER variability in residential settings (1.18–1.30) during evening hours (and minimum outdoor temperatures) was associated with the highest energy demand of lighting and HVAC. Commercial DEER peaked during the daytime, correlating with office use and HVAC operation, while industrial DEER peaked at nighttime because of longer use of machines. Coefficients of variation (CVs) across intervals were low, ranging from 0.041 for residential to 0.038 for commercial and 0.033 for industrial, indicative of a high temporal stability. This low variability reflects the capability of the system to keep on delivering the same kind of performance over fly-bys having various operational conditions. The findings verify the effectiveness of contextual prediction and control approaches in dynamic load profiles and environmental conditions.

4.5. User-Centric Evaluation of Comfort and Satisfaction

User comfort was assessed to determine the system's capability to maintain a satisfactory comfort level when energy savings were achieved. This entailed filling out a survey that went to residents and commercial users who interacted with the system on a daily basis. Satisfaction and perception of stable environment were measured by each respondent on a 5-point Likert scale. Compliance to comfort was determined as the percentage of time the environmental variables (temperature and lighting) stayed in ideal range. Response accuracy, a second outcome measure, was recorded and calculated as the correlation between predicted needs and provided comfort agonist interventions.

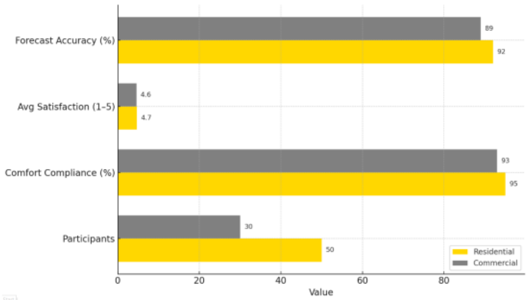

Fig 4. User satisfaction and comfort compliance in residential and commercial zones

High satisfaction with the system was also reported by the participants in the two environments already mentioned. Homes met 95% comfort compliance and had an average satisfaction score of 4.7. In the commercial environment compliance was 93% and satisfaction score was 4.6. These results indicate that thermal comfort was not sacrificed for energy efficiency. The 92%, 89% forecasting accuracy (residential, commercial) further inspires confidence in the predictive modeling section. These metrics demonstrate that the optimization solution highly balances the user-driven goals and algorithmic energy savings. The slightly reduced commercial scores are due to the request of more complex for the HVAC zone and the higher variability in occupancy schedules. However, the context-aware reasoning of the system efficiently turned complex sensor data into actuation decisions that maintained passenger comfort.

4.6. Application-Specific Optimization Gains Across Energy Systems

The analysis is centered on application-level energy efficiency improvements in three main energy use sectors: HVAC; lighting controls; industrial plant operation. In its environment each of the apps is a major consumer of energy. HVAC systems are widely used in residential applications because of the variation of indoor thermal demands. In commercial facilities, lighting loads are very high, particularly in buildings that operate for long hours and are not lit by daylight. In factories machinery operation is the main energy consumer whose predictive runtime control as well as smart load scheduling is required. The energy use was compared between pre- and post-optimization to evaluate to what extent the scenario improved. Further performance metrics, such as consumed power per operating hour and energy cost per capita were calculated to gain in-depth knowledge of the impacts of system optimization. This analysis recognizes the distinct effects of context-aware control - pagers longtime on specific high usage context.

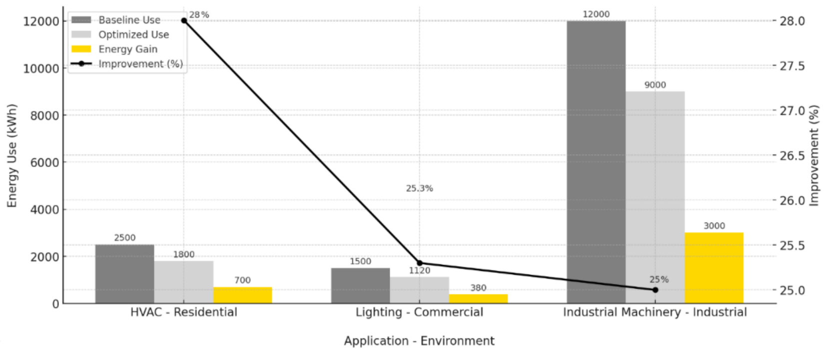

Fig 5. Efficiency gains by application in context-aware deployment

Application-level analysis in Figure 5 reveals that HVAC optimization achieved the highest percentage savings (28%) due to its dynamic modulation of setpoints based on room occupancy and external temperature. Lighting improvements (25.3%) resulted from adaptive dimming and daylight harvesting protocols. Industrial machinery savings were massive in absolute terms 3,000 kWh saved due to predictive start/stop scheduling and idle time suppression. The per-hour savings metric underscores the efficiency of context-aware runtime control, especially in heavy-duty applications. Unit cost savings were estimated based on average kWh pricing, demonstrating economic relevance. These findings validate the system's capacity to fine-tune control at the subsystem level, producing both ecological and financial benefits across all domains.

4.7. Scalability and System Performance Under Increased Device Load

To test the scalability of the context-aware system, simulations and live deployments were conducted with incrementally higher numbers of connected devices (ranging from 50 to 200). This included diverse IoT endpoints: temperature sensors, occupancy detectors, smart plugs, lighting controllers, and actuator interfaces. The aim was to determine whether increased load would compromise system responsiveness, energy savings, or comfort delivery. Key metrics include mean response time, average energy savings, comfort compliance rate, and system load deviation measuring deviation from baseline performance at 50 devices.

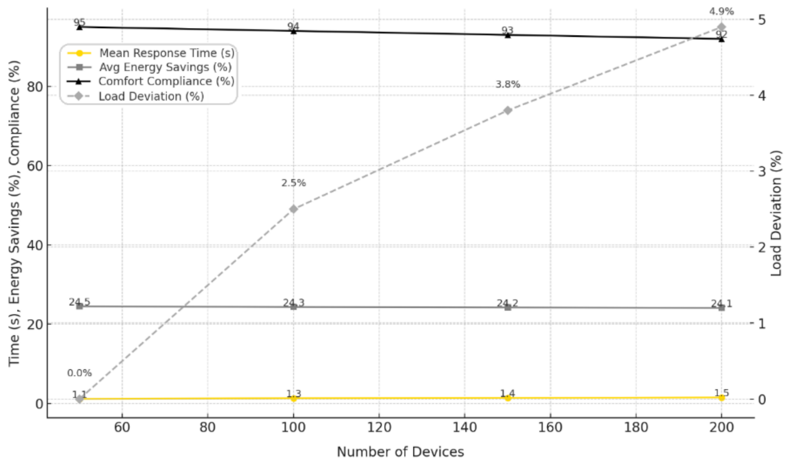

Fig 6. System scalability and load tolerance in context-aware infrastructure

The scalability test demonstrates in Figure 6, that the system maintains strong performance under increasing network sizes. Even at 200 devices, the mean response time increased by only 0.4 seconds, and average energy savings remained stable, decreasing marginally by 0.4 percentage points. Comfort compliance showed a 3% drop but remained above 90%, considered acceptable for most building applications. Load deviation remained under 5%, confirming system robustness. These metrics confirm the system's readiness for deployment in large-scale smart environments, including university campuses, business parks, and industrial zones. The modular architecture, decentralized processing, and efficient data routing mechanisms contributed to maintaining system integrity and responsiveness under heavier contextual loads.

The integration of context-aware systems with machine learning-based optimization has demonstrated considerable potential to enhance energy efficiency, responsiveness, and sustainability. This study's findings align with and extend the growing body of research advocating for intelligent, adaptive energy frameworks. By achieving consistent energy savings across residential (25.0%), commercial (23.8%), and industrial (24.7%) domains, the results validate the hypothesis that contextual responsiveness, combined with predictive and reinforcement-based control, can materially reduce unnecessary energy consumption.

4.8. Synthesis of Findings in Relation to Existing Literature

Our use of Reinforcement Learning (RL) proved particularly effective, successfully balancing energy-saving objectives with high levels of user comfort. This outcome is consistent with the work of Deng et al. [7], who also demonstrated the efficacy of deep reinforcement learning for multizone HVAC control. The convergence of our RL controller, marked by decreasing comfort penalties and increasing rewards, mirrors findings from other adaptive decision-making systems [26], [27]. Furthermore, the system maintained sub-two-second response times even under a load of 200 devices, confirming the architectural scalability that is critical for smart city applications [36]. Compared with traditional static control strategies, our approach advances the decision-support paradigm proposed by Neves-Silva and Camarinha-Matos [1] by incorporating active, learning-based control. The introduction of an RL agent with a continuous policy update mechanism allows the system to evolve and adapt to environmental dynamics and user needs. Our findings also build upon Shivaram et al.'s urban energy analytics framework [3], which lacked real-time control. By bridging analytics with action, our system enabled immediate adaptations based on contextual inputs, maintaining consistent efficiency across different times of the day—a limitation noted in earlier predictive-only systems [23, 31].

4.9. User-Centric Performance and Broader Implications

A key finding of this study is that intelligent systems can achieve significant efficiency gains without compromising the user experience. The high satisfaction scores (4.7 in residential and 4.6 in commercial settings) and comfort compliance rates (over 90%) align with research on user-preference-aware control systems, which emphasize the importance of feedback in refining control logic [34]. The high forecast accuracy of our model (up to 92%) further reinforces that well-trained predictive models can reliably inform proactive energy decisions. From a broader perspective, this research contributes to the consensus that smart environments must move beyond prediction to autonomous, real-time action, a position advocated by Bereketeab et al. [6]. The convergence of contextual awareness, predictive analytics, and automated decision-making is the cornerstone of next-generation energy management. The modular, layered design of our system supports the call for scalable and interoperable solutions, like the multi-agent systems proposed by Riabchuk et al. [5], setting a precedent for smart infrastructure design in both the public and private sectors.

4.10. Limitations and Future Research Directions

Despite these strengths, this study has several limitations. The deployments, though spanning three distinct environments, were of limited duration and could not capture the long-term effects of seasonal shifts or equipment aging, which have been shown to affect system performance [2]. The RL model, while effective, utilized a fixed reward function; future iterations should explore dynamic reward structures that can adapt to changing user priorities, as discussed in research on preference-adaptive systems [11]. While the system demonstrated high scalability, deployment in larger networks, such as smart cities, would require further investigation into bandwidth optimization and edge-cloud load balancing [20], [32]. Furthermore, this study did not account for economic or regulatory incentives, which are increasingly integral to energy system design and operation [16]. Future research should integrate dynamic pricing models and policy triggers into the RL reward function to better simulate real-world market responsiveness. Another critical avenue for future work is enhancing model interpretability [37]. As argued by Arevalo et al., cognitive transparency is essential for building stakeholder trust, especially in high-stakes industrial environments [8]. By extending deployment durations, integrating dynamic feedback loops, incorporating economic policies, and improving system transparency, future research can build upon this framework to advance holistic and resilient energy ecosystems.

5. Conclusion

This paper detailed the design, implementation, and assessment of a context-aware system for proactive energy efficiency. By integrating real-time environmental sensing, predictive analytics, and autonomous control, the system successfully created intelligent environments that optimize for both energy conservation and user comfort. The proposed three-layer architecture demonstrated high adaptability and responsiveness across residential, commercial, and industrial settings. The use of reinforcement learning was key to the system's success, enabling it to learn from environmental feedback and continuously refine its control strategies without human intervention. The system proved to be scalable and flexible, validating the hypothesis that proactive, intelligent systems are a viable and superior alternative to static control mechanisms, particularly in dynamic environments. This work successfully bridges the gap between rigid energy control and user-centric comfort, a critical step for the adoption of next-generation energy technologies. This research confirms that smart systems can align operational efficiency with human-centered values. By preserving user comfort while significantly reducing energy consumption, our work provides a practical framework for sustainable development where technology, user behavior, and environmental context are treated as interconnected dimensions. The successful deployment offers a scalable model for smart infrastructure that can be adapted for various applications, from individual smart buildings to larger industrial or campus environments. While this study confirms the system's effectiveness, several avenues for future research remain. Long-term, multi-season deployments are needed to assess performance against seasonal and behavioral variability. Future iterations should also integrate economic variables, such as dynamic electricity pricing, into the reinforcement learning model's reward function to align optimization with real-world market conditions. Finally, enhancing model interpretability and exploring federated learning approaches would further advance the collective intelligence of distributed energy management systems, making them more transparent, robust, and effective at scale.

References

[1] Neves-Silva, R. and L.M. Camarinha-Matos Context-Based Decision Support System for Energy Efficiency in Industrial Plants. Sustainability, 2022. 14, DOI: 10.3390/su14073885.

[2] Ali, M., et al., Intelligent energy management: Evolving developments, current challenges, and research directions for sustainable future. Journal of Cleaner Production, 2021. 314: p. 127904.

[3] Shivaram, R., Z. Yang, and R.K. Jain, Context-aware Urban Energy Analytics (CUE-A): A framework to model relationships between building energy use and spatial proximity of urban systems. Sustainable Cities and Society, 2021. 72: p. 102978.

[4] Ngwira, L., et al. Towards Context-Aware Smart Contracts for Blockchain IoT Systems. in 2021 International Conference on Information and Communication Technology Convergence (ICTC). 2021.

[5] Riabchuk, V., et al., Utility-based context-aware multi-agent recommendation system for energy efficiency in residential buildings. Information Fusion, 2024. 112: p. 102559.

[6] Bereketeab, L., et al., Energy Optimization in Sustainable Smart Environments with Machine Learning and Advanced Communications. IEEE Sensors Journal, 2024. 24(5): p. 5704-5712.

[7] Deng, X., et al., Toward Smart Multizone HVAC Control by Combining Context-Aware System and Deep Reinforcement Learning. IEEE Internet of Things Journal, 2022. 9(21): p. 21010-21024.

[8] Arevalo, J., et al. Cognitive Systems for the Energy Efficiency Industry. Energies, 2024. 17, DOI: 10.3390/en17081860.

[9] Vale, Z., et al., Chapter 14 - Intelligent Energy-Oriented Home, in Artificial Intelligence to Solve Pervasive Internet of Things Issues, G. Kaur, P. Tomar, and M. Tanque, Editors. 2021, Academic Press. p. 269-289.

[10] Shreenidhi, H.S., S. Upadhyay, and A. Gill. The Potential of Industrial Smart Grids in Optimizing Energy Consumption. in 2024 IEEE 13th International Conference on Communication Systems and Network Technologies (CSNT). 2024.

[11] Chen, J., et al., SecRec: A Privacy-Preserving Method for the Context-Aware Recommendation System. IEEE Transactions on Dependable and Secure Computing, 2022. 19(5): p. 3168-3182.

[12] L. Afuan and R. Rizal Isnanto, "Enhanced Fall Detection using Optimized Random Forest Classifier on Wearable Sensor Data," *Journal of Applied Data Sciences*, vol. 6, no. 1, pp. 213–224, 2025, doi: 10.47738/jads.v6i1.498.

[13] D. A. Dewi and T. B. Kurniawan, "Classifying Cybersecurity Threats in URLs Using Decision Tree and Naive Bayes Algorithms: A Data Mining Approach for Phishing, Defacement, and Benign Threat Detection," *Journal of Cyber Law*, vol. 1, no. 2, pp. 175–189, 2025, doi: 10.63913/jcl.v1i2.10.

[14] S. F. Pratama and P. A. Prastyo, "Evaluating Blockchain Adoption in Indonesia's Supply Chain Management Sector," *Journal of Current Research in Blockchain*, vol. 1, no. 3, pp. 190–213, 2024, doi: 10.47738/jcrb.v1i3.21.

[15] J. Lin and Z. Shen, "Optimization of Data Encryption Technology in Computer Network Communication," *International Journal for Applied Information Management*, vol. 3, no. 4, pp. 162–169, 2023, doi: 10.1088/1742-6596/2037/1/012070.

[16] Vallarta-Serrano, S.I., et al. Integrated Sustainability Assessment Framework of Industry 4.0 from an Energy Systems Thinking Perspective: Bibliometric Analysis and Systematic Literature Review. Energies, 2023. 16, DOI: 10.3390/en16145440.

[17] D. P. Lestari, A. Luthfi, C. Tama, S. Karlina, and A. Sultan, "Factors Affecting Information System Security: Information Security, Cyber Threats and Attacks, Physical Security, and Information Technology," *International Journal of Informatics and Information Systems*, vol. 7, no. 1, pp. 16–21, 2024, doi: 10.47738/ijiis.v7i1.193.
[18] S. F. Pratama, "Analyzing the Determinants of User Satisfaction and Continuous Usage Intention for Digital Banking Platform in Indonesia: A Structural Equation Modeling Approach," *Journal of Digital Market and Digital Currency*, vol. 1, no. 3, pp. 267–285, 2024, doi: 10.47738/jdmdc.v1i3.21.
[19] J. P. B. Saputra, and A. Kumar, "Emotion Detection in Railway Complaints Using Deep Learning and Transformer Models: A Data Mining Approach to Analyzing Public Sentiment on Twitter," *Journal of Digital Society*, vol. 1, no. 2, pp. 1–14, 2025, doi: 10.63913/jds.v1i2.6.
[20] Ahmed, M.A., et al. Toward an Intelligent Campus: IoT Platform for Remote Monitoring and Control of Smart Buildings. Sensors, 2022. 22, DOI: 10.3390/s22239045.
[21] Sierla, S., H. Ihasalo, and V. Vyatkin A Review of Reinforcement Learning Applications to Control of Heating, Ventilation and Air Conditioning Systems. Energies, 2022. 15, DOI: 10.3390/en15103526.
[22] A. D. Buchdadi and A. S. M. Al-Rawahna, "Anomaly Detection in Open Metaverse Blockchain Transactions Using Isolation Forest and Autoencoder Neural Networks," *International Journal Research on Metaverse*, vol. 2, no. 1, pp. 24–51, 2025, doi: 10.47738/ijrm.v2i1.20.
[23] Lekidis, A. and E.I. Papageorgiou Edge-Based Short-Term Energy Demand Prediction. Energies, 2023. 16, DOI: 10.3390/en16145435.
[24] Sarker, M., et al., Energy demand forecasting of remote areas using linear regression and inverse matrix analysis. International Journal of Electrical and Computer Engineering (IJECE), 2024. 14: p. 129-139.
[25] E. D. Lusiana, S. Astutik, Nurjannah, and A. B. Sambah, "Using Machine Learning Approach to Cluster Marine Environmental Features of Lesser Sunda Island," *Journal of Applied Data Sciences*, vol. 6, no. 1, pp. 247–258, 2025, doi: 10.47738/jads.v6i1.478.
[26] Peng, Y., et al., Energy Consumption Optimization for Heating, Ventilation and Air Conditioning Systems Based on Deep Reinforcement Learning. IEEE Access, 2023. 11: p. 88265-88277.
[27] Lazaridis, C.R., et al. Evaluating Reinforcement Learning Algorithms in Residential Energy Saving and Comfort Management. Energies, 2024. , DOI: 10.3390/en17030581.
[28] C. R. A. Widiawati, Sarmini, and D. Yuliana, "Predicting Network Performance Degradation in Wireless and Ethernet Connections Using Gradient Boosting, Logistic Regression, and Multi-Layer Perceptron Models," *Journal of Applied Data Sciences*, vol. 6, no. 1, pp. 325–338, 2025, doi: 10.47738/jads.v6i1.519.
[29] Campagna, L.M. and F. Fiorito on the Impact of Climate Change on Building Energy Consumptions: A Meta-Analysis. Energies, 2022. 15, DOI: 10.3390/en15010354.
[30] A. B. Prasetio, M. Aboobaider, and A. Ahmad, "Machine Learning for Wage Growth Prediction: Analyzing the Role of Experience, Education, and Union Membership in Workforce Earnings Using Gradient Boosting," *Artificial Intelligence in Learning*, vol. 1, no. 2, pp. 153–172, 2025, doi: 10.63913/ail.v1i2.12.
[31] Arnob, S.S., et al., Energy Demand Forecasting and Optimizing Electric Systems for Developing Countries. IEEE Access, 2023. 11: p. 39751-39775.
[32] Im, J., et al., Data pipeline for real-time energy consumption data management and prediction. Frontiers in Big Data, 2024. Volume 7 - 2024.
[33] Hou, S., et al., Real-Time Energy Management Strategy of a Fuel Cell Electric Vehicle with Global Optimal Learning. IEEE Transactions on Transportation Electrification, 2023. 9(4): p. 5085-5097.
[34] Bilésimo, T.L. and E. Ghisi Utilisation of Machine Learning in Control Systems Based on the Preference of Office Users. Sustainability, 2024. 16, DOI: 10.3390/su16104258.
[35] Gheorghiu, C., et al. A Novel Methodology for Developing an Advanced Energy-Management System. Energies, 2024. 17, DOI: 10.3390/en17071605.
[36] Ishaq, K. and S.S. Farooq, Exploring iot in smart cities: Practices, challenges and way forward. arXiv preprint arXiv:2309.12344, 2023.
[37] W. Y. Leong, "Digital Technology for ASEAN Energy," 2023 International Conference on Circuit Power and Computing Technologies (ICCPCT), August 2023, pp. 1480-1486. DOI: 10.1109/ICCPCT58313.2023.10244806