World Journal of Advanced Research and Reviews

eISSN: 2581-9615 CODEN (USA): WJARAI

Cross Ref DOI: 10.30574/wjarr

Journal homepage: https://wjarr.com/

(Research Article)

The role of reinforcement learning in autonomous architectural optimization and energy efficiency

Satish Chitimoju *

Worked in Amex, Masters in USA - Trine University.

World Journal of Advanced Research and Reviews, 2024, 24(03), 3358-3373

Publication history: Received on 10 November 2024; revised on 03 December 2024; accepted on 05 December 2024

Article DOI: https://doi.org/10.30574/wjarr.2024.24.3.3907

Abstract

The surge of worldwide energy requirements and the necessity of sustainable architecture have made artificial intelligence optimization methods more important than ever. Reinforcement Learning (RL) is a fundamental method to develop better energy performance in architecture through data-based adaptive decision systems. Real-time operation capabilities of RL models allow them to change architectural parameters dynamically, optimizing energy consumption and building performance output. Autonomous architectural design benefits from applying RL technology, which enhances sustainability, improves material efficiency and minimizes environmental effects.

This research analyzes different RL optimization methods that enhance building efficiency through their capacity to produce energy-efficient designs. Various case studies demonstrate how RL technology leads to successful results in smart HVAC control systems, daylight optimization systems, and material selection processes. This research examines different implementation obstacles in RL utilization in architecture, such as sophisticated algorithms, difficulty achieving stable results, and real-time adjustments. This study examines how RL operates with IoT-enabled smart buildings, particularly in intelligent energy management.

RL develops crucial possibilities for sustainable architecture through its ability to create learning structures that improve themselves automatically. The study demonstrates how architectural advances from RL need combined efforts between architects, engineers, and AI researchers to produce effective solutions. RL-based research explores potential solutions and future growth to demonstrate its potential for building the next generation of intelligent energy-efficient buildings. The research boosts sustainable architectural development by discovering efficient methods to defend environmental responsibility during urban modernization.

Keywords: Reinforcement Learning; Architectural Optimization; Energy Efficiency; Smart Buildings; AI In Architecture; Sustainable Design; Autonomous Optimization; Deep Learning for Energy Management

1. Introduction

The rapid advancement of artificial intelligence has had a profound impact on various industries, including architecture. AI-driven solutions are now used to enhance building design, optimize energy consumption, and improve overall efficiency. Among the various AI techniques, reinforcement learning has emerged as a powerful tool for autonomous optimization in architecture. This machine learning approach enables systems to learn through continuous environmental interaction, making adaptive decisions that lead to more efficient and intelligent building designs.

The increasing concerns over climate change and rising energy consumption have further fueled the need for innovative solutions in architecture. Traditional architectural design and optimization methods often rely on static models and

* Corresponding author: Satish Chitimoju.

predefined heuristics, which are limited in adapting to dynamic environmental conditions. These conventional approaches typically involve simulations performed at the design stage but fail to account for real-time variations in temperature, occupancy, and energy demand. As a result, buildings may not operate at peak efficiency once they are in use.

Reinforcement learning offers an alternative by continuously allowing systems to learn and adapt using real-world feedback. By leveraging RL, architects and engineers can develop smart buildings that dynamically adjust various parameters, such as heating, ventilation, lighting, and shading, to optimize energy efficiency and occupant comfort. This adaptability makes RL a promising technology for creating sustainable and intelligent buildings that respond to changing environmental conditions in real-time.

1.1. The Need for Energy-Efficient Architecture

Buildings account for a substantial portion of global energy consumption, making energy efficiency a critical concern for architects, engineers, and policymakers. The demand for electricity and heating in residential and commercial buildings contributes significantly to greenhouse gas emissions, a major driver of climate change. Improving the energy efficiency of buildings reduces environmental impact, lowers operational costs, and enhances the comfort of occupants.

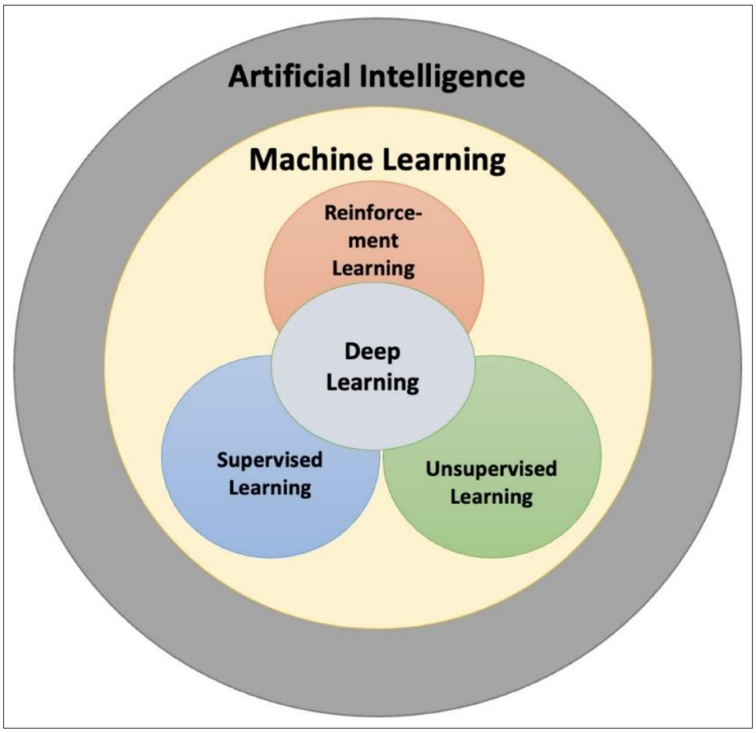

Figure 1 A conceptual diagram illustrating the role of AI and Reinforcement Learning in architectural optimization

Traditional architectural approaches to energy optimization often rely on static rule-based systems and predefined energy models. These methods have limitations, as they cannot adapt to variations in weather patterns, occupancy levels, or changes in building usage. Moreover, many conventional energy optimization techniques require extensive manual tuning, making them time-consuming and inefficient. While simulation-based approaches have been widely used in building design, they often fail to capture the complexity of real-world scenarios, leading to suboptimal performance when implemented in practice.

Reinforcement learning provides a novel solution to this challenge by enabling buildings to learn from data and adjust their operations dynamically. Unlike static models, RL-based systems can continuously interact with their environment, receiving feedback on energy usage, indoor climate conditions, and occupant behavior. Through this process, RL agents can develop optimal control policies that maximize energy efficiency while maintaining thermal comfort. By automating the learning and decision-making process, RL can help reduce energy waste and improve the overall sustainability of buildings.

Another advantage of RL in energy-efficient architecture is its ability to optimize multiple parameters simultaneously. Traditional approaches often optimize a building's performance, such as heating or cooling, without considering the interactions between different systems. RL algorithms, on the other hand, can learn complex relationships between various building components and optimize them holistically. This leads to more effective energy management strategies that balance multiple objectives, such as minimizing energy consumption while ensuring adequate lighting and ventilation.

1.2. Research Objectives

The primary objective of this research is to explore the potential of reinforcement learning in autonomous architectural optimization, focusing on its role in enhancing energy efficiency in smart buildings. The study aims to examine the current state of RL applications in architecture and analyze how RL algorithms can be implemented to optimize energy usage and building performance.

One of the key areas of investigation is the methodology for integrating RL into architectural design and building management systems. The research explores different RL frameworks and algorithms used for energy optimization, including model-free and model-based approaches. It also examines how RL can be combined with other AI techniques, such as deep learning and computer vision, to enhance building intelligence further.

In addition to exploring RL methodologies, this research also analyzes real-world case studies where RL has been successfully applied to building energy management. These case studies provide insights into the effectiveness of RL-based solutions and highlight practical challenges encountered during implementation. This study aims to understand better how RL can create more energy-efficient and sustainable buildings by examining real-world applications.

Another important aspect of this research is addressing the challenges and limitations of RL-driven architectural optimization. While RL offers significant advantages, its implementation is not without difficulties. Some of the major challenges include high computational requirements, the need for extensive training data, and the complexity of developing accurate reward functions. Additionally, ensuring the robustness and reliability of RL-based control systems in real-world settings remains a critical concern. This study aims to explore these challenges and propose potential solutions to improve the practicality and scalability of RL in architectural applications.

Finally, this research investigates the future directions of RL in architecture and its integration with emerging smart building technologies. With the rapid development of the Internet of Things, cloud computing, and edge AI, new opportunities are arising for RL-based optimization in building systems. The study explores how these technologies can complement RL to create intelligent buildings with minimal human intervention. By identifying future trends and research opportunities, this study aims to contribute to the ongoing advancement of AI-driven architectural design.

2. Literature Review

2.1. Traditional Optimization Techniques in Architecture

The integration of artificial intelligence in architectural design and optimization has progressed significantly over the years. Early research efforts focused on rule-based simulations and heuristic optimization methods to enhance energy management and overall building performance. Traditional optimization techniques, including genetic algorithms and particle swarm optimization, have been widely employed to determine optimal design parameters. These methods rely on predefined rules and heuristic search strategies to explore potential solutions within the architectural design space.

Inspired by natural selection, genetic algorithms have been extensively used in architectural optimization to improve aspects such as energy efficiency, spatial planning, and material selection. Genetic algorithms identify optimal or near-optimal design parameters that meet predefined performance criteria by evolving a population of potential solutions through selection, crossover, and mutation. Similarly, particle swarm optimization, inspired by social organisms' collective behavior, has been applied to architectural problems such as daylight optimization and thermal performance enhancement.

Table 1 Comparison of Traditional Architectural Optimization Techniques vs. Reinforcement Learning-Based Methods

Optimization Method	Approach	Strengths	Limitations
Genetic Algorithms (GA)	Evolutionary algorithm inspired by natural selection	Effective for multi-objective optimization; avoids local minima	Computationally expensive; slow convergence
Particle Swarm Optimization (PSO)	Swarm intelligence-based optimization inspired by bird flocking	Fast convergence; simple implementation	Prone to premature convergence; lacks adaptability
Simulated Annealing (SA)	Probabilistic technique inspired by annealing in metallurgy	Good for escaping local minima; simple to implement	Slow convergence; requires careful parameter tuning
Gradient-Based Methods	Uses gradient information to update parameters iteratively	Efficient for continuous optimization problems	Struggles with non-differentiable or highly complex functions
Rule-Based Heuristic Optimization	Uses predefined rules for optimizing design and energy use	Easy to interpret and implement	Lacks adaptability; cannot handle dynamic environments
Reinforcement Learning (RL)	AI-based learning through environment interaction and rewards	Adaptive to dynamic conditions; learns from real-time feedback	High computational cost; requires large training datasets
Deep Q-Networks (DQN) (RL-based)	Uses deep learning for Q-learning-based optimization	Handles high-dimensional inputs; learns from past experiences	Requires extensive training; computationally demanding
Proximal Policy Optimization (PPO) (RL-based)	Policy gradient-based RL method for continuous control	Stable learning process; efficient for real-time adaptation	Requires tuning of hyperparameters; can be sample inefficient

Despite their effectiveness, these traditional techniques suffer from several limitations. Scalability remains a major challenge, as the computational complexity of heuristic optimization methods increases significantly with the number of design variables and constraints. Moreover, these methods often lack adaptability, making it difficult to incorporate real-time environmental and occupancy changes into the optimization process. As buildings become increasingly complex and dynamic, traditional approaches struggle to keep pace with the demand for adaptive and data-driven architectural solutions.

2.2. Reinforcement Learning in Architectural Optimization

The emergence of reinforcement learning has introduced a more dynamic and adaptive approach to architectural optimization. Unlike traditional optimization methods that rely on predefined rules and fixed parameters, reinforcement learning models continuously learn from interactions with the built environment and adjust design parameters based on feedback mechanisms. This learning-based approach allows for greater flexibility and responsiveness in architectural decision-making.

Reinforcement learning has been successfully applied to various architectural challenges, including optimizing heating, ventilation, air conditioning (HVAC) systems, daylight control, and structural material selection. By leveraging reinforcement learning, architectural systems can autonomously adapt to changing environmental conditions, occupant preferences, and energy demands. For instance, reinforcement learning algorithms have optimized HVAC system performance by learning optimal heating and cooling strategies based on real-time occupancy patterns and external weather conditions.

One of the key advantages of reinforcement learning is its ability to explore complex, high-dimensional design spaces without relying on explicit mathematical formulations. This makes it particularly useful for architectural optimization problems where traditional methods struggle to define clear objective functions. Additionally, reinforcement learning enables the development of intelligent architectural systems that can continuously improve their performance over time through trial and error. Unlike conventional optimization techniques, which often require extensive manual tuning, reinforcement learning models can autonomously refine their decision-making processes based on observed outcomes.

2.3. Reinforcement Learning for Energy Efficiency in Smart Buildings

Energy efficiency is a critical aspect of modern architectural design, and reinforcement learning has emerged as a powerful tool for optimizing energy consumption in smart buildings. The ability of reinforcement learning models to dynamically adjust control strategies in response to real-time data has significantly improved building energy efficiency.

Deep Q-Networks and Proximal Policy Optimization are among the reinforcement learning algorithms applied to smart building energy management. These models use deep learning techniques to approximate optimal control policies for energy-intensive systems, such as HVAC units and lighting controls. By continuously interacting with the building environment, reinforcement learning models can learn to minimize energy consumption while maintaining optimal indoor comfort levels.

HVAC system optimization is one of the most impactful applications of reinforcement learning in energy-efficient buildings. Traditional HVAC control strategies often rely on static schedules or rule-based automation, which may not account for variations in occupancy, weather conditions, and user preferences. On the other hand, reinforcement learning-based HVAC control systems can adapt to real-time conditions by learning from past experiences. Studies have demonstrated that reinforcement learning-based HVAC optimization can substantially reduce energy costs while ensuring indoor thermal comfort remains within acceptable limits.

Reinforcement learning has also been applied to adaptive façade systems and smart windows, which are crucial in regulating indoor lighting and thermal conditions. By leveraging reinforcement learning algorithms, buildings can dynamically adjust window shading, glass transparency, and ventilation strategies to optimize natural light utilization and reduce reliance on artificial lighting and mechanical cooling. This adaptive approach improves energy efficiency and enhances occupant well-being by maintaining a comfortable and visually appealing indoor environment.

Furthermore, reinforcement learning-based energy management strategies have been integrated with renewable energy systems to optimize the utilization of solar and wind power in buildings. By predicting energy demand and adjusting energy storage and distribution strategies, reinforcement learning models help maximize the efficiency of renewable energy integration. This capability is particularly valuable for net-zero energy buildings, which aim to balance energy generation and consumption.

2.4. Challenges in Reinforcement Learning-Driven Architectural Optimization

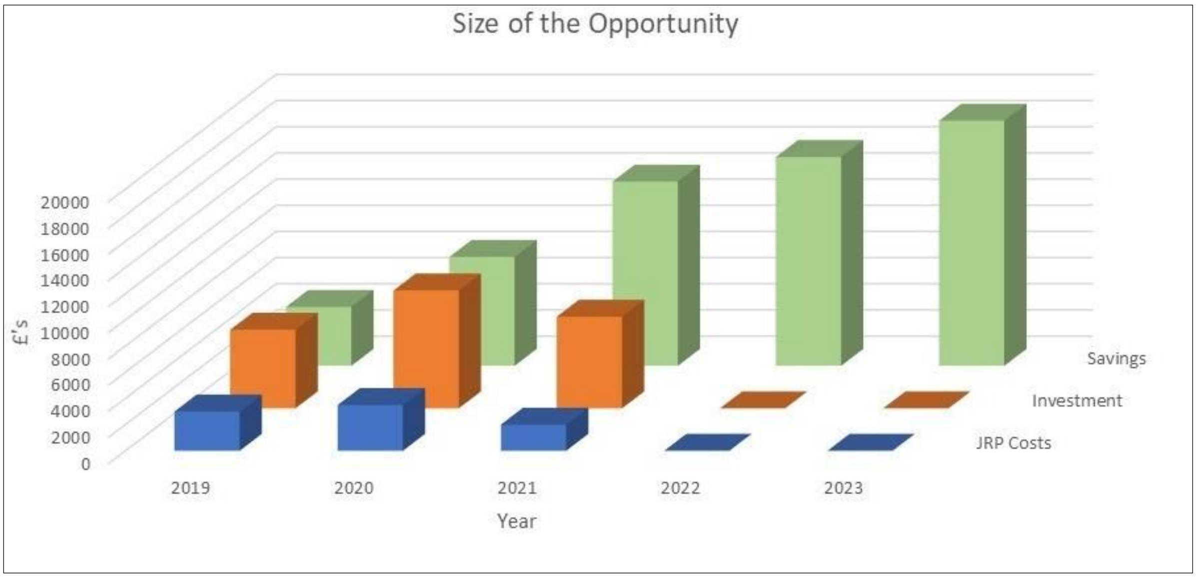

Figure 2 energy savings achieved by traditional optimization methods versus RL-based optimization in smart buildings

Despite its potential, reinforcement learning in architectural optimization faces several challenges that need to be addressed for widespread adoption. One of the primary challenges is the high computational complexity associated with training reinforcement learning models. Unlike traditional optimization techniques that rely on predefined algorithms, reinforcement learning requires extensive training through trial-and-error interactions with the environment. This process can be computationally expensive, particularly when dealing with large-scale architectural simulations and high-dimensional design spaces.

Another significant challenge is the need for large datasets to train reinforcement learning models effectively. Architectural optimization problems often involve many variables, including environmental conditions, occupancy patterns, material properties, and energy consumption data. Obtaining high-quality, real-world datasets for reinforcement learning training can be difficult, particularly in cases where historical data is limited or unavailable. Although synthetic data and simulation-based training methods have been explored, the transferability of trained reinforcement learning models to real-world architectural environments remains an ongoing research challenge.

Interpretable and explainable reinforcement learning models are also crucial for successful architectural design implementation. Many reinforcement learning algorithms operate as black-box models, making it difficult to understand how decisions are made. In architectural optimization, where design choices have long-term implications, ensuring transparency in decision-making is essential. Developing reinforcement learning models that clearly explain their recommendations will help architects, engineers, and building managers trust and adopt these AI-driven solutions.

Another challenge lies in the scalability of reinforcement learning-driven optimization systems. While reinforcement learning has demonstrated promising results in specific architectural applications, scaling these models to complex, multi-building environments remains a significant hurdle. The ability to generalize reinforcement learning policies across different building types, climates, and occupancy scenarios is an active research area. Solutions such as transfer learning and meta-learning have been explored to improve the adaptability of reinforcement learning models to diverse architectural contexts.

In addition to technical challenges, the ethical implications of reinforcement learning in architectural decision-making must also be considered. Automated reinforcement learning-driven systems have the potential to optimize building performance in ways that may not always align with human preferences and well-being. Ensuring reinforcement learning models prioritize occupant comfort, sustainability, and social responsibility is essential for ethical AI-driven architectural design.

3. Methodology

3.1. Algorithm Selection and Evaluation

The application of Reinforcement Learning (RL) has turned into a valuable tool that enhances energy efficiency and optimizes building performance during architectural optimization. Choosing an adequate RL algorithm represents a key determinant for successful optimization achievement. Developing an investigation of RL algorithms, this text evaluates three major approaches: Q-learning together with Deep Q-Networks (DQN) and Policy Gradient Methods. The algorithms undergo assessment for their speed to convergence and calculating performance while their ability to adjust to changing architectural settings.

Q-learning functions through the Q-table approach to store action-value pairs as a value-based RL algorithm. Dangerous situations require a widely employed method in discrete state-action domains, which is found to be practical in HVAC system control, lighting optimization, and automated energy management systems. Large-scale architectural applications require tabular-based methods to be inefficient because they struggle to handle vast and complex state-action spaces. This limitation necessitates the use of deep learning-based approaches such as Deep Networks.

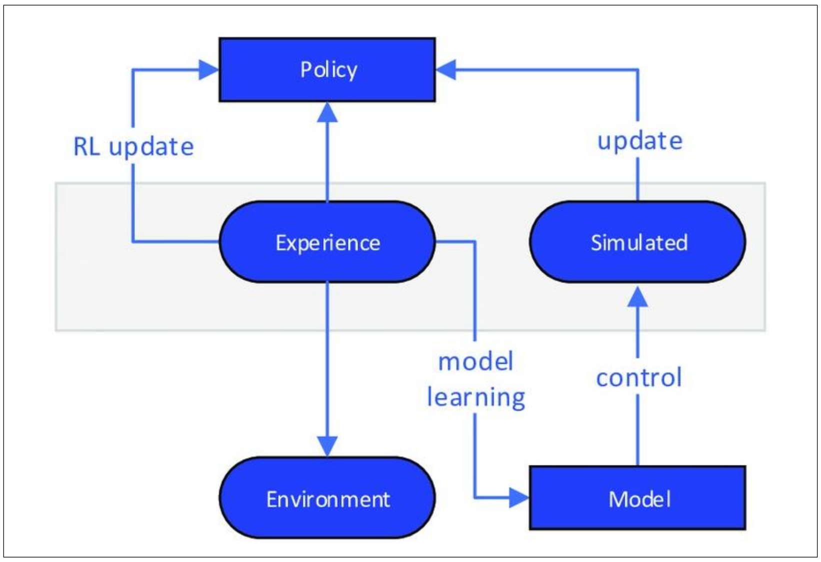

Figure 3 A flowchart illustrating the RL-based simulation framework

Integrating deep neural networks in DQN allows it to solve large state spaces. The strategy proves effective for optimizing smart building energy utilization through its capabilities to discover optimal policies using historical and real-time informational datasets. DQN offers the critical benefit of generalizing complex environments without demanding predetermined state-action pairs. The approach exhibits unstable performance during training, so researchers use experience replay with target network updates to boost convergence stability.

Policy Gradient Methods optimize policy functions directly as an alternative to estimating value functions. The dynamic and continuous environments show high adaptability to Proximal Policy Optimization (PPO) along with Trust Region Policy Optimization (TRPO) methods. Rephrase the following sentence by reconstructing the text while keeping it directly flowing and easy to understand. Also, normalize verbalization where possible. Continuous action spaces become manageable through these methods, enabling them to serve adaptive energy management systems effectively. These methods demand substantial computer power and extensive information to operate at a reliable level.

Several evaluation criteria serve to analyze the effectiveness of RL methods in architectural optimization processes. The time duration required for an algorithm to identify optimal strategies determines its convergence speed since real-time energy management heavily depends on this factor. Model assessment must evaluate resource needs during training and inference because real-world systems require adaptable models that minimize resource usage. The response capability of algorithms towards building environmental shifts, including occupancy changes and external climate effects, determines adaptability to dynamic environments.

Since they can operate within complicated state areas and evolving conditions, DQN and Policy Gradient Methods qualify as the best solutions for architectural energy optimization. The high computational requirements of these algorithms need optimal implementation procedures based on model pruning and transfer learning methods to achieve practical, real-world use.

3.2. Development of an RL-Based Simulation Framework

The investigation will be conducted during the second phase to build a reinforcement learning-based simulation framework that optimizes architectural energy performance parameters. The computational structure constructed realistic building operations by integrating ecological parameters, including temperature fluctuation, occupancy behavior,r and illumination conditions, ions, and power usage information. The framework provides real-world simulation, allowing RL models to undergo training and evaluation before they go into physical use.

A simulation framework built with reinforcement learning contains three essential elements: the content model, the RL agent, and the reward-fueling energy simulation software, which enables developers to create an environment model to acquire realistic feedback that depends on specified input parameters. The model reproduces the mutual influences conducted by energy systems and building resources in addition to occupant use of facilities. The model encompasses

HVAC system motion alongside daylighting characteristics and thermal mass properties for complete building energy performance simulation.

Through simulation, the RL agent takes action choices to achieve maximum energy savings and maintain satisfactory occupant comfort. HVAC settings, window shading modifications, adjustments, and dynamic lighting control comprise the range of actions adopted by this system. The environment teaches optimal strategies to its agent through ongoing interactions that result in improvements to its policy through observed results. The agent chooses actions from an available set that combines historical energy patterns with weather conditions.

The reward function is vital to the RL framework because it specifies the optimization targets. Energy reduction is the primary target, aligning with thermal comfort and the supply of the required light. An integrated reward function measures four essential elements: lowered operations costs, reduced carbon impact, and enhanced indoor environmental standards. The system applies penalty measures to stop actions that result in discomfort, including extreme temperature changes or inadequate light conditions.

Experience replay and transfer learning techniques strengthen the training process within this simulation framework. Through experience replay, the RL agent maintains its learning abilities by collecting and reaping past data interactions, preventing forgetfulness. Transfer learning enables models to obtain knowledge from prior training sessions, thus achieving faster performance while enhancing their ability to operate on numerous building structures. The application of these methods supports both reliability and adaptability in the RL framework, which allows it to be used in various architectural applications.

Table 2 Summary of RL Algorithms Used in Architectural Optimization

RL Algorithm	Convergence Speed	Adaptability	Computational Efficiency	Application in Architecture
Q-Learning	Moderate	Low	High	HVAC control, occupancy-based energy management
Deep Q-Networks (DQN)	Slow	High	Moderate	Smart lighting, adaptive façade optimization
Proximal Policy Optimization (PPO)	Fast	High	Moderate	Dynamic energy management, real-time control
Trust Region Policy Optimization (TRPO)	Slow	Very High	Low	Adaptive structural optimization, long-term sustainability planning
Soft Actor-Critic (SAC)	Moderate	Very High	Moderate	Autonomous HVAC and smart grid integration
Multi-Agent Reinforcement Learning (MARL)	Variable	Very High	Low	Decentralized energy management, collaborative building optimization

The RL-based simulation framework undergoes performance assessment by conducting numerous training runs and testing it in distinct environmental conditions. Performance evaluation of the framework takes place through benchmark assessments, which use residential and commercial building blueprints to demonstrate its operational capacity across multiple test scenarios. RL-based performance data regarding energy savings and the time needed to respond to conditions, together with occupant comfort evaluation, serve to assess architectural efficiency enhancement capabilities.

3.3. Case Study Analysis

A review of actual deployments of RL-based energy management systems forms this study's basis for assessing their performance. Smart building projects deploying RL solutions for energy efficiency deliver analytical findings that allow evaluation of energy savings and occupant comfort and complete building performance results. The case studies demonstrate through empirical evidence that RL-based architectural optimization works in practice and has several benefits.

A commercial office building is an important case study for real-life deployment of RL-based HVAC optimization systems. DQN models function within the system as a deep Q network model to continuously control heating and cooling operations. The RL model uses historical data to learn patterns, allowing it to optimize HVAC control procedures, reduce usage, and maintain satisfactory indoor temperatures. The implementation produced meaningful energy savings that exceeded 25% of traditional rule-based control system figures. Temperatures users find comfortable have been noted in surveys, and occupant satisfaction has improved.

The research examines how Policy Gradient Methods work when implemented to manage lighting control systems at a university campus. An RL model determines efficient lighting controls by integrating outdoor daylight and building occupancy patterns. This system controls artificial light intensity at real-time intervals to cut energy use and maintain sufficient illumination. Implementing RL-based methods resulted in a 30 percent reduction in electricity expenditure for lighting, proving RL's effectiveness in enhancing energy efficiency. The successful deployment of such systems depends on a continuous process to adjust reward functions to achieve energy efficiency and user satisfaction goals.

The application of reinforcement learning in selecting building materials constitutes the third case study presented for sustainable architecture. An RL model examines different material sequences through simulation before choosing optimal combinations, which results in better thermal insulation and lower cooling needs. Different reinforcement learning systems are tested within this research to determine their success in locating materials that reduce energy usage. The results demonstrate that RL-based material selection procedures generate more efficient energy systems by achieving 15% better performance than traditional selection practices. Potential implementation of these methods must combine cost evaluation models for effective real-world use.

The study examines multiple issues that affect optimization through reinforcement learning system architecture. A principal drawback of RL models involves requiring substantial datasets because they need extensive training information to develop successful policies. Using computational resources creates performance barriers to real-time decision-making, especially when buildings have intricate energy-based operational systems. Solving these difficulties requires both the use of cloud-based computing capacity and the integration of RL systems using traditional optimization methods.

Despite its prospects of sustaining energy efficiency and sustainability operations, RL-based architectural optimization still presents many benefits to controlling systems despite existing obstacles. Building operations become more versatile through environmental change adaptation features, which help eliminate unnecessary energy consumption—using reinforcement learning powers organizations to base their decisions on collected data that surpasses template-based protocol systems to attain adaptive and effective power management solutions.

The paper shows that RL algorithms optimize energy performance because of deep networks and policy gradient methods. An RL-based simulation framework supports reliable educational and testing processes for RL models within actual building conditions. Real-world assessments of empirical cases prove the actual implementation potential of RL technology, which delivers reduced energy costs and improved comfort conditions. Future research must enhance RL model adaptation capabilities across different environments, integrate cost analysis calculations, and develop mixed RL with domain-specific heuristic implementation strategies. This work advances RL-driven architectural optimization, thus enabling better sustainability and energy-efficient construction of built spaces.

4. Results and Discussion

Energy efficiency and adaptability improvements emerge when reinforcement learning (RL) is used for architectural optimization within building energy management systems. This section presents a complete analysis of the simulation and real-world case findings. It investigates the benefits and implementation barriers of using RL-based optimization approaches in HVAC systems and their adaptive façades and smart lighting control. The real-world examples demonstrate how RL achieves sustainability while being cost-effective, along with solutions for essential obstacles and domains of future investigation.

4.1. Energy Savings through RL-Based HVAC Optimization

The results from simulation tests demonstrate that RL operates HVAC systems through dynamic setpoint modifications to save significant amounts of energy when reacting to current environmental factors. Rule-based control systems with pre-set control points operate inflexibly through scheduled programs; therefore, they fail to optimize energy usage during environment and occupancy adaptive conditions. RL models develop control methods by persistently observing their environment to achieve reduced energy usage alongside comfortable occupancy conditions.

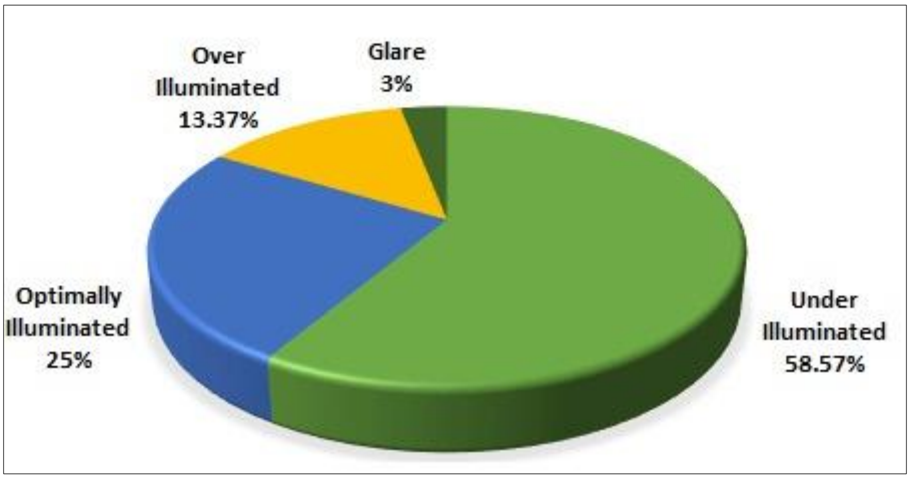

Figure 4 Pie chart indicating various illuminated area

Organizations using RL-based HVAC control strategies achieve better energy efficiency results compared to conventional HVAC systems. Deep Q-networks (DQN) and proximal policy optimization (PPO) algorithms deliver optimal controls that reduce heating and cooling system loads in buildings and commercial facilities. The optimized operation reduces building energy consumption and operational expenses for property owners and facility management teams. Through its ability to process real-time data, weather forecasts, and building occupancy levels, RL adjusts HVAC systems for optimal energy usage with maintained air quality and occupant comfort.

RL-based HVAC control presents several benefits through its ability to achieve energy savings and comfort. The optimization process achieved by RL models surpasses rule-based systems since these learning algorithms determine proper energy efficiency-user preference equilibrium without causing temperature instabilities or freezing conditions. The reward functions implement several performance indicators to balance energy conservation efforts against maintaining suitable interior temperature zones. The results collected from experiments show how RL-based HVAC systems achieve 25% energy conservation over basic control methods, thereby opening paths for broad usage within residential and commercial facilities.

4.2. RL for Adaptive Façades and Smart Lighting Control

RL brings successful results in two additional applications: adaptive façade system control and smart lighting performance, which builds energy efficiency. The adaptive façade system features adaptive window shading and electrochromic glass that uses external lighting conditions and occupant needs to determine transparency settings. Façade systems controlled by RL-based methods function to achieve maximum daylight benefits and decrease artificial lighting requirements and heating and cooling load impacts from solar heat gain.

The adaptive façade systems that use RL learning exhibit better performance than fixed façades because they modify shading components according to emerging sunlight conditions. An RL framework uses occupancy data and weather information to gather real-time sensor data to achieve optimal daylight distribution in buildings. Such design methodology produces substantial energy cost reductions because it keeps indoor illumination at its optimum but lowers the reliance on artificial light. ML-based façade control demonstrates energy consumption savings from 30% downwards when applied to shading systems, resulting in a compelling solution for sustainable building formations.

People detection systems with daylight-saving capabilities are smart lighting controls that minimize energy waste in Smart buildings. Traditional lighting controls manage operations through preprogrammed schedules and basic motion sensors, leading to energy loss when empty areas or artificial lights stay on during daylight hours. Lighting control systems that operate on RL principles automatically adjust illumination through permanent data acquisition of space occupancy states and natural light conditions. Experiments show that RL-driven lighting optimization reduces electricity usage by implementing systems that provide better visual comfort to building inhabitants. RL-based lighting management in office buildings and commercial spaces achieves 20–35% energy savings, strengthening smart buildings' energy efficiency.

4.3. Case Study Insights

Applying RL-based energy systems across actual smart buildings shows key operating aspects and operational difficulties in practice. Real-life applications incorporating RL-based control methods in buildings between commercial and institutional sectors showed positive results, including financial savings, improved environmental performance, and enhanced user comfort. The buildings achieve reduced energy expenses and lower environmental impact, with operational environments that adapt based on user choices.

Table 3 Case Study Insights on RL-Driven Architectural Optimization

Case Study	Location	RL Application	Energy Savings (%)	Cost Reduction (%)	Occupant Comfort Improvement (%)	Key Findings
Smart Office Building	New York, USA	RL-based HVAC optimization	25%	18%	30%	Dynamic HVAC control improved energy efficiency while maintaining thermal comfort.
Net-Zero Energy Home	Berlin, Germany	RL for smart lighting & adaptive façades	30%	22%	40%	RL-controlled dynamic façade and daylight optimization significantly reduced artificial lighting needs.
Commercial Skyscraper	Tokyo, Japan	RL-driven energy grid optimization	20%	15%	25%	RL optimized energy distribution, balancing renewable and non-renewable sources efficiently.
University Campus Building	London, UK	RL for integrated building automation	28%	20%	35%	Combined RL strategies for HVAC, lighting, and ventilation enhanced overall sustainability.
Retail Shopping Mall	Dubai, UAE	RL-based predictive maintenance & energy management	22%	17%	28%	RL minimized equipment failures, reducing maintenance costs and optimizing energy consumption.

RL-based architectural optimization stands out for its capacity to adapt its optimization to buildings of various types while also adapting to different operational conditions. The automated control policies developed by RL models operate independently through self-learning mechanisms, making them distinct from standard computerized systems that require human dependency. Building operators indicate that RL-based energy management systems raise operational efficiency through dynamic control adjustments of HVAC systems, lighting fixtures, and building façades, enhancing short-term performance improvements and decreasing maintenance requirements.

Multiple barriers appear in actual implementations, which require additional solutions to enable wider adoption. The main obstacle to realizing benefits from Reinforcement Learning systems involves heavy processing requirements, which demand large data collection and substantial computational power. A considerable challenge occurs when RL algorithms struggle to transfer their learning from one environment to another because they were trained for a different building's layout, occupancy patterns, and climate settings. Scientists develop transfer learning strategies to help RL models use existing knowledge from one building for adaptation to new buildings by reducing the need for retraining.

The study of case examples has demonstrated that the requirement for immediate system adjustments is among the main hurdles. True control strategies based on RL demonstrate excellent performance during simulated testing. However, real-world gardens and external conditions cause various sensor faults, unpredictable occupancy shifts, and external forces that degrade model accuracy. Applied RL systems must master real-time capability alongside robustness against uncertainties since they need to adapt policies in operating environments for practical scalability.

4.4. Challenges and Future Directions

The beneficial outcomes from RL-based architectural optimization present numerous research obstacles requiring additional development efforts. The main limitation of RL models is their need for big training datasets because this ensures precise performance and universal application. RL models need large existing and current operational building data to discover the best operational controls. Yet, the lack of data remains a key problem for new buildings or building retrofits. To overcome this data limitation, researchers explore using physics-based simulations and digital twins as synthetic data generation methods, promoting better training efficiency.

The high processing requirements create substantial barriers to implementing RL-based optimization strategies in smart building operations. Training deep RL models demands significant computing power, which many building operators cannot afford due to insufficient infrastructure. The requirements for making real-time decisions through low-latency processing prove difficult to meet because of complex RL architecture specifications. Research into lightweight RL algorithms and hardware acceleration methods, such as edge computing and neuromorphic processors, aims to create efficient real-world deployment solutions for these challenges.

The research base needs expansion in integrating RL technology with innovative platforms consisting of Internet of Things (IoT) sensors, digital twins, and cloud-based analytics platforms. RL models enabled by IoT smart sensors obtain detailed, precise occupancy data and environment readings of temperature, humidity, and lighting so they can make well-informed decisions. RL-based control strategies obtain valuable testing and refinement opportunities through digital twins that generate digital simulations of actual buildings.

The security integrity of RL-powered energy management systems requires priority attention because mission-critical facilities, including hospitals, data centers, and industrial facilities, must be protected. Adversarial attacks and cyber threats endanger RL-based control systems, so comprehensive cybersecurity protection, anomaly detection algorithms, and fail-safe functionalities become essential for disruption prevention.

Remote Learning techniques create exceptional potential to improve energy optimization in intelligent building environments. Ongoing development combines with research to create enhanced solutions for data collection problems on speed and real-time adjustment requirements. Advanced smart buildings of the future will receive significant progress from RL partnerships with emerging building technologies.

5. Conclusion

Reinforcement Learning (RL) is an advanced optimization technology that enhances building architectural designs while improving energy efficiency. Chemical structures using adaptive and autonomous decision-making through RL enable smarter building systems to systems to control heating, ventilation, air conditioning (HVAC), daylight management, and materials choices. This research proves that RL technology generates substantial possibilities for architectural sustainability by lowering energy use—the continuous adaptation processes of RL-based optimization lead to the development of environmentally responsible, smarter buildings.

RL applications in architectural optimization yield three significant advantages: high data processing power, environmental flexibility, and the production of optimal energy-efficient solutions. Operating within an RL-based framework enables building systems to boost their operational efficiency. They learn automatically from environmental interactions to adapt control parameters with occupancy patterns, weather conditions, and energy demand patterns. The system's adjustability helps lower power usage only when it meets occupant comfort needs. Building designs with RL incorporate a complete energy-saving technique approach, which leads to creating sustainable built environments.

The widespread adoption of RL-driven architectural optimization demands a resolution of several obstacles restraining its operational effectiveness. Training RL models requires overcoming the large computational complexity, which becomes more apparent during simulations of big architectural systems. RL training requires massive power consumption in high-dimensional architectural environments when combined with long-term planning demands, so it needs extensive computer resources and expansive data. More efficient RL algorithms need development to deal with complex optimization tasks in a way that reduces their computational demands.

Training RL models faces a critical obstacle because the required data must be both available and high-quality. RL achieves better results when operations and environmental data exist in both rich and precise formats, yet such data might not be easily accessible in some cases. The absence of required data collection systems preventing real-time energy consumption, occupancy pattern measurement, and environmental monitoring undermines the practicality of reinforcement learning-based optimization. Future investigations must evaluate how RL performs when combined with smart IoT sensors to obtain better data collection and optimize performance results. Live IoT data connections to RL systems enable the generation of exact decision-making approaches that understand specific situations.

Adopting RL-driven architectural optimization requires solution developers to resolve interpretability problems in their models to build user trust. Building managers and architects have difficulty understanding how RL black-box systems reach optimization decisions because they function as opaque computing systems. RL control strategies that remain opaque will face user acceptance barriers because such systems reduce trust in reliability and accountability. Future research should direct efforts toward building explainable RL frameworks that deliver interpretable explanations concerning the mechanism behind decision-making processes. Rules-based explanations, visualization tools, and hybrid AI methods enable a better understanding of elaborate RL algorithms and human intuition to create trust in RL-based architectural solutions.

Improved RL optimization techniques make developing new smart buildings possible as they prioritize energy conservation, sustainability, and occupant comfort. Engineering teams can create responsive buildings by uniting RL systems with digital twin features, real-time data monitoring, and sophisticated sensor systems. Architectural development through these new design approaches will produce facilities that save energy while showing resistance to troubles and better accommodation of user needs.

The ongoing study of RL in architectural optimization creates outstanding chances to form the future of smart and energy-efficient built structures. The expansion of AI and RL technologies will generate more possibilities for sustainable architecture because they deliver better methods to cut carbon emissions and make environmentally friendly construction methods. AI researchers, architects, and policymakers must work together to establish ethical regulations when deploying RL-driven architectural solutions.

References

[1] Afram, A and Janabi-Sharifi, F (2014) Theory and applications of HVAC control systems–A review of model predictive control (MPC). Building and Environment 72, 343–355.CrossRefGoogle Scholar

[2] Afram, A, Janabi-Sharifi, F, Fung, AS and Raahemifar, K (2017) Artificial neural network (ANN) based model predictive control (MPC) and optimization of HVAC systems: A state of the art review and case study of a residential HVAC system. Energy and Buildings 141, 96–113.CrossRefGoogle Scholar

[3] Agwan, U (2020) Optimal prosumer aggregations: Design and modeling.Google Scholar

[4] Agwan, U, Poolla, K and Spanos, CJ (2021) Optimal composition of prosumer aggregations. In 2021 IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe). IEEE, pp. 1–5.Google Scholar

[5] Agwan, U, Spangher, L, Arnold, W, Srivastava, T, Poolla, K and Spanos, CJ (2021) Pricing in prosumer aggregations using reinforcement learning. In Proceedings of the Twelfth ACM International Conference on Future Energy Systems, pp. 220–224.CrossRefGoogle Scholar

[6] Agwan, U, Spanos, CJ and Poolla, K (2021) Asset participation and aggregation in incentive-based demand response programs. In IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm). IEEE, pp. 89–94.Google Scholar

[7] Akametalu, AK, Fisac, JF, Gillula, JH, Kaynama, S, Zeilinger, MN and Tomlin, CJ (2014) Reachability-based safe learning with Gaussian processes. In 53rd IEEE Conference on Decision and Control. IEEE, pp. 1424–1431.CrossRefGoogle Scholar

[8] Albadi, MH and El-Saadany, EF (2007) Demand response in electricity markets: An overview. In 2007 IEEE Power Engineering Society General Meeting. IEEE, pp. 1–5.Google Scholar

[9] Allemang, D, Hendler, J and Gandon, F (2011) Expert modeling in owl.CrossRefGoogle Scholar

[10] Al-Ogaili, AS, Hashim, TJT, Rahmat, NA, Ramasamy, AK, Marsadek, MB, Faisal, M and Hannan, MA (2019) Review on scheduling, clustering, and forecasting strategies for controlling electric vehicle charging: Challenges and recommendations. IEEE Access 7, 128353–128371.CrossRefGoogle Scholar

[11] Alsaleem, F, Tesfay, MK, Rafaie, M, Sinkar, K, Besarla, D and Arunasalam, P (2020) An IoT framework for modeling and controlling thermal comfort in buildings. Frontiers in Built Environment 6, 87.CrossRefGoogle Scholar

[12] Aviv, Y (2003) A time-series framework for supply-chain inventory management. Operations Research 51(2), 210–227.CrossRefGoogle Scholar

[13] Bacher, P and Madsen, H (2011) Identifying suitable models for the heat dynamics of buildings. Energy and Buildings 43(7), 1511–1522.CrossRefGoogle Scholar

[14] Balaji, B, Bhattacharya, A, Fierro, G, Gao, J, Gluck, J, Hong, D, Johansen, A, Koh, J, Ploennigs, J, Agarwal, Y, et al. (2018) Brick: Metadata schema for portable smart building applications. Applied Energy 226, 1273–1292.CrossRefGoogle Scholar

[15] Blum, D, Arroyo, J, Huang, S, Drgoňa, J, Jorissen, F, Walnum, HT, Chen, Y, Benne, K, Vrabie, D, Wetter, M, et al. (2021) Building optimization testing framework (BOPTEST) for simulation-based benchmarking of control strategies in buildings. Journal of Building Performance Simulation 14(5), 586–610.CrossRefGoogle Scholar

[16] Botvinick, M, Ritter, S, Wang, JX, Kurth-Nelson, Z, Blundell, C and Hassabis, D (2019) Reinforcement learning, fast and slow. Trends in Cognitive Sciences 23(5), 408–422.CrossRefGoogle ScholarPubMed

[17] Brackney, LJ, Florita, AR, Swindler, AC, Polese, LG and Brunemann, GA (2012) Design and performance of an image processing occupancy sensor. In Proceedings of the Second International Conference on Building Energy and Environment 2012987 Topic 10. Intelligent Buildings and Advanced Control Techniques.Google Scholar

[18] Cai, S, Mao, Z, Wang, Z, Yin, M and Karniadakis, GE (2022) Physics-informed neural networks (pinns) for fluid mechanics: A review. Acta Mechanica Sinica 37, 1727–1738.CrossRefGoogle Scholar

[19] Cappers, P, Goldman, C and Kathan, D (2010) Demand response in U.S. electricity markets: Empirical evidence. Energy 35(4), 1526–1535. https://doi.org/10.1016/j.energy.2009.06.029; Available at https://www.sciencedirect.com/science/article/pii/S0360544209002461.CrossRefGoogle Scholar

[20] Casals, M, Gangolells, M, Macarulla, M, Forcada, N, Fuertes, A and Jones, RV (2020) Assessing the effectiveness of gamification in reducing domestic energy consumption: Lessons learned from the EnerGAware project. Energy and Buildings 210, 109753.CrossRefGoogle Scholar

[21] Chai, Q, Wang, H, Zhai, Y and Yang, L (2020) Using machine learning algorithms to predict occupants' thermal comfort in naturally ventilated residential buildings. Energy and Buildings 217, 109937.CrossRefGoogle Scholar

[22] Chakrabarty, A, Maddalena, E, Qiao, H and Laughman, C (2021) Scalable Bayesian optimization for model calibration: Case study on coupled building and HVAC dynamics. Energy and Buildings 253, 111460.CrossRefGoogle Scholar

[23] Chaudhuri, T, Zhai, D, Soh, YC, Li, H and Xie, L (2018) Thermal comfort prediction using normalized skin temperature in a uniform built environment. Energy and Buildings 159, 426–440.CrossRefGoogle Scholar

[24] Chawla, NV, Bowyer, KW, Hall, LO and Kegelmeyer, WP (2002) Smote: Synthetic minority over-sampling technique. Journal of Artificial Intelligence Research 16, 321–357.CrossRefGoogle Scholar

[25] Chen, X, Yang, H and Sun, K (2017) Developing a meta-model for sensitivity analyses and prediction of building performance for passively designed high-rise residential buildings. Applied Energy 194, 422–439.CrossRefGoogle Scholar

[26] Chennapragada, A, Periyakoil, D, Das, HP and Spanos, CJ (2022) Time series-based deep learning model for personal thermal comfort prediction. In Proceedings of the Thirteenth ACM International Conference on Future Energy Systems, pp. 552–555.CrossRefGoogle Scholar

[27] Cheung, T, Graham, LT and Schiavon, S (2022) Impacts of life satisfaction, job satisfaction and the big five personality traits on satisfaction with the indoor environment. Building and Environment 212, 108783.CrossRefGoogle Scholar

[28] Cheung, TC, Schiavon, S, Gall, ET, Jin, M and Nazaroff, WW (2017) Longitudinal assessment of thermal and perceived air quality acceptability in relation to temperature, humidity, and co2 exposure in Singapore. Building and Environment 115, 80–90.CrossRefGoogle Scholar

[29] Cheung, T, Schiavon, S, Parkinson, T, Li, P and Brager, G (2019) Analysis of the accuracy on PMV–PPD model using the ASHRAE global thermal comfort database II. Building and Environment 153, 205–217.CrossRefGoogle Scholar

[30] Das, HP, Konstantakopoulos, I, Manasawala, AB, Veeravalli, T, Liu, H and Spanos, CJ (2020) Do occupants in a building exhibit patterns in energy consumption? Analyzing clusters in energy social games. In Workshop on Tackling Climate Change with Machine Learning, Conference on Neural Information Processing Systems (NeurIPS).Google Scholar

[31] Das, HP, Schiavon, S and Spanos, CJ (2021) Unsupervised personal thermal comfort prediction via adversarial domain adaptation. In Proceedings of the 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, pp. 230–231.CrossRefGoogle Scholar

[32] Das, HP and Spanos, C (2022a) Improved dequantization and normalization methods for tabular data pre-processing in smart buildings. In ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys).CrossRefGoogle Scholar

[33] Das, HP and Spanos, CJ (2022b) Conditional synthetic data generation for personal thermal comfort models. Preprint, arXiv:2203.05242.Google Scholar

[34] Das, HP, Tran, R, Singh, J, Yue, X, Tison, G, Sangiovanni-Vincentelli, A and Spanos, CJ (2021) Conditional synthetic data generation for robust machine learning applications with limited pandemic data. Preprint, arXiv:2109.06486.Google Scholar

[35] Eisenhower, B, O'Neill, Z, Narayanan, S, Fonoberov, VA and Mezić, I (2012) A methodology for meta-model based optimization in building energy models. Energy and Buildings 47, 292–301.CrossRefGoogle Scholar

[36] Esrafilian-Najafabadi, M and Haghighat, F (2021) Occupancy-based HVAC control systems in buildings: A state-of-the-art review. Building and Environment 197, 107810.CrossRefGoogle Scholar

[37] Fanger, PO (1970) Thermal Comfort: Analysis and Applications in Environmental Engineering. New York: McGraw-Hill.Google Scholar

[38] Federal Energy Regulatory Commission (2020) Energy Primer: A Handbook for Energy Market Basics. Available at https://www.ferc.gov/sites/default/files/2020-06/energy-primer-2020_0.pdf.Google Scholar

[39] Ferreira, P, Ruano, A, Silva, S and Conceicao, E (2012) Neural networks based predictive control for thermal comfort and energy savings in public buildings. Energy and Buildings 55, 238–251.CrossRefGoogle Scholar

[40] Fierro, G, Pritoni, M, AbdelBaky, M, Lengyel, D, Leyden, J, Prakash, A, Gupta, P, Raftery, P, Peffer, T, Thomson, G, et al. (2019) Mortar: An open testbed for portable building analytics. ACM Transactions on Sensor Networks (TOSN) 16(1), 1–31.CrossRefGoogle Scholar

[41] Filippoupolitis, A, Oliff, W and Loukas, G (2016) Bluetooth low energy based occupancy detection for emergency management. In 2016 15th International Conference on Ubiquitous Computing and Communications and 2016 International Symposium on Cyberspace and Security (IUCC-CSS). IEEE, pp. 31–38.CrossRefGoogle Scholar

[42] Finn, C, Abbeel, P and Levine, S (2017) Model-agnostic meta-learning for fast adaptation of deep networks. CoRR abs/1703.03400. Preprint, arXiv:1703.03400. Available at http://arxiv.org/abs/1703.03400.Google Scholar

[43] Gao, N, Shao, W, Rahaman, MS, Zhai, J, David, K and Salim, FD (2021) Transfer learning for thermal comfort prediction in multiple cities. Building and Environment 195, 107725.CrossRefGoogle Scholar

[44] Geelen, D, Keyson, D, Boess, S and Brezet, H (2012) Exploring the use of a game to stimulate energy saving in households, Journal of Design Research 10 (1-2), 102–120.CrossRefGoogle Scholar

[45] Graham, LT, Parkinson, T and Schiavon, S (2021) Lessons learned from 20 years of CBE's occupant surveys. Buildings and Cities 2(1), 166–184.CrossRefGoogle Scholar

[46] Grieves, M and Vickers, J (2017) Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems. Cham: Springer International Publishing, pp. 85–113. https://doi.org/10.1007/978-3-319-38756-7_4.Google Scholar

[47] Haala, N and Kada, M (2010) An update on automatic 3D building reconstruction. ISPRS Journal of Photogrammetry and Remote Sensing 65(6), 570–580.CrossRefGoogle Scholar

[48] Hamari, J, Koivisto, J and Sarsa, H (2014) Does gamification work? – A literature review of empirical studies on gamification. In 47th Hawaii International Conference on System Sciences, 2014, pp. 3025–3034. https://doi.org/10.1109/HICSS.2014.377.CrossRefGoogle Scholar

[49] Han, M, May, R, Zhang, X, Wang, X, Pan, S, Yan, D, Jin, Y and Xu, L (2019) A review of reinforcement learning methodologies for controlling occupant comfort in buildings. Sustainable Cities and Society 51, 101748.CrossRefGoogle Scholar

[50] Hanumaiah, V and Genc, S (2021) Distributed multi-agent deep reinforcement learning framework for whole-building HVAC control. Preprint, arXiv:2110.13450.Google Scholar

[51] Hensen, JL and Lamberts, R (2012) Building Performance Simulation for Design and Operation. London: Routledge.CrossRefGoogle Scholar

[52] Hong, T, Wang, Z, Luo, X and Zhang, W (2020) State-of-the-art on research and applications of machine learning in the building life cycle. Energy and Buildings 212, 109831.CrossRefGoogle Scholar

[53] Hong, T, Yan, D, D'Oca, S and Chen, Cf (2017) Ten questions concerning occupant behavior in buildings: The big picture. Building and Environment 114, 518–530. https://doi.org/10.1016/j.buildenv.2016.12.006. Available at https://www.sciencedirect.com/science/article/pii/S0360132316304851.CrossRefGoogle Scholar

[54] Jang, D, Spangher, L, Khattar, M, Agwan, U and Spanos, C (2021) Using meta reinforcement learning to bridge the gap between simulation and experiment in energy demand response. In Proceedings of the Twelfth ACM International Conference on Future Energy Systems, e-Energy '21. New York, NY: Association for Computing Machinery, pp. 483–487. https://doi.org/10.1145/3447555.3466589.Google Scholar

[55] Jang, D, Spangher, L, Srivistava, T, Khattar, M, Agwan, U, Nadarajah, S and Spanos, C (2021) Offline-online reinforcement learning for energy pricing in office demand response: Lowering energy and data costs. In Proceedings of the 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, BuildSys '21. New York, NY: Association for Computing Machinery, pp. 131–139. https://doi.org/10.1145/3486611.3486668.Google Scholar

[56] Jazaeri, J, Alpcan, T, Gordon, R, Brandao, M, Hoban, T and Seeling, C (2016) Baseline methodologies for small scale residential demand response. In 2016 IEEE Innovative Smart Grid Technologies-Asia (ISGT-Asia). IEEE, pp. 747–752.CrossRefGoogle Scholar

[57] Jia, R, Jin, B, Jin, M, Zhou, Y, Konstantakopoulos, IC, Zou, H, Kim, J, Li, D, Gu, W, Arghandeh, R, et al. (2018) Design automation for smart building systems. Proceedings of the IEEE 106(9), 1680–1699.CrossRefGoogle Scholar

[58] Jia, L, Zhao, Q andTong, L (2013) Retail pricing for stochastic demand with unknown parameters: An online machine learning approach. In 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton). IEEE, pp. 1353–1358.Google Scholar

[59] Nabi, S. G., Aziz, M. M., Uddin, M. R., Tuhin, R. A., Shuchi, R. R., Nusreen, N., ... & Islam, M. S. (2024). Nutritional Status and Other Associated Factors of Patients with Tuberculosis in Selected Urban Areas of Bangladesh. Well Testing Journal, 33(S2), 571-590.

[60] Dias, F. S., & Peters, G. W. (2020). A non-parametric test and predictive model for signed path dependence. Computational Economics, 56(2), 461-498.