buildings

MDPI

Review

A Systematic Review of Building Energy Management and Optimization Using the Artificial Intelligence of Things (AIoT)

Yunzhi Tian [1,*], Yuan Tian [2], Yi Jiang [1] and Vedran Mrzljak [3,*]

[1] School of Architectural Engineering, Huanghuai University, Zhumadian 463000, China; 18137402211@163.com
[2] Xi'an ChangDa Asset Management Co., Ltd., Chang'an University, Xi'an 710064, China; tianyuan@chd.edu.cn
[3] Faculty of Engineering, University of Rijeka, Vukovarska 58, 51000 Rijeka, Croatia
* Correspondence: tianyz1983@163.com (Y.T.); vedran.mrzljak@riteh.uniri.hr (V.M.)

Abstract

The transition toward a net-zero economy requires buildings to evolve from passive consumers into Grid-Interactive Efficient Buildings (GEBs). Traditional Building Energy Management Systems (BEMSs) lack the dynamic intelligence needed to control stochastic energy flows and solve multi-objective optimization problems. To systematically map this technological shift, this study conducts a Systematic Literature Review (SLR) following PRISMA guidelines, analyzing a curated corpus of 144 studies (135 primary technical papers and 9 review articles). Due to the significant diversity in methodological approaches within cyber-physical testbeds and IoT architectures discovered through the literature review process, a qualitative narrative and architectural synthesis was conducted rather than a quantitative meta-analysis. Based on this framework, this review examines emerging paradigms for Cognitive Buildings based on Artificial Intelligence of Things (AIoT), edge computing, and semantic interoperability. This review discusses the evolution of algorithms from predictive Deep Learning (DL) and Deep Reinforcement Learning (DRL) to newer approaches such as Agentic AI and Physics-Informed Neural Networks (PINNs). These new methods address the fundamental "sim-to-real" gap while ensuring thermodynamic consistency and safety in physical actuation. It also presents strategic applications in multi-objective optimization of HVAC systems, demand response, energy arbitrage, and predictive maintenance. Moreover, this review tackles major real-world deployment issues by introducing Federated Learning for data privacy, Transfer Learning for portfolio scaling, and TinyML for overcoming the computational carbon paradox of "Green AI." By quantifying this paradox, the review contrasts the massive computational carbon footprint of cloud-based model training against the milliwatt-class efficiency of localized edge deployments. Overall, this review outlines potential research directions toward the development of autonomous Cognitive Digital Twins (CDTs) and Human-Centric Personal Comfort Models (PCMs).

Keywords: Artificial Intelligence of Things (AIoT); Building Energy Management Systems (BEMS); grid-interactive efficient buildings (GEBs); deep reinforcement learning (DRL); Physics-Informed Machine Learning (PIML); cognitive digital twins (CDTs)

Check for updates

Received: 21 May 2026
Revised: 20 June 2026
Accepted: 23 June 2026
Published: 27 June 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

The paradigm shift to a net-zero economy has radically reshaped the future for the building sector. Buildings are no longer just passive electricity consumers; they are active, strategic participants in the intelligent grid system. Currently, as of 2019, the building

sector is responsible for releasing 38% of global CO_2 emissions, with 28% released from buildings and 10% released due to construction materials [1], making the sector the largest single solution to mitigate climate change. According to forecasts, the peak of global building sector emissions will be between 2030 and 2032, with the possibility of reducing these by as much as 76% by 2060 through the application of new technologies and policy interventions [2,3]. However, the imperativeness of this mandate is not just an issue related to efficiency. With renewable energy resources, solar as well as wind power stochastically add volatility to the power system; buildings need to change from just being consumers to being "prosumers," meaning they should have the ability to both consume and create, store, and trade [4]. This paradigm shift necessitates the creation of Grid-Interactive Efficient Buildings (GEBs) that are capable of providing flexibility services such as Demand Response (DR) and peak shaving [5]. Traditional Building Energy Management Systems (BEMSs) often rely on rigid, rule-based control strategies that are ill-equipped to handle the dynamic, non-linear complexities introduced by bi-directional energy flows and diverse DERs. Consequently, there is a critical need to transition toward advanced, data-driven solutions [6]. Artificial Intelligence of Things (AIoT) technologies bridge this gap by synergizing the real-time sensing and communication capabilities of IoT with the predictive and adaptive control power of AI, thereby enabling proactive energy management without compromising occupant comfort.

The technological evolution of the built environment can be broadly described as a progression through three stages, with the first stage characterized by a deterministic and rule-based Building Energy Management System (BEMS) [6], the second stage marked by the increasing capabilities for sensing, communication, and inter-operability facilitated by the Internet of Things (IoT) [7], and the third stage marked by the increasing integration of artificial intelligence (AI) with IoT, also known as the Artificial Intelligence of Things (AIoT) [8]. Traditionally, BEMSs have been used as a type of Supervisory Control and Data Acquisition (SCADA) technology. These systems have been based on a combination of Rule-Based Control (RBC) and Proportional–Integral–Derivative (PID) control loops and have operated on a pre-programmed or scheduled basis (e.g., fixed temperature setpoints based on the time of day) [9,10]. Although these traditional BEMSs have been successful in their role as automation and control technologies, they are by their very nature reactive and deterministic and are unable to cope with stochastic variables such as variable occupancy levels, changing weather patterns, or time-of-use prices. As a result, these types of BEMSs are often found to be wasteful and are described as "heating empty rooms" [11]. The second phase of evolution, which involves the integration of the IoT, transformed buildings from static structures into data-rich environments. The proliferation of Wireless Sensor Networks (WSNs) and low-power communication protocols, such as Zigbee [12], Long Range Wide Area Network (LoRaWAN) [12], and Message Queuing Telemetry Transport (MQTT) [13], has allowed for the fine-grained monitoring of environmental factors, from CO2 concentration at the zone level to density at the occupancy level. This era solved the problem of "visibility"; however, it has now given birth to the "Data Rich, Information Poor" (DRIP) paradox [14], where IoT infrastructures produce terabytes of high-frequency data, yet traditional control strategies fail to effectively utilize this high-dimensional data to produce actionable optimization strategies. The current state-of-the-art represents the integration of these two concepts: connectivity and intelligence. Accordingly, the concept is referred to as the AIoT [8]. In the context of the AIoT concept, the role of AI is not only limited to the analytical capabilities but is integrated into the operation of the "brain" of the building. To put it differently, the building is transformed into a Cyber-Physical System (CPS). Unlike the previous generations of building automation systems, the AIoT concept is based on a shift from "passive response" to "proactive prediction" through the application

of Deep Learning (DL) for load forecasting and Deep Reinforcement Learning (DRL) for adaptive control. According to the literature, the shift from BEMSs to AIoT results in a potential saving of 20–40% [15,16].

The integration of AI and the IoT in the built environment has led to a significant growth in the secondary literature. A thorough overview of the subject domain revealed the presence of more than 50 review articles published from 2020 to 2025, reflecting the high level of research interest in the subject domain, also known as AIoT. A critical synthesis of the existing knowledge base, as presented in Table 1, revealed the presence of several silos.

Table 1. Comparative analysis of recent review articles (2020–2025) on AIoT in buildings, highlighting the gaps related to agentic control, physical deployment, and the limitations of "Green AI" addressed in this work.

Reference/Year	Application Scenario	Algorithm Type	Implementation Strategy	Performance Focus/Key Findings	Identified Limitations
[7]/2024	Commercial & Residential: HVAC/Lighting automation, Predictive Maintenance, Security.	ML (Automation), Predictive Maintenance	4-Layer Architecture (Sensing, Network, Processing, App)	Highlights automation gains but notes high initial capital costs consume up to 15% of budget.	Integration Complexity: Interoperability issues between devices; prohibitive initial costs.
[17]/2021	Smart Buildings & Microgrids: Energy cost reduction, Thermal comfort, Carbon emissions.	Deep Reinforcement Learning (DQN, DDPG, PPO, MADDPG)	Simulation (EnergyPlus, MATLAB)	Evaluated DRL's superior multi-objective performance over traditional rule-based methods.	Sim-to-Real Gap: Safety and robustness concerns in real-world deployment; poor data efficiency.
[18]/2024	Res/Com/Educational: Load forecasting, HVAC control, Occupancy detection.	ML (SVM, DT), DL (LSTM, CNN), RL	Digital Technologies (SRI), IoT Monitoring	Showed high theoretical accuracy in forecasting, but poor translation to field performance.	Deployment Gap: Only 11% of works had real-world validation; "Black box" interpretability issues.
[19]/2025	Sustainable Building Design: Operational energy, Daylight, Embodied carbon.	Surrogate Models (CNN, GAN, LSTM), Meta-heuristics	Physics-based Simulation	Accelerated design-stage optimization and reduced computational load for embodied carbon.	Lifecycle Gap: Focuses strictly on the design phase rather than real-time operational control.
[20]/2021	Res/Com/Educational: Thermal Comfort (PMV), Energy Efficiency, IAQ.	ANN, Fuzzy Logic, Neuro-Fuzzy, Genetic Algorithms	WSN (ZigBee, Wi-Fi), Mixed testbeds	Benchmarked PMV-based control strategies balancing comfort against energy savings.	Data Scarcity: Lack of high-quality real-world datasets; complexity of human behavior modeling.
[21]/2021	Urban Smart Buildings: Demand response, Grid stability, User comfort.	Fuzzy Logic, Multi-Agent Systems (MAS), Neural Networks	CPS, Cloud Computing (Theoretical)	Demonstrated MAS efficacy for large-scale grid stability and aggregated demand response.	Privacy/Security: Severe data privacy concerns; high cost of retrofitting existing infrastructure.

Table 1. *Cont.*

Reference/Year	Application Scenario	Algorithm Type	Implementation Strategy	Performance Focus/Key Findings	Identified Limitations
[22]/2024	Smart Buildings: Energy optimization, Environmental impact reduction.	ML, IoT Analytics	Sensor Networks, Automation Systems	Summarized energy optimization potentials via sensor-driven IoT analytics.	User Acceptance: Data privacy trust issues; complexity of integrating legacy systems.
[23]/2023	Grid-Interactive Buildings: Generation/Load forecasting, Energy trading.	Hybrid Models (Neuro-Fuzzy), Big Data Mining, Blockchain	Edge, Fog, & Cloud Computing	Showed hybrid models yield higher forecasting accuracy for complex energy trading.	Computational Cost: High resource demand for training hybrid models; cybersecurity risks.
[24]/2024	Smart Buildings: HVAC optimization, Demand response.	Predictive Analytics, Adaptive Control	Sensing & Data Acquisition Layer	Evaluated predictive analytics for adaptive HVAC energy reductions and peak shifting.	Skill Gap: Shortage of skilled personnel to operate and maintain sophisticated AI systems.
Proposed Work	Cognitive Buildings (Prosumers): Multi-objective (Energy + Comfort + Grid + Green AI ROI).	Hybrid AI (Agentic LLMs + PINNs + TinyML)	Semantic IoT (Brick/Haystack) + Edge/TinyML	Targets holistic optimization, bridging the cognitive gap with measurable carbon payback ROI.	Addresses Gaps: Directly tackles the "Cognitive Gap," Semantic Interoperability, and the "Green AI" Paradox.

First of all, the field is dominated by algorithm-centric reviews. For example, seminal works in [17] and recent surveys by Ali et al. [18] provide detailed reviews on Machine Learning (ML) and Deep Reinforcement Learning (DRL). Although these works are successful in benchmarking the algorithms' performance, with energy savings ranging from 20 to 30% in HVAC control applications, they tend to focus on the predictive and adaptive capabilities of the algorithms within simulation environments. The "Sim-to-Real" gap is often ignored, which is a major failure mode for "black box" approaches when they are faced with the stochastic noise and safety constraints of physical buildings. Second, reviews of IoT and infrastructure, such as the one by Poyyamozhi et al. [7], offer a thorough examination of the sensory nervous system in smart buildings. These reviews are excellent in their treatment of the four-layer IoT stack (Sensing, Network, Processing, Application), as well as their treatment of protocols such as MQTT and LoRa. However, they often treat the "intelligence" layer of IoT systems superficially, considering AI as a tool only for data-logging or automation, as opposed to an independent reasoning entity capable of semantic interoperability. Third, Design and Lifecycle Reviews are newly introduced, as shown in [19]. These studies are focused on Surrogate Models and Building Information Modeling (BIM) in optimizing sustainability during the design and construction stages. Although these are important in minimizing carbon footprint, they are often not integrated with real-time control, thereby creating a disconnect between the design intent and the actual operation under AI control.

In spite of the plethora of literature available on the subject, a significant gap has been identified. As illustrated in Table 1, most surveys do not address the conceptual 'Cognitive Gap' required to progress from Predictive AI (forecasting load) to what is theoretically proposed as Agentic AI (reasoning about trade-offs and coordinating systems). It must be noted that while Agentic AI offers a promising framework, empirical evidence of its deployment in physical building environments remains highly limited. In addition, most surveys have yet to examine the "Green AI Paradox" that compares the computational energy costs of training these models with the building energy savings that are achieved. Despite the common claim that "data scarcity" is an issue in the field, the solution that is offered by using Physics-Informed Machine Learning (PIML) to satisfy the laws of thermodynamics in data-scarce environments is rarely discussed in surveys of the broader AIoT literature.

In this paper, the above-mentioned fragmentations will be addressed by proposing a comprehensive "Cognitive Building" approach, which, unlike previous works, will achieve the following:

- Unifying the physical constraints of IoT devices (Edge/TinyML) with the reasoning capabilities of Agentic AI, moving the field beyond the boundaries of mere automation towards autonomous orchestration.
- A clear discussion of Physics-Informed methods (PINNs), which satisfy the laws of thermodynamics, thus directly addressing the safety concerns that have been holding back the application of real-world reinforcement learning (RL).
- A critical discussion of the Energy Return on Investment (EROI). of AI models themselves, promoting the use of efficient TinyML-based solutions instead of energy-inefficient cloud computing.

In order to address systematically the so-called "Cognitive Gap" and the deployment barriers identified in the literature, this review aims to synthesize the most recent advances in the intersection of Artificial Intelligence, Internet of Things (IoT), and building energy management. In particular, by defining the evolutionary path from automated and rule-based systems to autonomous and cognitive cyber-physical systems, this paper aims to address the following core research questions (RQs):

- RQ1 (Infrastructure and Semantics): How can we evolve from a built environment based on heterogeneous high-frequency IoT telemetry into an actionable and semantic-driven architecture that exploits edge computing and standardized ontologies?
- RQ2 (Algorithmic Intelligence and Safety): What are some of the novel algorithmic paradigms, such as Agentic AI and Physics-Informed Machine Learning (PIML), that effectively bridge the "sim-to-real" divide while ensuring thermodynamically compliant actuation?
- RQ3 (Cyber-Physical Applications): How do these cognitive frameworks address complex and multi-objective trade-offs, such as HVAC Pareto Optimization, Fault Diagnosis, and Grid-Interactive Efficient Building (GEB) Energy Arbitrage?
- RQ4 (Deployment Viability): What are some of the important commercialization-related issues related to data privacy, portfolio-level scalability, and computational carbon footprint or "Green AI" paradoxes, and which architectures can alleviate these concerns?
- RQ5 (Future Trajectory): What are some of the important components of the ultimate evolutionary roadmap for smart buildings as they evolve toward real-time Cognitive Digital Twins and Human-Centric Personal Comfort Models?

To address the gaps, this review synthesizes the primary literature through several targeted thematic lenses, including AI and machine learning paradigms, IoT sensing and semantic interoperability architectures, cyber-physical energy management applications (e.g., HVAC optimization and grid-interactive control), and the emerging concept of cognitive digital twins. The detailed characteristics and thematic classification of all included studies are provided in Supplementary Table S2, which organizes the literature according to these dimensions and supports the synthesis presented in the subsequent sections.

The paper is divided into sections with the intent of creating a cohesive flow from the research methodologies to the latest advancements in cyber-physical applications and potential deployment paradigms. After the introduction, Section 2 discusses the systematic literature review methodology adopted from the PRISMA protocol. Section 3 outlines the physical and logical foundations of cognitive buildings by discussing the latest advancements in multimodal IoT sensing technologies, interoperability schemas, and edge computing paradigms. Section 4 outlines the foundations of the algorithmic intelligence core by discussing the shift from predictive deep learning and reinforcement learning paradigms toward Agentic AI and Physics-Informed Machine Learning. Section 5 outlines the strategic cyber-physical applications by discussing the latest advancements in multi-objective HVAC optimization, grid-interactive energy arbitrage, and predictive fault diagnostics. Section 6 outlines the principal implementation challenges and proposes solutions by discussing Federated Learning for data privacy, Transfer Learning for portfolio scalability, and model compression for the "Green AI" computational paradox. Section 7 outlines the potential future roadmap toward the realization of autonomous cognitive digital twins and human-centric personal comfort models, followed by a synthesis of the paradigm shift in Section 8.

To ensure conceptual coherence and maintain scholarly rigor, this review systematically distinguishes between three tiers of technological maturity in smart building operations. First, it analyzes methods and approaches that have been empirically validated in building applications, such as traditional Deep Reinforcement Learning (DRL) for HVAC control and established Fault Detection and Diagnostics (FDD). Second, it evaluates promising concepts borrowed from adjacent fields, such as Large Language Models (LLMs) and Agentic AI, acknowledging that while these show theoretical potential, they lack extensive demonstration in real building contexts. Finally, the manuscript synthesizes these findings to propose forward-looking perspectives and concrete research directions, highlighting the critical pathways required to bridge the gap between simulation and real-world deployment.

2. Research Methodology

With a view to ensuring a rigorous and unbiased synthesis of the state of the art of knowledge in this domain, a systematic literature review (SLR) methodology was used following PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) guidelines (Figure 1). This methodology was designed to be dual-pronged: (1) a scientometric study of existing review literature to identify knowledge gaps (as explained in Section 1), and (2) a critical technical study of primary literature to propose the "Cognitive Building" framework.

This systematic review was not prospectively registered, and no formal protocol was prepared prior to conducting the review. No amendments were therefore applicable.

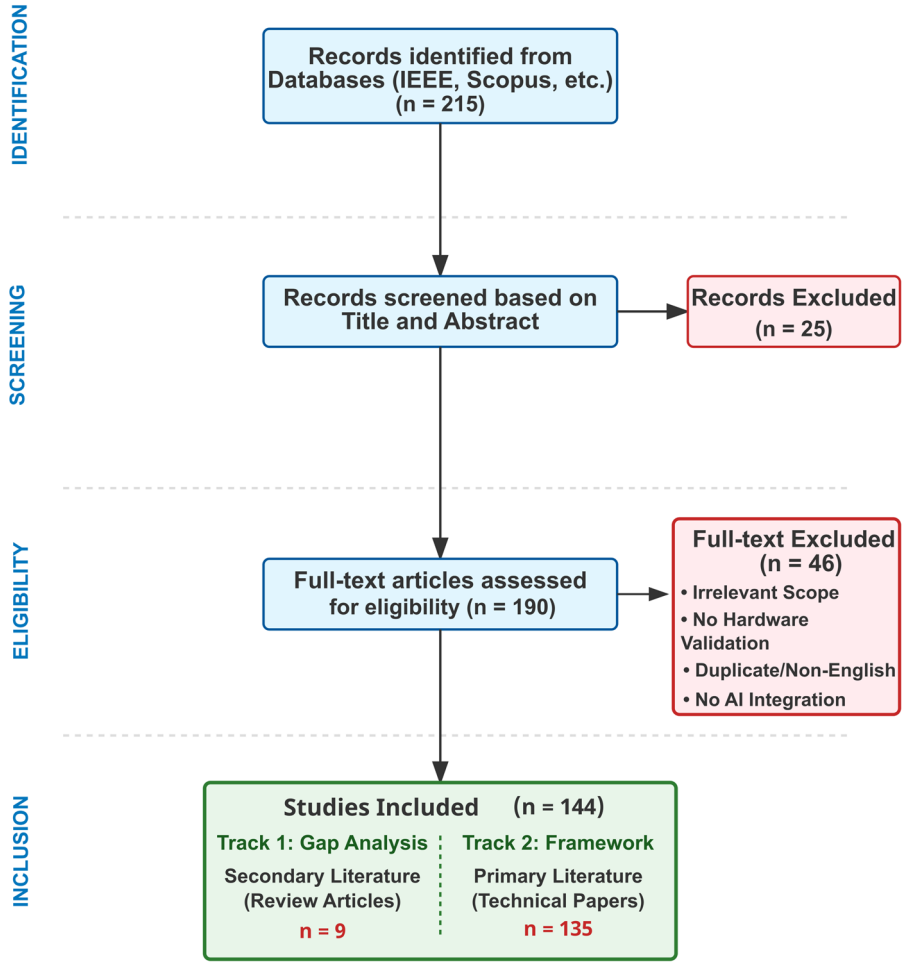


Figure 1. PRISMA flow diagram for the systematic literature search and selection methodology, with *n* representing the number of papers. The process splits at the inclusion stage to accommodate the two different research objectives: (1) the scientometric study of the secondary literature for the identification of research gaps (Track 1), and (2) the critical technical synthesis for the proposed cognitive building framework (Track 2).

2.1. Literature Search Strategy and Eligibility Criteria

A comprehensive search was executed across five major scientific databases: Web of Science, Scopus, IEEE Xplore, ScienceDirect, and Google Scholar. The primary Boolean search string was: ("Artificial Intelligence" OR "Deep Learning" OR "Reinforcement Learning" OR "Large Language Models") AND ("Internet of Things" OR "IoT" OR "Edge Computing") AND ("Building Energy Management" OR "HVAC Control" OR "Smart Buildings"). Supplementary keywords (e.g., "Physics-Informed Machine Learning," "Federated Learning") were utilized to target emerging technologies.

Moreover, to ensure relevance and quality, the inclusion criteria were strictly defined:

- Peer-reviewed journal articles and highly cited conference papers published mainly between 2020 and 2025. (Seminal foundational papers published prior to 2020 were retained where necessary).
- Studies demonstrating a clear convergence of AI algorithms, IoT sensor networks, and building energy applications.

On the other hand, editorials, whitepapers, non-English language articles, preprints from arXiv, and works which only focus on grid supply-side management without considering the effect of the demand side were excluded.

To ensure methodological transparency and reproducibility, the precise Boolean search strings adapted for each database, alongside the initial retrieval counts and the exact search commencement date (24 November 2025), are detailed in Table 2.

Table 2. Database search strategy, exact Boolean strings, and retrieval count.

Database	Exact Boolean Search String	Articles Retrieved
Web of Science	TS = ((“Artificial Intelligence” OR “Deep Learning” OR “Reinforcement Learning” OR “Large Language Models”) AND (“Internet of Things” OR “IoT” OR “Edge Computing”) AND (“Building Energy Management” OR “HVAC Control” OR “Smart Buildings”))	86
Scopus	TITLE-ABS-KEY((“Artificial Intelligence” OR “Deep Learning” OR “Reinforcement Learning” OR “Large Language Models”) AND (“Internet of Things” OR “IoT” OR “Edge Computing”) AND (“Building Energy Management” OR “HVAC Control” OR “Smart Buildings”))	118
IEEE Xplore	((“All Metadata”:”Artificial Intelligence” OR “All Metadata”:”Deep Learning” OR “All Metadata”:”Reinforcement Learning” OR “All Metadata”:”Large Language Models”) AND (“All Metadata”:”Internet of Things” OR “All Metadata”:”IoT” OR “All Metadata”:”Edge Computing”) AND (“All Metadata”:”Building Energy Management” OR “All Metadata”:”HVAC Control” OR “All Metadata”:”Smart Buildings”))	71
ScienceDirect	(“Artificial Intelligence” OR “Deep Learning” OR “Reinforcement Learning” OR “Large Language Models”) AND (“Internet of Things” OR “IoT” OR “Edge Computing”) AND (“Building Energy Management” OR “HVAC Control” OR “Smart Buildings”)	45
Google Scholar	(“Artificial Intelligence” OR “Deep Learning”) AND (“Internet of Things” OR “Edge Computing”) AND “Building Energy Management” (Note: adapted due to character limits)	22
Total Initial Harvest		342

2.2. Data Collection, Extraction, and Quality Assessment

A systematic approach to the selection of literature involved the use of a well-defined three-tiered filter (Identification, Screening, and Eligibility) that started off from a larger pool of 342 articles. After aggregating the initial search results from the five databases, duplicate publications were systematically identified and removed using Zotero’s built-in duplicate detection function, which cross-references Digital Object Identifiers (DOIs), title strings, and author lists, followed by a manual verification step. Deduplication generated 215 unique papers for the initial screening stage. The articles went through an analysis of titles and abstracts for relevance, followed by an evaluation of their technical contents and capability to address the Sim-to-Real issue. A total of 46 articles were excluded for reasons such as irrelevant scope, lack of hardware validation or AI integration, and being non-peer-reviewed or non-English, resulting in a final selection of 144 studies (see Supplementary Table S1 to check the 46 excluded full-text studies, and Supplementary Table S2 for the characteristics of included studies based on the thematic summary). This curated corpus comprises 9 secondary review articles utilized for gap analysis (Track 1) and 135 primary technical papers utilized for framework development (Track 2).

Data extraction was independently conducted by two reviewers using a standardized coding sheet to classify the literature into three framework layers: Infrastructure (sensing/semantics), Intelligence (algorithms), and Application (deployment/testing). Any discrepancies between reviewers were resolved through consensus or by consulting a third independent reviewer. Considering the heterogeneity that exists between the various technical approaches, architectures, and performance measures in the studies included, a narrative and table synthesis approach was adopted instead of quantitative meta-analysis.

Consequently, formal statistical effect measures were not calculated, missing data was explicitly recorded as "Not Reported" (NR) without imputation, and the synthesis instead focused on presenting technical performance indicators such as prediction accuracy, energy savings, and thermal comfort indices.

Two independent reviewers evaluated each primary study's research design, baselines, and uncertainty handling using a quality checklist. Since the type of analysis performed is both qualitative and technical, any quantitative assessment of publication bias or certainty of evidence, including approaches such as GRADE, was considered inappropriate. In their place, a qualitative assessment of the strength of evidence was performed, taking into account consistency, validation, and soundness of the baseline results.

3. Results Overview

Following the PRISMA guidelines, the systematic analysis of the literature included reveals that there are four main themes in the process of transitioning traditional Building Energy Management Systems (BEMS) to cognitive and AI-based structures. Rather than a traditional statistical summary, the findings are structured to represent the lifecycle of cognitive buildings, including (1) the underlying cognitive infrastructure, (2) the core algorithmic paradigms, (3) significant cyber-physical applications in energy management, and (4) the major implementation challenges bridging theoretical models to practical commercial deployment. A detailed thematic analysis of these areas is provided in the subsequent sections beginning with Section 4, which focuses on examining of the foundational physical and logical infrastructure required to enable these advanced cognitive operations.

Quality Assessment and Risk of Bias

As outlined in Section 2.2, all included studies underwent a quality assessment to evaluate the clarity of their research design, the establishment of robust baselines, including, Rule-Based Control versus AI), and the handling of uncertainty. Overall, the methodological quality of the included studies was high, particularly in recent publications (2023–2025) which demonstrated rigorous baseline comparisons. However, a recurring risk of bias identified across several predictive modeling studies was the lack of standardized validation in stochastic, real-world environments, with many algorithms performing well in idealized simulations but lacking robust uncertainty quantification when transitioned to physical testbeds.

4. The Cognitive Infrastructure: From Sensing to Semantics

The transition from automated to cognitive buildings is not based on the amount of information gathered, but rather the granularity and relevance of the information. Unlike the conventional BEMS, which operate with sparse and low-frequency data sets such as dry bulb temperature and humidity, the cognitive systems require a high-fidelity digital nervous system to sense the behavior of occupants, thermodynamic inertia, and equipment condition in real time. From a literature synthesis perspective, the studies reviewed in this section correspond to the foundational IoT infrastructure and architectural layer of AIoT-enabled buildings. As detailed in the thematic classifications in Supplementary Table S2, these works primarily focus on the sensing technologies, semantic interoperability frameworks, and edge-to-cloud data architectures required to enable higher-level algorithmic intelligence.

4.1. Advanced IoT Sensing

Passive Infrared (PIR) sensors form an important class of commonly employed sensors for the purpose of monitoring the presence of occupants in buildings, whether they are

commercial or residential. This is due to their low cost, privacy-preserving properties, and low computational demands. However, PIR sensors provide binary information regarding the presence of occupants and tend to produce false negatives for persons at rest or with low levels of motion [25]. This can cause aggressive setbacks in HVAC systems, reducing occupant comfort. In an effort to overcome this "observability gap," the IoT sensing layer is moving towards multimodal sensing technologies that are also privacy-preserving. One notable improvement in NI sensing is the utilization of existing infrastructure as a sensing mechanism. Unlike Received Signal Strength Indicator (RSSI) technology, which provides general proximity data, Channel State Information (CSI) allows for detailed data on multipath propagation of wireless signals [26]. As the signal passes through individuals and objects, the change in signal strength and phase can be analyzed using ML to identify individuals and even their activities and respiratory rates, without any invasion of privacy that camera surveillance often entails [27,28]. This allows the building to distinguish between a transient individual and an employee who spends most of their time there, so that appropriate energy consumption can be provided.

Moreover, to overcome the drawbacks of using air temperature sensors, which often do not correlate with the Mean Radiant Temperature (MRT) perceived by the people inside the space, low-cost infrared thermography is being integrated in the edge nodes to create heat maps for the space, thus identifying areas where there are hot and cold spots due to solar gain or inadequate insulation. It also allows the estimation of the metabolic rate and clothing level (clo values) to control the HVAC based on the perceived physiological comfort (PMV) rather than fixed setpoints [29,30]. On the other hand, the ubiquitous presence of these advanced sensors results in the occurrence of the data heterogeneity problem. This heterogeneity arises from multiple dimensions including differences in sensor types, data ontologies, temporal irregularities, spatial diversity, and varying data quality across devices and locations [31]. To put it differently, a cognitive building has to consume and integrate different data sources. These sources include 1D high-frequency time series data from smart meters, 2D spatial data from thermal images, and unstructured semantic logs. This results in an exacerbation of the "Data Rich, Information Poor" paradox, as the inherent complexity of synchronizing the data from different vendors results in the negation of the energy savings achievable through the utilization of the real-time control [32].

4.2. Semantic Interoperability or the Language of Cognitive Buildings

The first barrier to the implementation of a scalable AI system is the lack of a standard ontology [33]. In traditional Building Energy Management Systems (BEMS), data is typically given a label that is a cryptic code (for instance, AHU_1_SAT or Z1_T_Stpt). A human might understand that AHU_1_SAT refers to the temperature of the supply air from Air Handling Unit 1, but an AI system sees only a series of arbitrary strings. It is not apparent that AHU_1 supplies air to VAV_Box_12, which in turn supplies Conference_Room_B. Without semantic interoperability, the integration of these interdependent applications becomes technically challenging, labor-intensive, and costly, which can hinder optimization, fault diagnosis, and grid service delivery for grid-interactive efficient buildings [34]. An algorithm cannot perform causal reasoning or fault propagation analysis without a unified understanding of these diverse data streams. To solve this problem, the field is shifting from flat-based naming to graph-based semantic models like Brick Schema, Project Haystack 4.0, and Smart Applications REFerence ontology (SAREF). By using formal ontologies based on standards like Resource Description Framework (RDF) and Web Ontology Language (OWL) [34,35], these models create directed graphs where nodes represent physical, logical, and virtual entities (e.g., Equipment, Location, Point), and edges define relationships between them (e.g., feeds, isLocationOf, controls). This machine-readable context can be used by inference

engines to derive implicit information from various sources. For instance, the AI agent might ask the building structure itself, "Identify all Variable Air Volume (VAV) boxes served by AHU-2, which are currently reporting a cooling error." This would be useful for scalable applications like context-aware control and automated fault detection and diagnostics [36,37].

Enabling Agentic Reasoning

This semantic foundation provides the underlying rationale for Agentic AI (as discussed in Section 4). An LLM-based assistant for facility management cannot perform a natural language command such as "Optimize cooling in the north wing" unless it can perform a semantic mapping from the vague concept of "north wing" to a set of specific HVAC zones and "cooling" to the corresponding set of actuator setpoints [38]. Semantic graphs bridge the gap by relating abstract user intent to specific device-level operations. This interoperability is equally important for portability [39]. In fact, Model Predictive Control (MPC) frameworks leverage such models to parameterize the physical profiles of HVAC zones, thereby enabling modular control sequences to be ported across different buildings without the need for reconfiguration [40].

4.3. Edge Computing and TinyML

In spite of the vast computational powers offered by cloud-centric architectures, there are considerable bottlenecks in latency, bandwidth, and privacy, making them inappropriate for real-time control and safe building management. The round-trip transmission of high-frequency sensor data, such as 10 kHz voltage sampling, to a remote server causes deterministic latency, making immediate fault mitigation impossible [41]. To overcome this drawback, a new paradigm has been proposed, moving from cloud-centric to Edge Computing and TinyML, which shifts the computation directly to the sensor node.

Tiny Machine Learning (TinyML) describes a paradigm shift in which quantized neural networks, usually reduced from 32-bit floating-point precision to 8-bit integer precision, are deployed on resource-constrained microcontrollers (MCUs) that only contain kilobytes of RAM [42,43]. Unlike the often-used term "Edge AI," which often requires powerful gateway devices (e.g., Raspberry Pi, NVIDIA Jetson), TinyML is used in milliwatt-class end devices (e.g., ESP32, ARM Cortex-M). This allows a vibration sensor in a chiller to process the data from the accelerometers locally and only send a semantic message, such as "Bearing Failure Detected," rather than gigabytes of raw waveforms. This is referred to as a "Compute-forward" approach [43,44].

Another important, yet often neglected, reason for decentralization is the environmental impact of AI itself. The term "Green AI" emphasizes the energy efficiency of the entire model development cycle. The inference in the cloud is a major contributor to greenhouse emissions due to the cost of transmission as well as data center cooling [45]. TinyML eliminates the energy cost of transmission. Studies have shown that local data processing on edge devices is significantly more energy-efficient than transmitting sensor data to the cloud, thereby aligning the AI architecture with the building's net-zero goals. Specifically, research suggests that edge computing helps save energy by reducing the need for continuous and voluminous data transmission through the network, as it forms one of the major areas of energy consumption. In one research, it was observed that the overall energy consumption of the system, as well as the emission of greenhouse gases, could be cut down by 36% and 43%, respectively, with the implementation of the edge–fog–cloud system in smart agriculture, as compared to the cloud system alone, with the added advantage of reducing the network traffic by up to 86% [46]. Another research showed the reduction in the overall energy consumption of the system by 65.7% and the time taken for completing

the tasks by 37%, as compared to the cloud system, by implementing local feature extraction in the edge IoT sensor nodes [47]. The intelligent management of tasks and filtering of the data by the edge computing systems also helps in the reduction in energy consumption by up to 40% [48,49].

4.4. Hierarchical Control Architecture

To minimize latency and energy consumption while preserving data privacy at the source, the proposed cognitive framework utilizes a bi-directional, hierarchical topology, as illustrated in Figure 2. The communication flow relies on bottom–up data aggregation and top–down model and control updates across three interconnected tiers:

- Device Layer (TinyML): Manages sub-second reflex actions (e.g., adjusting lighting levels based on real-time occupancy) and privacy-critical feature extraction [50].
- Edge Layer (Fog): Manages zone-level data aggregation to optimize the operation of various subsystems (e.g., VAV systems) and supports short-term storage [51].
- Cloud Layer: Reserved for intensive training, long-term pattern recognition (seasonality), and global fleet management [52].

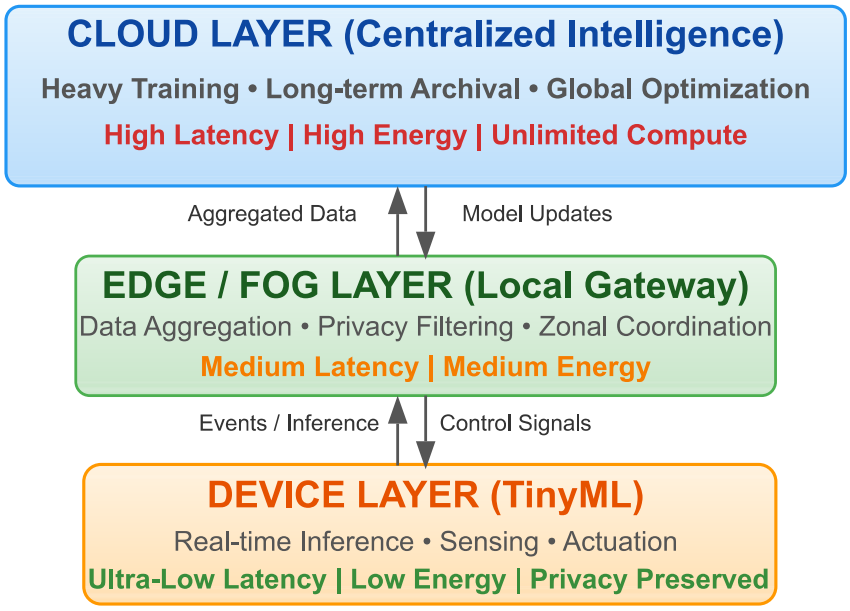


Figure 2. The proposed hierarchical AIoT architecture for cognitive buildings. This conceptual framework illustrates a three-tier topology designed to offload computational tasks and optimize energy efficiency. (1) The Cloud Layer acts as centralized intelligence, handling heavy training, global optimization, and long-term archival (characterized by high latency and unlimited compute). (2) The Edge/Fog Layer serves as a local gateway for zonal coordination, privacy filtering, and data aggregation. (3) The Device Layer (TinyML) operates at the extreme edge, managing real-time inference, sensing, and actuation with ultra-low latency and minimal energy consumption. The arrows demonstrate the bi-directional communication flow: bottom–up transmission of "Events/Inference" and "Aggregated Data," and top–down distribution of "Model Updates" and "Control Signals.

5. Algorithmic Paradigms and the Intelligence Core

Following the establishment of the cognitive framework in Section 3, the focus now shifts to the system's core: the algorithmic paradigms that process raw sensor data to provide actionable control strategies. This section examines the progression from predictive models (Deep Learning) to adaptive control (Reinforcement Learning), and finally to the newly emerging field of generative and reasoning agents.

5.1. Predictive Modeling in Cognitive Buildings

Essentially, the core capability that all cognitive buildings are built upon is that of forecasting, in terms of predicting the load demand, renewable energy (PV/Wind), and occupancy for the purpose of proactive control. While the first deployments relied on statistical techniques, such as ARIMA and Support Vector Regression (SVR), the non-linear and stochastic nature of modern building energy dynamics has necessitated the move from statistical techniques to Deep Neural Networks (DNN).

For the past decade, Long Short-Term Memory (LSTM) networks have served as the industry standard for time series forecasting. In contrast to standard Recurrent Neural Networks (RNNs), LSTM networks address the "vanishing gradient problem" via gating mechanisms (input, output, and forget gates) and are thus capable of learning temporal dependencies for longer time horizons, such as a 24 h profile. Yet, LSTM networks face a major "sequential bottleneck": time series are processed sequentially ($t_1 \rightarrow t_2 \rightarrow t_3$), and thus cannot be parallelized, limiting their ability to learn ultra-long-term dependencies, such as yearly seasonality or complex holiday effects [53,54].

Recent advances (2022–2025) have facilitated the migration of Transformer architectures from their origins in Natural Language Processing (NLP) to the energy domain, as demonstrated in [55–57]. The LSTM transmits data via bottleneck states through linear processing, which often leads to the loss of information after prolonged time periods. On the other hand, the use of a global attention mechanism in the Transformer allows the model to apply different numerical values to each point in time in one go, thus eliminating sequential limitations of recurrent networks.

More advanced models, such as the Temporal Fusion Transformer (TFT) and Informer, have been shown to have high performance in handling multi-scale seasonality [58] (e.g., recognizing that Monday's load profile resembles last Monday's more than yesterday's Sunday profile). By ingesting static metadata (building type) alongside dynamic inputs (weather), TFTs provide interpretable forecasts where the "attention weights" reveal exactly which features (e.g., humidity vs. hour-of-day) drove a specific prediction peak [59].

While supervised DL is highly effective for forecasting exogenous factors such as weather and occupancy, it is nonetheless a passive technique that predicts the future but does not influence it. The move towards cognitive buildings, however, calls for active and adaptive control strategies, which can navigate the complexities of multi-objective trade-offs in real time [60]. This is where Deep Reinforcement Learning (DRL) comes into play. DRL differs from RBC, which is based on static "if–then" logic, and also from MPC, which is based on computationally expensive physical models [61]. DRL is, instead, a model-free technique. The DRL agent learns optimal control policies (π), through trial and error, with the aim of maximizing a cumulative reward signal, which balances the objectives of minimizing energy consumption with the constraints of thermal comfort (e.g., Predicted Mean Vote) [62].

The algorithmic landscape for HVAC control is currently dominated by policy-gradient methods, most notably Proximal Policy Optimization (PPO) and Soft Actor–Critic (SAC). The PPO has been established as the new standard based on the training properties of the algorithm, which is stable and less sensitive to the tuning of hyperparameters. This makes it highly suitable for scenarios with large amounts of simulated interactions. The SAC, on the other hand, is an off-policy algorithm that focuses on maximizing sample efficiency through the use of past experiences in the replay buffer to enhance learning with reduced interactions [63]. The algorithm also incorporates an entropy maximization component to encourage exploration, which is highly important considering the non-convex optimization problem in HVAC control [64].

The academic literature predominantly describes the theoretical benefits of DRL agents over commonly used proportional–integral–derivative (PID) controllers. In simulated environments for HVAC systems, deep reinforcement learning (DRL) controllers have shown energy savings compared to baseline control strategies, which vary between 6% and over 40%. Several research studies have shown that energy savings between 10% and 25% can be achieved, including a 25% reduction in energy intensity for low-energy buildings and 10% to 15% compared to rule-based and model predictive controls [64,65]. Other research shows DRL achieving approximately 21.4% energy savings over rule-based controllers, with some model predictive control methods slightly outperforming DRL in specific scenarios [15]. Moreover, advanced configurations of DRL, such as multi-zone or multi-input configurations, can achieve 37% savings while maintaining the thermal comfort level [66]. Notably, the use of the SAC algorithm can achieve 43% savings in the total annual energy consumption while maintaining the balance between the comfort level of the occupants and the energy savings [67]. In addition, model-based RL can achieve 8–10% more savings than the state-of-the-art methods while requiring reduced data for the model [68,69]. In these virtual sandbox conditions, the agent is permitted to explore extreme control strategies, such as momentarily disabling the chiller during peak demand periods to gauge the building's inertia response, at millisecond resolutions over millions of simulated scenarios without actual-world repercussions.

However, this impressive simulation performance rarely translates directly to physical deployment, a phenomenon known as the "Sim-to-Real" gap (Figure 3). Most DRL HVAC studies remain validated in simulation, with comparatively few real-world deployments [70,71]. As depicted in the diagram, an agent can easily achieve optimal performance during the training phase within a deterministic, risk-free simulation. However, transferring this policy directly to a physical building often leads to erratic behavior. The diagram visually captures this environmental mismatch that the clean feedback loop of the simulation is disrupted in reality by stochastic real-world dynamics. This stochasticity (represented by the visual noise applied to the physical building block and the dashed lines indicating unpredictable state/reward feedback) causes the overfitted agent to execute unsafe exploratory actions, ultimately resulting in control oscillations and comfort violations.

Although some recent studies have begun bridging this gap by deploying DRL controllers in operational buildings, demonstrating successful transfer from simulation to reality under certain conditions, such as static thermal comfort limits [72–75], the gap exists because no simulator, regardless of its fidelity, can perfectly capture the stochastic dynamics of a real building. Sensor measurement noises, actuator aging/saturation effects, unmodeled air infiltration, and random occupant behaviors all contribute to disturbances that are difficult to model with perfect accuracy. This creates a dynamic mismatch between the training scenario and the real world. An agent trained with DRL techniques to perform well in a deterministic training scenario may perform poorly in the presence of real-world uncertainties. This may manifest as control oscillations, actuator chattering, mechanical degradation, and violation of thermal comfort constraints [75].

Furthermore, the underlying principle of the training methodology of RL, i.e., exploration, is inherently unsafe for application in real-world buildings. A DRL agent needs to explore pseudo-randomly to find better policies. In a real-world building, this means that a model that has not been trained cannot be allowed to randomly overheat a conference room or over-pressurize the ducts in order to learn that it was a bad policy. As a result, the current research is moving in the direction of methodologies that can bridge this gap, such as Offline RL, where the agent is trained on historical data from a Building Energy Management System, and Domain Randomization, where the agent is trained across a wide range of simulation parameters [76]. To address this, safe reinforcement learning uses

tools like safety layers, formal guarantees (such as Lyapunov functions), and constraint enforcement to ensure that the actions remain within safe limits during training as well as at runtime [76].

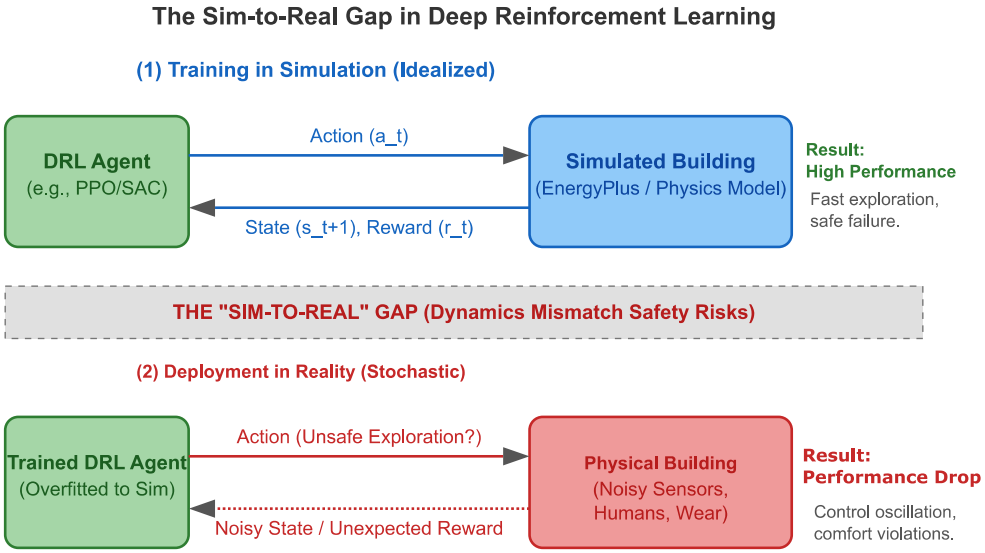


Figure 3. Visual representation of the "Sim-to-Real" gap in Deep Reinforcement Learning (DRL) for HVAC control. This diagram contrasts the training and deployment phases of DRL agents to highlight environmental mismatch. (1) Training in Simulation (Idealized): Shows a deterministic feedback loop where an agent (e.g., PPO/SAC) interacts with a physics-based model. The solid arrows representing action (a_t), state (S_{t+1}), and reward (r_t) denote a clean, risk-free environment leading to high performance. (2) Deployment in Reality (Stochastic): Depicts the physical building environment, visually emphasized with a textured noise filter, characterized by unpredictable elements like sensor noise, wear, and human presence. The dashed return arrow signifies the reception of noisy states and unexpected rewards, illustrating how an agent overfitted to simulation experiences severe performance drops, control oscillations, and safety risks in the real world.

5.2. The New Frontier, Generative and Agentic AI

A notable recent development in the last few years (2023–2025) has been to transition from simple Chatbots to true Reasoning Agents that can make autonomous decisions [77]. Whereas DL thrives at predictions (pattern recognition) and reinforcement learning excels at optimization problems (learning to maximize rewards), traditional task-specific DL and RL systems are not inherently designed for open-ended semantic reasoning [78]. Agentic AI addresses this by using the large language model as a cognitive engine to perceive the environment, break down complex goals into sub-tasks, and take action via API-driven control of the building's physical systems [79]. Figure 4 maps the architecture of an Agentic Control Loop. As visually depicted, this framework relies on a continuous three-phase cycle (Perceive, Reason, Act). A technical bridge shown in the transition from the *Reason* to the *Act* phase is the model's ability to output structured data. By formulating outputs as strictly typed JSON objects (illustrated in Figure 4 action block), the non-deterministic LLM can safely and predictably trigger API commands within the deterministic physical building infrastructure. Furthermore, as indicated by the dashed feedback line, this is not a one-way command; the agent continuously ingests new state feedback to evaluate the physical consequences of its actions and adjust subsequent iterations.

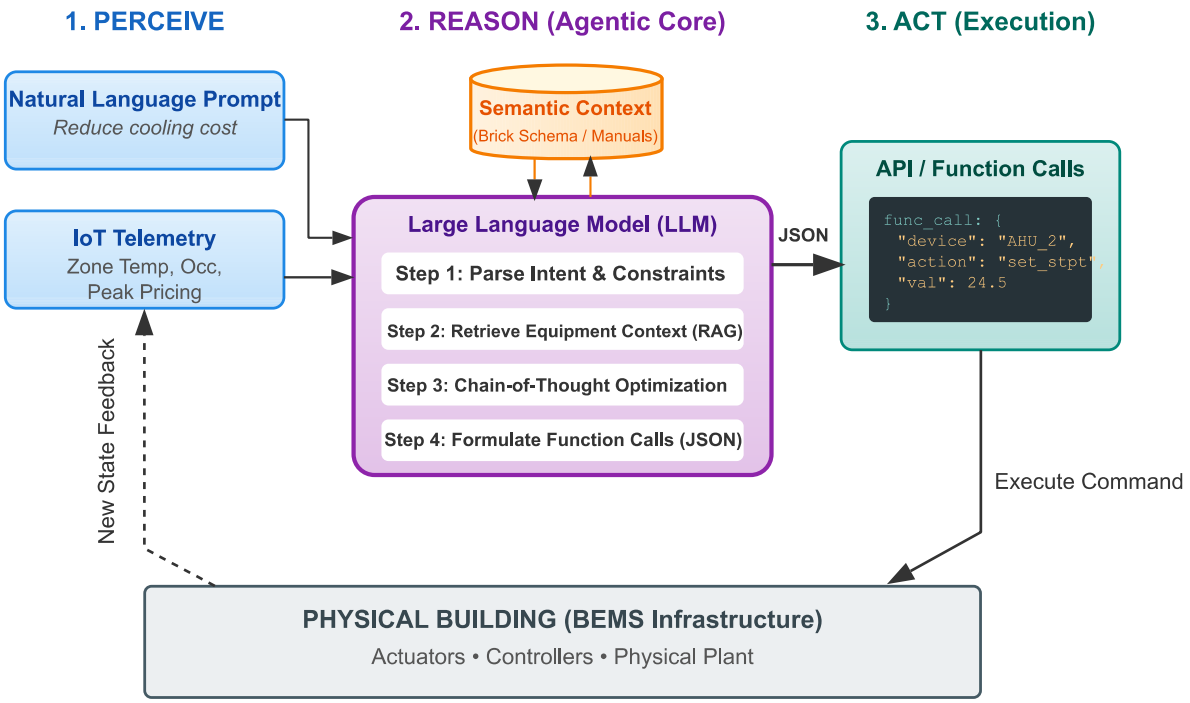


Figure 4. Architectural flow of the Agentic Control Loop. This diagram illustrates LLMs acting as closed-loop orchestrators across three phases. 1. PERCEIVE: The system fuses natural language user intents with real-time IoT telemetry. 2. REASON: The LLM processes inputs via RAG (for spatial context) and Chain-of-Thought (for optimization) to formulate actions. 3. ACT: Conceptual reasoning is translated into deterministic API function calls (visualized by the JSON block) executed by the BEMS. The dashed arrow denotes continuous state feedback, closing the loop for real-time adaptation.

The essential feature of an agentic system is its ability to act on the physical world. This ability to act on the physical world is enabled by Function Calling or Tool Use. In this model, it is not just a matter of the LLM responding with natural language. Rather, it has a set of executable APIs at its disposal, such as *get_zone_temp()*, *set_damper_position()*, or *query_energy_price()* under the Building Management System API [80]. If a facilities manager gives an order such as "Reduce peak load by 10% without impacting the executive meeting in room 304," it does not simply acknowledge receipt of the order. Rather, it analyzes what it has to do to execute on that order. It then creates a structured set of machine-readable commands in JSON format to be sent to the BEMS [81,82]. These are operational changes to the building's systems while at the same time maintaining set points in a given conference room.

In contrast to many task-specific deep neural networks, which can be thought of as black box approximators, agentic models based on large language models can often incorporate intermediate reasoning steps, often in the form of Chain-of-Thought (CoT) Prompting. In the event of conflicting objectives, such as increased energy cost signals with comfort complaints, the agent can generate a reasoning trace to evaluate possible trade-offs prior to executing any action. In one example, instead of immediately adjusting temperature setpoints, the system can query contextual variables, such as solar gains or shading, to explore less impactful actions. Moreover, they enable multi-objective infeasibility diagnosis in routing problems, generating diverse diagnostic suggestions that involve trade-offs between path cost and constraint violations. In some cases, intermediate reasoning can be logged to aid human evaluation, which may offer increased transparency compared to traditional deep reinforcement learning approaches, often plagued by low interpretability [83]. However, the accuracy of such intermediate reasoning traces with regard to the internal model decision process is still an area of ongoing study. Studies

show that CoT outputs can be plausible yet misleading rationalizations that do not reveal the actual factors driving the model's predictions, especially when models are biased or manipulated, leading to unfaithful explanations [84].

To function in domain-specific environments, agentic systems often use Retrieval-Augmented Generation (RAG) to tap into external knowledge sources during inference. Large language models do not have built-in knowledge of the most current or site-specific operational knowledge, such as which Air Handling Unit (AHU) serves a particular zone [85]. By tapping into structured data repositories like vector databases containing operations and maintenance (O&M) documentation or semantic building representations like the Brick Schema, the system may retrieve relevant contextual information before generating its response or action. The retrieval step allows the system to condition its responses on the specific topology and metadata of the building [34]. Thus, general engineering knowledge (such as the relationship between pressure and fan energy consumption) may be grounded in site-specific system configurations.

In this architecture, Agentic AI is positioned as a high-level coordinator that interacts with specialized analytical or control modules. Instead of replacing existing techniques, it can orchestrate existing lower-level components, for example, by using a transformer-based forecasting approach to estimate building loads, then using an MPC solver to compute an optimal control action, or by sending control commands to the building automation systems while monitoring feedback from operations [86]. This hierarchical structure decouples semantic planning from numerical optimization, allowing for the integration of data-driven forecasting, physics-based control, or language-based decision-making. These hybrid architectures can be seen as conceptually related to recently proposed Neuro-Symbolic AI frameworks, in which neural networks are combined with reasoning or optimization components [87]. While there is little evidence on how this orchestration approach can be applied in real-world building settings, it is intended to benefit from the synergies between different families of ML and control techniques [88].

5.3. Physics-Informed Machine Learning (PIML)

Despite their flexibility and capacity to model complex relationships within their data distributions, DL and RL models have been shown to lack reliability in extrapolating beyond their learned distributions. As such models do not inherently encode any physical conservation relationships, their predictions may become inconsistent with physics in under-represented regimes (e.g., extreme weather conditions or fault regimes) [89]. In building energy systems, this may correspond to thermodynamic constraints or unrealistic behavior. This has led to a 'trust gap' among facility operators in safety-critical/high-reliability settings.

Physics-Informed Machine Learning (PIML) seeks to address this problem by leveraging prior knowledge in the form of conservation relationships, thermodynamic constraints, or differential equation residuals within data-driven learning frameworks. By encoding prior knowledge in data-driven learning models, PIML seeks to improve upon aspects of physics consistency, generalizability beyond their learned distributions, and interpretability of learned relationships relative to their 'black box' counterparts [90,91]. A widely studied realization of physics-informed learning is the Physics-Informed Neural Network (PINN). Unlike traditional neural networks that are solely driven by data and only focus on minimizing the data term, PINNs are designed to incorporate physics constraints into the loss function. In this regard, the overall objective function of physics-informed neural networks is given by the combination of the data term and the physics term, which represents the residual of the governing equations such as partial differential equations (PDEs) that govern the physical phenomenon of interest, e.g., the PDE of heat conduction or the PDE of

conservation of energy [92]. By enforcing penalties for violations of the aforementioned governing equations during the optimization process, the model is effectively motivated to produce outputs that are not only consistent with the data but also physically plausible. Figure 5 details the architecture of a PINN. Unlike purely data-driven models, the PINN framework utilizes a dual-objective optimization process. As visually depicted by the summation (Σ) node, the network simultaneously evaluates a standard "Data Loss" and a "Physics Residual". By heavily penalizing deviations from governing partial differential equations (such as the PDE $\partial u/\partial t - \alpha \nabla^2 u = 0$), the dashed backpropagation loop forces the neural network to update its weights in a manner that strictly adheres to the physical laws of building thermodynamics.

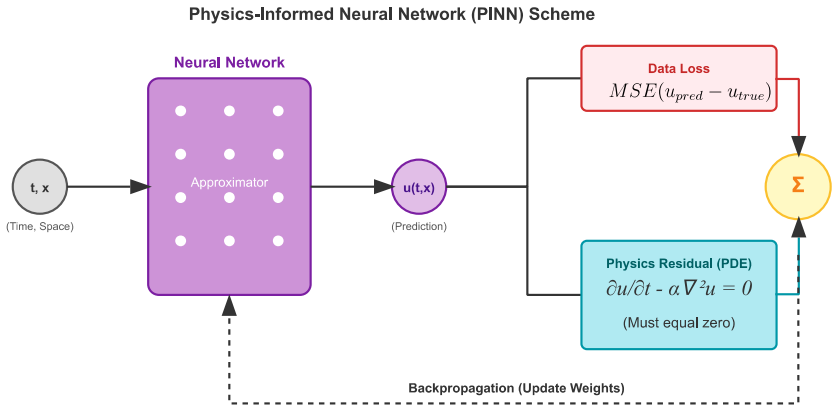


Figure 5. Schematic of a Physics-Informed Neural Network (PINN). Spatiotemporal inputs $(\boldsymbol{t}, \boldsymbol{x})$ are fed into a Neural Network Approximator to generate a prediction $\boldsymbol{u}(\boldsymbol{t}, \boldsymbol{x})$. The total loss (Σ) is computed by aggregating two distinct components: the standard Data Loss (fitting observational data) and the Physics Residual (enforcing the governing partial differential equation, shown here as the diffusion equation $\partial u/\partial t - \alpha \nabla^2 u = 0$). The dashed arrow represents the backpropagation step, which updates network weights to minimize this combined loss, ensuring predictions remain physically consistent.

The overall objective function can therefore be defined as a weighted sum of the data residual and physics residual terms [93]:

$$\mathcal{L} = \omega_{\mathrm{d}} \mathcal{L}_{\mathrm{d}} + \omega_f \mathcal{L}_{\mathrm{f}} + \omega_{\mathrm{b}} \mathcal{L}_{\mathrm{b}} \quad (1)$$

where ω_{d}, ω_{f}, ω_{b} are weighting coefficients balancing data fidelity and physical consistency, $\mathcal{L}_{\mathrm{d}}$ is the Loss from data/observations, and $\mathcal{L}_{\mathrm{f}}$ represents the Loss from PDE residuals.

The integration of the governing equations as soft constraints can allow the PIML method to be used as a form of regularization in the development of AI-based control models. For example, in the development of building thermal dynamics models, the learning process can be constrained in such a way that the energy balance equations, which implicitly include the effect of thermal inertia, are used to ensure that the indoor temperatures predicted by the model vary in a smooth, rather than discontinuous, fashion [94]. This is due to the fact that, when the neural network predicts discontinuous behavior, the residual part of the loss function increases, which helps in preventing physically inconsistent behavior. Such a factor is important when developing AI-based controllers, especially when it comes to buildings, as this can impact the safety of occupants.

One of the most commonly cited benefits of PIML relates to its increased sample efficiency in particular scenarios. In general, data-centric models require large amounts of data in order to learn system dynamics directly from available data. By contrast, PIML approaches, as exemplified by PINNs, constrain the hypothesis space to solutions that are consistent with well-established laws of physics. This inductive bias can reduce the

amount of data required to achieve acceptable predictive performance, particularly when measurements are sparse [95]. However, while physics-informed machine learning (PIML) techniques improve data efficiency and speed up training, there are difficulties in achieving computational efficiency and accuracy in predictions [96]. Similarly, purely data-driven neural networks often show reduced performance when extrapolating outside the training distribution. PINNs have the advantage that they are based on invariant physical laws like conservation laws or heat transfer equations. This means that physics-informed neural networks may perform better when there are distributional shifts, as long as the physics are valid. However, this does not mean that extrapolation performance is guaranteed outside the range where the physics are valid [96].

With respect to building thermal modeling, the energy balance equations represent the heat transfer relations, such as $Q = U \cdot A \cdot \Delta T$, and these relations remain structurally valid throughout the seasons. Incorporating the relations into the learning process can help alleviate the cold-start problem in newly constructed buildings or those with sparse monitoring, though the performance still relies on the calibration of parameters and data quality [97] (see Table 3 for a detailed comparison of these modeling paradigms).

Table 3. Comparative Analysis of Modeling Paradigms.

Feature	White Box (First Principles)	Black Box (Pure Data-Driven)	Physics-Informed (PIML)
Basis	Physical Laws (Thermodynamics)	Statistical Correlations	Hybrid (Laws + Data)
Data Requirement	Low (only for calibration)	Very High (Big Data)	Moderate (Data-Efficient)
Interpretability	High (Explicit parameters)	Low (Opaque weights)	High (Physically consistent)
Generalization	Excellent	Poor (Overfitting risk)	Acceptable (Constrained by physics)
Computational Cost	High (Complex solvers)	Low (Fast inference)	Low (Fast inference)

6. Strategic Applications in Energy Management

The transition from prediction and reasoning algorithms into actual physical execution is most clearly visible in the active control of building subsystems. This section will discuss the application of the cognitive framework, as described in Sections 3 and 4, in actual-world cyber-physical problems, and will begin with the most energy-intensive component of the built environment, Heating, Ventilation, and Air Conditioning (HVAC).

6.1. HVAC Optimization, the Trade-Off Between Energy and Comfort

HVAC systems typically account for approximately 40% of total building energy consumption [66], making them a major focus of operational optimization efforts. In many legacy BEMS deployments, HVAC control is implemented using static RBC strategies and fixed temperature deadbands. Although deterministic control schemes are simple to compute and reliable, they might fail to appropriately describe the nonlinear and stochastic nature of building thermal characteristics and occupant behavior. As a result, the BEMS might operate suboptimally and cause problems such as energy inefficiencies (e.g., reheating after overcooling) and reduced comfort levels simultaneously [98,99]. New "cognitive building" paradigms redefine the HVAC control task as a constrained, possibly non-convex Multi-Objective Optimization (MOO) problem. In this context, the objectives (e.g., energy savings or costs, comfort levels) are simultaneously optimized under a set of constraints [100].

To rate thermal comfort, the most common approach in the industry is to use Fanger's PMV and the associated Predicted Percentage of Dissatisfied (PPD) indices, as defined in ASHRAE 55 and ISO 7730, respectively [101]. PMV estimates the mean thermal sensation vote of a large population on a seven-point scale ranging from −3 (cold) to +3 (hot) [102]. For a large number of mechanically conditioned buildings, comfort zones are generally

defined as $-0.5 \leq PMV \leq +0.5$, but adaptive comfort standards can extend the range in various ways depending on climate and building type [101]. PMV is calculated as a non-linear function of six primary parameters, encompassing two physiological variables (such as metabolic rate and clothing insulation) and four environmental variables (such as air temperature, mean radiant temperature, air velocity, and humidity) [103]. In the traditional implementation of HVAC systems, personal parameters are often treated as constant values, as it is difficult to measure them. In the context of the new concept of "cognitive building," the personal parameters are determined using various IoT sensors, and the comfort levels are achieved based on the calculated parameters rather than the temperature [102]. Sophisticated control schemes, including MPC and DRL, often formalize the energy-comfort trade-off as a constrained or multi-objective optimization problem. A typical cost function minimizes total energy consumption while penalizing deviations from established thermal comfort bounds, utilizing dynamically adjustable weighting factors to balance these competing objectives [104].

In the predictive control schemes, the effects of external disturbances like weather and internal heat gains are predicted over a finite period. Then the building's thermal mass is used in a strategic way, e.g., by precooling during low-cost periods and allowing temperature to drift during peak periods, satisfying comfort constraints like maintaining the PMV within prescribed bounds [105]. Nevertheless, the strict enforcement of PMV conformity is subject to the model's fidelity, the formulation of the constraints, and the consideration of the associated uncertainties; in practice, comfort is often ensured in a probabilistic manner or in relation to a range of margins rather than as a strict requirement. As the objectives of energy reduction and comfort maximization are in direct conflict with each other, there is no notion of a perfect solution; instead, the AI develops a Pareto Optimal Front (Figure 6). A control strategy that falls on this front signifies that it is an optimal solution where any further reduction in energy consumption will be accompanied by a concomitant decrease in human comfort. The Agentic AI, by constantly calculating this optimal front, will be able to move the building's point along the curve based on the semantic context. For example, in a normal workday scenario, it will be in "Balanced Mode" on the Pareto Optimal Front, but upon receiving an automated DR signal from the smart grid, it will move along the curve to reach "Energy Saving Mode" while allowing the PMV to reach the absolute edge of the ASHRAE limit without compromising compliance [106].

6.2. Grid-Interactive Efficient Buildings

The progression towards intelligent buildings inherently leads to the concept of Grid-Interactive Efficient Buildings (GEBs). Traditionally, the built environment has functioned as a passive consumer of electric energy, utilizing power without consideration for macroeconomic grid conditions. However, the adoption of intermittent renewable energy sources, such as solar and wind, has disrupted traditional supply-side macroeconomic grid stability [107]. GEBs address this disruption by transforming buildings from passive loads to dynamic, dispatchable energy producers (prosumers) that can supply bidirectional grid services. This is largely a function of the cognitive architecture discussed within the previous sections. Without the predictive accuracy of DL algorithms or the semantic interoperability of standardized ontologies, it would be computationally intractable to coordinate real-time grid signals with building-level constraints [107,108]. The most advanced use case within the GEB framework is automated DR and peak shaving. In the traditional DR approach, there is a lack of sophisticated and intelligent control, and it is often achieved by simplistic and unrefined means, such as shutting down chillers or dimming lights using a single, unified signal during an emergency, which results in considerable discomfort to the occupants [109]. On the other hand, cognitive GEBs employ sophisticated control

strategies like DRL and MPC to perform preemptive and optimized peak shaving, hence allowing for smart, flexible, and comfortable energy management [110]. These strategies allow buildings to adjust their loads and energy storage in response to anticipated conditions on the grid, thus improving grid stability and occupant comfort. According to studies, intelligent control strategies are able to significantly lower peak demand by over 40%, at times, through the use of multiple energy resources such as batteries, solar power, and electric vehicles [111,112].

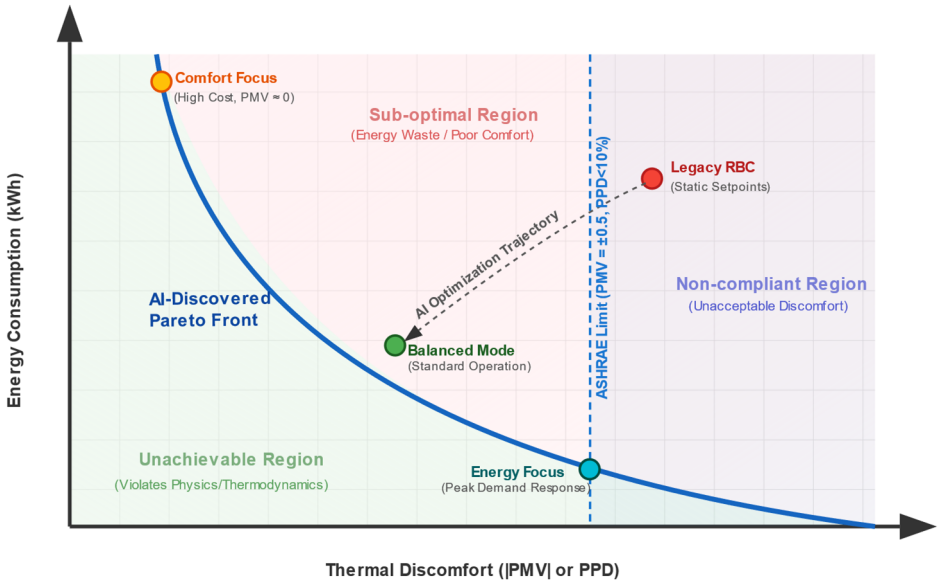


Figure 6. Multi-objective Pareto Front for HVAC optimization, including the fundamental trade-off between Energy Consumption and Thermal Discomfort (|PMV| or PPD). Traditional Rule-Based Control (RBC) operates in the "Sub-optimal Region," expending excess energy. AI architectures (DRL/MPC) drive the system along an optimization trajectory to the AI-Discovered Pareto Front. This allows for dynamic shifting between operating points (Comfort Focus, Balanced Mode, and Energy Focus) based on real-time grid needs. The vertical dashed line represents the strict ASHRAE compliance limit (PMV = ±0.5, PPD < 10%), preventing the system from entering the "Non-compliant Region.".

Beyond load curtailment, load curtailment, GEBs integrate distributed resources such as rooftop PV and battery energy storage systems (BESS). Energy arbitrage can be expressed as an optimization problem constrained by the goal of minimizing cost of operations. The objective function compares time-varying cost of electricity to cost of degradation of the battery, and constraints limit the state-of-charge and efficiencies of charging and discharging, available locally generated PV power, and the load demand of the building at any point in time [113,114].

The charging occurs during low prices or surplus PV generation, and discharge occurs during peak pricing periods. The strategy, therefore, enables cost minimization under operational and regulatory constraints. For further improvement of this process, more sophisticated cognitive systems could implement the high-frequency Maximum Power Point Tracking (MPPT) telemetry in the semantic knowledge graph itself. This will enable the agentic AI to adapt the load shifting and battery dispatching based on changes in efficiency and deterioration rates of the PV systems in real time [112,114]. To demonstrate the practical impact of AI-driven demand response, Figure 7 illustrates the daily energy arbitrage and load shifting capabilities of a GEB. Unlike the baseline RBC load, which sharply spikes during the highly punitive Peak Pricing Window, the AI-optimized controller systematically restructures the load profile. As visually detailed, the AI executes a multi-stage arbitrage strategy: it utilizes off-peak electricity for "Thermal Pre-cooling", absorbs

surplus onsite solar generation via "BESS Charging" at midday, and ultimately performs "Peak Shaving" by discharging batteries and coasting HVAC systems when grid prices peak. This active orchestration (shifting the blue solid line away from the red shaded region) strictly limits grid power draw during critical late-afternoon hours.

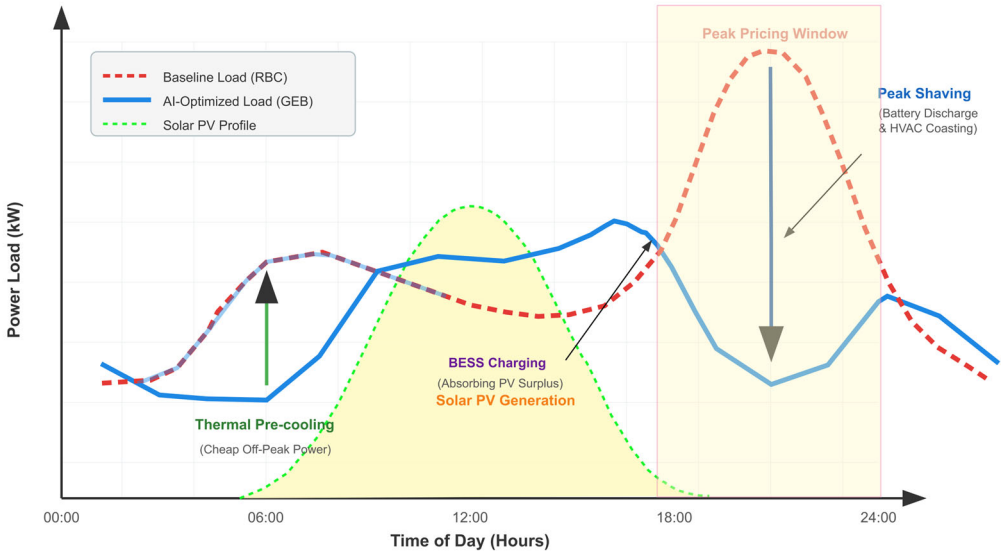


Figure 7. Load shifting and energy arbitrage mechanisms of a Grid-Interactive Efficient Building (GEB). The schematic depicts the modification of a building's daily power demand curve by AI-assisted Model Predictive Control (MPC). The baseline Rule-Based Control (RBC) demand (red dashed line) exhibits a sharp peak during the late-afternoon "Peak Pricing Window" (shaded red region). Conversely, the AI-Optimized GEB load (blue solid line) flattens this peak through three distinct operations: (1) Thermal Pre-cooling during off-peak hours utilizing low-cost power, (2) BESS Charging to absorb local Solar PV Generation (yellow dashed line) during midday, and (3) Peak Shaving (battery discharge and HVAC coasting) to drastically reduce grid reliance when electricity rates and carbon intensity are highest.

The successful implementation of these strategies requires interoperable communication between utility systems and building control systems. Such communication can be enabled through standards like OpenADR [115], which allows for automated communication between grid systems and building energy management systems. Advanced systems can employ supervisory AI agents to interpret these signals, identify flexible loads using semantic building models, and determine possible curtailment options. Such a level of coordination enables a building to participate in demand response programs, and in some cases, even in ancillary service markets. However, participation in such programs is dependent on response latency, measurement fidelity, and so forth [116]. A comprehensive taxonomy of these core operational strategies is provided in Table 4.

Table 4. Core Operational Strategies of Grid-Interactive Efficient Buildings (GEBs).

Operational Strategy	Description	AI/Algorithmic Enabler	Primary Grid & Building Benefit
Load Shedding	Temporary reduction in electricity usage to alleviate acute grid stress during peak periods.	Agentic AI (Semantic identification of non-critical loads).	Prevents grid blackouts; avoids peak demand charges for the facility.

Table 4. *Cont.*

Operational Strategy	Description	AI/Algorithmic Enabler	Primary Grid & Building Benefit
Load Shifting (Peak Shaving)	Moving electricity consumption from high-price/high-carbon periods to off-peak periods.	Transformer forecasting (pricing/weather) + MPC thermal modeling.	Flattens the building's load profile; significant reduction in utility OPEX.
Energy Arbitrage	Storing cheap, off-peak or excess PV energy to use or sell during expensive peak hours.	DRL (optimizing battery charge/discharge constraints).	Maximizes ROI of on-site renewables (PV) and Battery Energy Storage Systems.
Frequency Regulation	Rapid, sub-minute adjustments in power draw/export to maintain grid frequency at 50/60 Hz.	Edge Computing/TinyML (ultra-low latency execution).	Generates new revenue streams via participation in ancillary service markets.

6.3. Predictive Maintenance

The operational continuity of Grid-Interactive Efficient Buildings (GEBs) is not only dependent on supervisory optimization but is also heavily reliant on the reliability of its underlying cyber-physical infrastructure. Traditionally, facility management strategies utilize a variety of reactive (run-to-failure) and time-based preventive maintenance techniques, which can be less than ideal. Fault Detection and Diagnosis (FDD) systems, enabled by AI, are helping to drive the transition towards condition-based maintenance and predictive maintenance. In this context, machine learning algorithms are trained to identify patterns from the data streams collected from various sensors, thereby detecting abnormal patterns that may be indicative of potential faults. Such abnormal patterns could include sensor drift, degradation of machine actuators, or losses in efficiency [117,118].

However, the main difficulty in building strong FDD systems comes from the significant data imbalance that occurs during operations. In HVAC systems, the majority of the time is spent in a healthy operating status, which makes supervised learning methods impossible to apply because of the lack of available fault data. Therefore, the HVAC industry has largely turned to unsupervised deep learning methods, where Autoencoders are used most frequently. An Autoencoder is trained only on available healthy operating data to encode the input data (e.g., temperature, pressure, flow rates, power consumption) into a compact representation that is then decoded to reconstruct the original input data. When an incipient fault occurs, such as a failing motor or micro-leakage of refrigerant gas, the new sensor reading deviates from the learned data manifold [119,120]. The failure of the autoencoder to decode the new input data causes a high reconstruction error, which acts as an alarm for faults. The mechanics of an unsupervised anomaly detection approach are visually detailed in Figure 8. During normal operation, the Autoencoder successfully compresses and reconstructs the multidimensional sensor Input (X) with minimal deviation. However, when an incipient mechanical fault occurs, it introduces patterns unseen during the model's training phase. As depicted by the Δ comparator node in the architecture, the model fails to accurately decode these novel anomalies, leading to a substantial spike in the reconstruction error ($\|X - X'\|^2$). The system subsequently passes this metric to a Threshold Logic block, which automatically flags a system "FAULT" whenever the error exceeds the established tolerance baseline (τ). This system automatically and quickly finds problems just by checking reconstruction errors, completely avoiding the need for labeled examples.

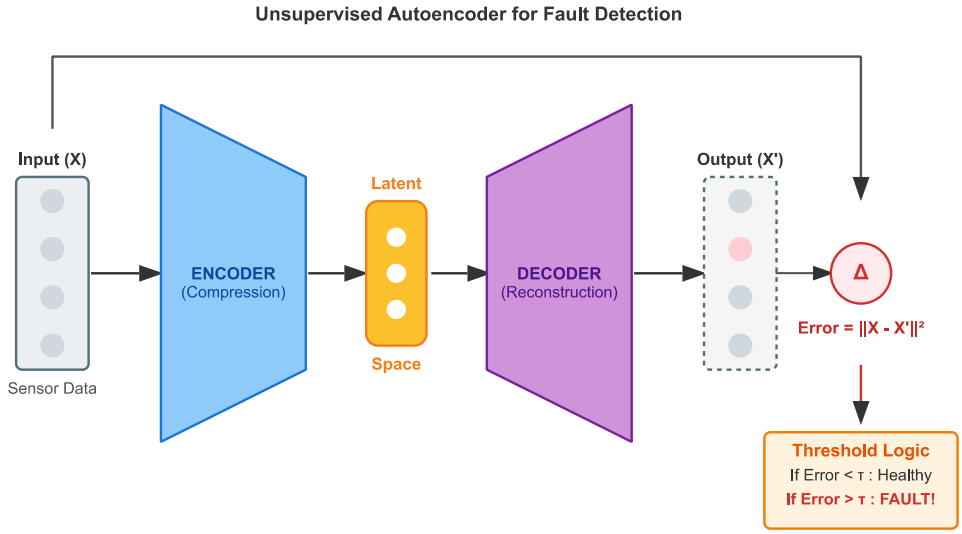


Figure 8. Pipeline for unsupervised Fault Detection and Diagnostics (FDD) using an Autoencoder. The Encoder compresses multi-modal sensor Input (X) into a Latent Space, while the Decoder reconstructs it as Output ($X\prime$). A comparator node (Δ) computes the reconstruction error ($Error = \left|X - X'\right|^2$). Because the network trains exclusively on healthy data, anomalies yield high reconstruction errors. A Threshold Logic module evaluates this error against a tolerance limit (τ), triggering a "FAULT" alert if $Error > \tau$, and maintaining a "Healthy" status otherwise.

While anomaly detection is recognized as a critical first step in FDD, fault isolation is significantly more complex because of the complex thermodynamic and control interdependencies within buildings [121]. For instance, localized faults such as a faulty VAV damper have the ability to spread to both airside and waterside systems, resulting in secondary effects such as changes to supply temperatures or increases to chiller power demand. To improve the level of resolution in diagnosis, modern FDD systems combine data-driven anomaly detection with semantic building representations like Brick Schema. The directed graph representation of sensor data and relationships allows for contextualization of the anomaly score with the building topology. Bayesian Belief Networks and Graph Neural Networks (GNNs) are used to infer the most likely root cause of faults given symptom patterns. While such approaches can improve fault localization accuracy, their performance depends on graph completeness, sensor coverage, and the validity of assumed causal relationships [122,123]. Moreover, bridging the gap between purely data-driven Fault FDD and the PIML framework, as discussed in Section 4.3, is an important research frontier. It is well known that purely data-driven FDD systems often suffer from false alarms due to sensor drift, seasonality of weather, or changes in usage patterns of buildings. However, physics-informed FDD, which uses constraints from the laws of physics, can clearly distinguish between real degradation and false alarms. PINN can be used to monitor real-time degradation of the overall heat transfer coefficient (U-value) of a heat exchanger and raise an alarm only when unequivocal fouling is indicated by the laws of physics [124].

7. Critical Implementation Challenges

While the dominance of the algorithmic paradigms discussed in Section 4 is theoretically assured, their practical application in commercial BEMSs is seen to be severely limited. The transition from idealized environments to stochastic environments with occupied buildings points to a set of critical implementation challenges. This section delineates these gaps, shifting the focus from accuracy to viability, and begins with the paramount concern of data governance.

7.1. Privacy, Security and the Federated Learning Paradigm

The performance of GEBs is closely related to the level of granularity and frequency of operational data. Advanced optimization and control strategies may be dependent on detailed telemetry, such as zone occupancy estimates, thermal preference inputs, appliance disaggregation, and detailed behavioral patterns [107]. However, the use of enhanced data granularity also brings forth significant privacy and data governance concerns. The centralized storage and processing of detailed information from buildings' telemetry data have the potential to compromise sensitive information, such as the schedules of occupancy or the usage patterns of space, thereby increasing the level of cybersecurity and privacy concerns. These issues must be addressed appropriately, adhering to data governance standards such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA) for the EU and US, respectively [125]. Hence, the balancing act between the need for high-resolution data and the need to maintain privacy remains a key issue.

To solve the problem of the conflict between the high demand for data and the constraints of the need to maintain privacy, there is an increasing focus of research and industry solutions on Federated Learning (FL). This is different from traditional machine learning, where aggregated data is stored in a cloud server. In the FL model, individual buildings or edge gateways act as local training clients in the federated learning structure. Each client improves the local model using the data it possesses, which remains locally stored. Once the training iteration is done, the model updates, such as weights or gradients, are sent to the server. The server, in turn, aggregates these updates, usually using Federated Averaging (FedAvg) algorithms, to produce a global model, which is then sent to the nodes. While this approach reduces direct sharing of raw telemetry, privacy guarantees rely on other techniques such as differential privacy and secure aggregation, and system performance might be affected by data diversity and communication constraints [126,127]. To address the stringent data privacy constraints of modern smart infrastructure, the Federated Learning (FL) paradigm replaces legacy centralized data aggregation. Unlike traditional methods that continuously stream vulnerable, raw telemetry to a central cloud, the FL approach completely isolates raw data. Training occurs strictly within the local building edge, and the network only exchanges mathematical weights with a Global Aggregation Server for Federated Averaging (FedAvg). This bi-directional exchange of parameters, rather than sensitive raw data, guarantees robust model optimization while strictly adhering to GDPR and CCPA privacy mandates.

While federated learning (FL) beautifully sidesteps the issue of direct data transmission, its application to the wide range of building portfolios poses a notable algorithmic problem: statistical heterogeneity. In traditional machine learning problems, the data are assumed to be independent and identically distributed (IID). In the built environment, this assumption is violated to a great degree. No two buildings have the same thermodynamic properties, HVAC system configurations, or occupant schedules. Training a global model on such highly non-IID data has been shown to lead to weight divergence, wherein the performance of the overall global model is poor compared to the local models. Thus, the concept of Personalized Federated Learning (pFL) has been introduced. In pFL, the global server learns a general representation of thermodynamics, while the last layer of the neural network is left local to account for the unique characteristics of each building [128,129].

Moreover, while Federated Learning (FL) is considered to safeguard telemetry data from cyber threats, it is still susceptible to attacks from sophisticated adversaries. They may resort to model inversion or gradient leakage attacks, which enable them to derive sensitive occupant information solely based on the weight update sent. To enhance GEB systems, it is critical to improve the FL pipeline with cryptographic mechanisms. The

inclusion of Differential Privacy (DP) mechanisms, which add statistical noise to local gradient information before it is sent, along with Secure Multi-Party Computation (SMPC) mechanisms, ensures that the server is incapable of extracting individual building updates, thus providing mathematical anonymity [130].

7.2. Scalability and Generalization and the Transfer Learning Imperative

The commercial viability of AI-based BEMS is heavily predicated upon scalability. However, it is recognized that there is a high level of domain heterogeneity in the built environment. This means that, unlike computer vision, where the fundamental features of an object are generally invariant across different datasets, the thermal mass of a building, the topology of the HVAC systems, the occupancy patterns, and the microclimate are highly unique [131]. As a result, it has been noticed that predictive models or DRL policies developed on one facility, i.e., the source building, may degrade significantly when applied to another facility, i.e., the target building. This is referred to as a domain shift, which can be mathematically defined by the divergence of input distributions ($P(X_S) \neq P(X_T)$), also known as covariate shift. The need to develop separate models for each facility usually requires large amounts of historical data, which can be computationally expensive and economically impractical [132].

To address the limitations of scalability related to the development of artificial intelligence models within buildings, Transfer Learning (TL) is being increasingly applied within the field of building energy management. TL is based on the understanding that, despite the diversity of mechanical and architectural configurations within buildings, the physical laws of heat transfer, fluid mechanics, and thermodynamics remain constant across buildings. The TL approach also enables the handling of varied state spaces and reward functions for different buildings that may have different energy infrastructures, occupancy levels, and environmental conditions, thus improving generalization for real-world scenarios [133]. In practice, a neural network trained on a source building learns a set of generalized representations in the early stages of the network (e.g., the thermal response to solar irradiance and the delay of the HVAC system). When deployed on a target building, these early stages of the network are "frozen" and the later stages are fine-tuned on a small dataset from the target building. This helps to overcome the cold start problem, speed up the training, and obtain a baseline of operation in a matter of days instead of months or seasons [134]. The concept of Transfer learning mitigates this issue by transferring knowledge from a data-rich source building to a new facility. Through the process of keeping the generalized thermodynamic layers frozen while tuning the building-specific layers with limited local data, transfer learning solves the cold start problem.

The advanced versions of this paradigm leverage Domain Adaptation (DA) techniques to address the statistical discrepancies between the source and target domains. In the case of HVAC predictive maintenance and load forecasting, the discrepancies between sensor calibrations, equipment sizes, and environmental conditions cause significant covariate shifts. The Unsupervised Domain Adaptation (UDA) methods, especially the adversarial training frameworks, leverage the discriminator network to force the feature extractor network to learn features that, by design, are domain-invariant [135]. Furthermore, the integration of Transfer Learning and the semantic ontologies, as described in Section 3, provides the theoretical basis for Zero-Shot Transfer. The AI system will be able to logically deduce the topological constraints of the target building before the receipt of any historical sensor telemetry, thus enabling the safe and conservative control execution from the start. The transition from localized, bespoke models toward more general, transferable architectures represents the fundamental catalyst for the widespread adoption of GEBs. The need for localized training loops forces occupants to endure extended periods of suboptimal

levels of thermal comfort and limits the facility manager's ability to rapidly scale up energy arbitrage strategies across a corporate real estate portfolio. Transfer Learning, as a unifying framework for a portfolio of buildings as a cohesive learning ecosystem, rather than separate data silos, enables AI to transition from a specialized engineering practice to a more general software solution. Research has demonstrated that transfer learning techniques can lower prediction errors by as much as 78% and enable real-time forecasting of energy consumption [136,137].

7.3. The Green AI Paradox

As the built environment progresses towards decarbonization, the algorithmic complexity required for the coordination of GEBs increases exponentially. This, in itself, presents a fundamental yet oft-overlooked paradox: the "Green AI" Paradox, which points to a fundamental irony that the deep learning techniques used to minimize the energy footprint of a building are, in fact, energy-intensive processes themselves. The training of complex models such as load forecasting with Transformers, or the training of complex DRL models for HVAC control, require millions of floating-point operations (FLOPs) that are performed on energy-intensive hardware such as GPUs and TPUs. In fact, if the energy footprint of the AI model's training and querying exceeds the thermodynamic energy savings within the building, then the AI model itself becomes a "Red AI" asset, meaning it is a net-negative contribution to mitigating the effects of climate change [138]. To evaluate the environmental sustainability of AI-based building management systems, it is necessary to consider not only the potential savings in operation-related energy consumption but also the carbon footprint of the AI model development process. The concept of AI Carbon Payback Period (CPP) can be used to describe the time required for the aggregated operation-related carbon savings to compensate for the carbon footprint of the AI model [139]. Figure 9 maps the cumulative net carbon of model deployment over a 24-month operational period. Traditional "Red AI" approaches, which require training complex models from scratch, generate massive initial compute emissions, placing the deployment deep within the Carbon Debt Zone. As the chart demonstrates, it takes approximately 15 months of building energy savings just to offset this initial training footprint. On the other hand, a "Green AI" framework, which leverages transfer learning and edge quantization, drastically reduces the initial compute penalty. By starting much closer to the Net Zero breakeven line, this system crosses into the Carbon Profit Zone in approximately 5 months, proving that AI-driven energy savings can outpace their own computational footprint rapidly when optimized. In quantitative terms, training a standard transformer-based prediction model within a centralized cloud environment may take exaFLOPs of computations and emit hundreds of kilos of CO_2 equivalent. However, when using a quantized TinyML model on an MCU at the edge, the amount of power required will be a mere few milliwatts [138]. The large disparity in computational energy expenditure has significant implications for the Carbon Payback Period [139].

Training large-scale DL models in the cloud data centers can result in significant greenhouse gas emissions, especially when the electricity used to power these centers is drawn from carbon-intensive grids. If the improvements in efficiency achieved through the deployed model are limited, the period required to achieve net carbon neutrality could extend over prolonged operational horizons [138]. On the contrary, techniques that reduce the computational load, such as transfer learning, model compression, and fine-tuning, have the potential to greatly reduce the carbon footprint of the model. This is because these techniques reduce the training time and resources required, thereby reducing the time to payback and increasing the chances of a positive climate impact for AI-based building optimization systems [140]. To reduce the environmental footprint associated with the

deployment of AI in building energy management, more emphasis is being placed on the computational efficiency of Edge AI. In this context, Tiny ML and Model Compression are being explored to enable constrained intelligence in BEMS. Some of these general compression strategies include quantization, which reduces precision (e.g., from 32-bit floating-point to 8-bit integer representations), hence reducing memory usage and energy consumption during inference [141]. In many cases, this reduction in precision results in marginal performance degradation for moderate precision reduction levels. Another popular approach is network pruning, which removes redundant or less important model parameters to produce a sparser model with reduced computational complexity [142]. Moreover, the training process may be fine-tuned to coincide with the computational workload and the macroeconomic grid conditions. Thus, within a paradigm referred to as Carbon Aware Computing, the offline training process for the BEMS neural networks may be coordinated to coincide with periods of high renewable energy production on the grid (such as solar peaks at noon or nocturnal wind production). In such a manner, the GEB principle is reflexively applied to the AI system itself [143].

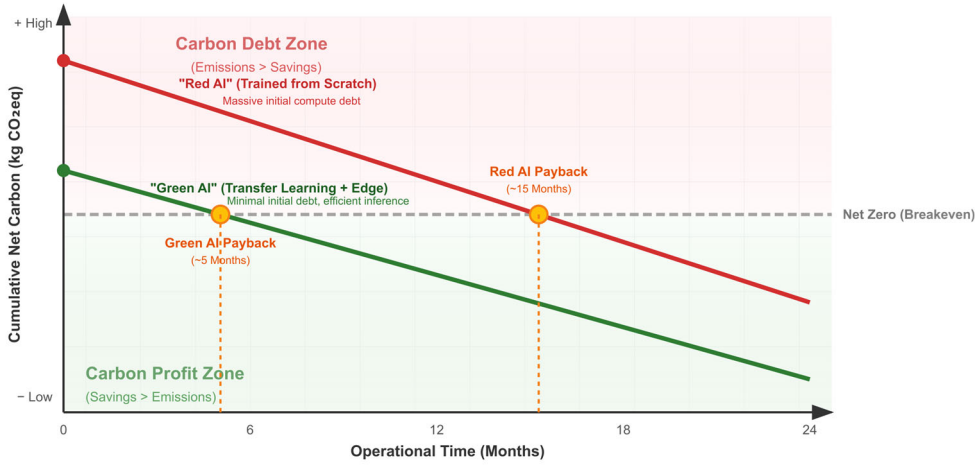


Figure 9. The carbon lifecycle of AI deployments in smart buildings, illustrating the "Green AI Paradox." The *y*-axis tracks Cumulative Net Carbon ($kgCO_2eq$), divided into a Carbon Debt Zone (top) and a Carbon Profit Zone (bottom). "Red AI" (trained from scratch) incurs a massive initial training debt, requiring approximately 15 months of operation to reach the Net Zero breakeven point. In contrast, our proposed "Green AI" approach (utilizing Transfer Learning and edge efficiency) minimizes initial compute debt, achieving environmental profitability in just 5 months.

8. Future Roadmap

While the following paradigms, such as Cognitive Digital Twins (CDTs), Agentic AI, and LLMs, represent the cutting-edge of AI research, they are discussed here as forward-looking concepts rather than proven field deployments. The building domain currently lacks extensive, deployment-level evidence for these technologies. Consequently, the capabilities described in this section should be viewed as hypothesized potential that requires rigorous real-world validation.

8.1. Cognitive Digital Twins

The logical end result of all the integrations described within this manuscript would be the development of a Cognitive Digital Twin (CDT). While Digital Twins are often used as a mere high-fidelity 3D dashboard or static Building Information Model (BIM) updated through unidirectional telemetry streams, the Cognitive Digital Twin represents a paradigm shift in autonomous cyber-physical system synchronization. As it would be imbued with all the capabilities described within Section 3 through its integration with Agentic AI, it would not simply represent or react to the physical state of the building; it

would sense, reason, and control it. It would be the ultimate cognitive bridge between digital computation and physical thermodynamics [144,145]. The key feature of the Cyber-Physical Digital Twin (CDT) is its "real-time, closed-loop, bi-directional" architecture, as indicated in Figure 10. In the forward direction of the physical-to-digital vector, the extensive multi-modal IoT data streams provide continuous calibration to the digital twin, such that the state matrix of the digital twin correctly reflects the physical world in near real-time. In the reverse direction of the digital-to-physical vector, the digital twin uses DRL and structured API function calls to directly transmit the optimal set points to the physical BEMS actuators. This architecture completely eliminates the human-in-the-loop bottleneck in building operations, allowing the building to react instantaneously to changing grid prices, weather, and spatial occupancies.

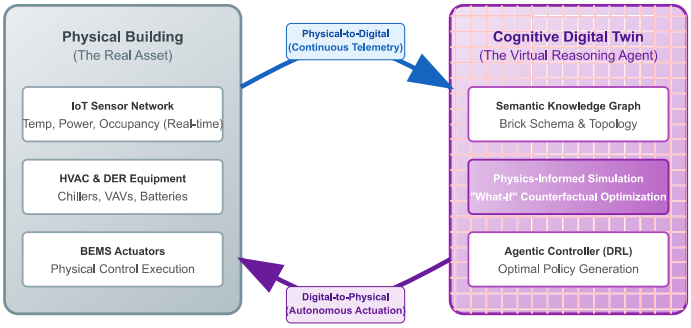


Figure 10. The continuous cyber-physical loop of a Cognitive Digital Twin (CDT). The physical building (left) acts as the real asset, providing real-time state telemetry via an IoT sensor network. In the digital realm (right), a Semantic Knowledge Graph (utilizing Brick Schema) contextualizes the data, feeding into Physics-Informed Simulations for "what-if" counterfactual optimization. Finally, the Agentic Controller utilizes Deep Reinforcement Learning (DRL) to reason about optimal strategies before autonomously pushing control signals back to the physical BEMS actuators.

Outside of real-time orchestration, the CDT functions as a risk-free, high-fidelity "sandbox" for counterfactual "what-if" simulations. By incorporating PINNs to regularize its neural approximations, the CDT automatically satisfies the laws of thermodynamics (Section 4.3). This allows the supervisory AI to safely simulate extreme exogenous events without negatively impacting occupant safety or physical equipment. To illustrate, in preparation for a heatwave event with a demand-response event, the CDT can simulate tens of thousands of pre-cooling and demand-response shedding scenarios in the cloud overnight. It analyzes the Pareto-optimal trade-offs (Section 5.1) in the digital space, determines the mathematical optimum, and sends the winning control policy to the physical building ahead of the event.

Ultimately, the broad adoption of Cognitive Digital Twins will usher in a new era in the building lifecycle, moving the industry from static point-in-time commissioning to Continuous Commissioning (CCx). Buildings have traditionally been observed to undergo rapid performance degradation in the early stages after occupancy due to the natural wear and tear of the mechanical systems. A CDT naturally resists this entropy. As it continuously cross-references the actual performance of the physical building against the optimal performance represented in the digital twin, it automatically identifies the early signs of anomalies (as discussed in the FDD frameworks in Section 5.3) and continuously adapts its control strategies [146]. The evolutionary development from static spatial models to fully autonomous cyber-physical twins is summarized in Table 5.

Table 5. The Evolutionary Roadmap of Building Representations, highlighting the progression of algorithms, autonomy, and the limitations that drive next-generation transitions.

Evolutionary Stage & Framework	Enabling Algorithms	Application Scenario	Data Flow & Autonomy	Identified Limitations (Driving Next Evolution)	References
Generation 1: Building Information Modeling (BIM)	None (Geometry/Physics rendering engines)	3D Spatial Coordination & Clash Detection during design/construction.	None (Static). Autonomy: Zero.	Static representation; completely lacks real-time operational data or temporal dynamics.	[147–149]
Generation 2: Traditional Digital Twin	Rule-based logic, Descriptive Statistics	Real-time monitoring, basic fault alarms, and historical dashboarding.	Unidirectional (Physical → Digital). Autonomy: Low.	Reactive only; lacks predictive capability and relies entirely on human operators to interpret data and act.	[148–152]
Generation 3: Predictive Digital Twin	Supervised ML (e.g., Autoencoders), Time series forecasting (LSTM)	Load forecasting, predictive maintenance, and operational energy simulations.	Unidirectional (Physical → Digital). Autonomy: Medium.	Open-loop control; AI only recommends actions, lacking autonomous physical actuation and semantic reasoning.	[148–153]
Generation 4: Cognitive Digital Twin (CDT)	DRL, Agentic AI, Semantic Knowledge Graphs (Brick/Haystack)	Counterfactual simulation, continuous commissioning, and dynamic microclimate control.	Bi-directional (Real-time synchronization). Autonomy: High.	Requires complex semantic ontology mapping to successfully deploy across diverse physical environments.	[153–157]

8.2. Personalized Comfort Models

For several decades, the macroscopic thermal comfort models of commercial HVAC systems have been based on the thermal comfort theories introduced by P.O. Fanger in the 1970s. As mentioned in Section 5.1, the PMV and PPD methods are, in fact, statistical averages, which should reflect the hypothetical "average" human. However, recent extensive field tests show that the accuracy of the PMV-based methods in predicting the thermal satisfaction of individuals in the field is only 34%. This is due to the fact that the thermal comfort level of humans is a very subjective phenomenon, which depends on various physiological and psychological factors that show great variation between genders, ages, metabolic rates, and acclimatization levels. With the advent of the cognitive buildings concept, the "one-size-fits-all" approach has to be abandoned in favor of Human Centric AI and Personal Comfort Models (PCMs) [158]. In contrast to traditional thermal comfort models such as PMV, which are based on steady-state heat balance theory, the Personal Comfort Model (PCM) takes a data-driven approach to predict thermal comfort based on individual thermal preferences. In PCM methods, ML algorithms are used instead of making assumptions based on physiological principles. It is envisioned that the application of a PCM system would require integration with IoT sensing infrastructure or participatory feedback mechanisms. Input features can be composed of environmental measurements specific to a given zone, such as zone air temperature, relative humidity, and radiant asymmetry, physiological measurements derived from wearable devices, and direct occupant feedback via a mobile application. As a result, it is envisioned that a combination of environmental, biometric, and subjective information can be amalgamated to develop an adaptive model of comfort specific to each occupant. The PCMs based on ML have been found to be consistently outperform the conventional PMV model in accurately predicting

personal comfort levels. Research suggests that PCMs can offer a prediction accuracy up to 74% better than the conventional method, and certain ML models can even achieve accuracy rates greater than 70%, whereas conventional models can achieve accuracy rates up to 34% [159,160]. Although this increased feature space allows for greater personalization, it also presents problems with regard to data quality, privacy, and model generalization, which have to be addressed appropriately [161].

From an algorithmic standpoint, Programmable Climate Management systems often employ supervised learning methods like Random Forest, Support Vector Machines (SVM), or a light-weight multi-layer perceptron with the occupant's discrete micro-feedback as the ground truth label. However, in a more sophisticated cognitive building, this is formalized as a Human-in-the-Loop DRL problem [162] (Figure 11). In this case, the occupant is the final reward signal. If the AI agent adjusts the micro-zone temperature and the occupant indicates a "too cold" condition, the AI agent will receive a negative reward signal, immediately triggering a mathematical update to its policy matrix. To overcome the "cold start" issue for new occupants, Transfer Learning methodologies (Section 6.2) are used, where the system is initialized with a generalized demographic comfort baseline and the upper layers of the neural network are adapted using only a few days of occupant feedback.

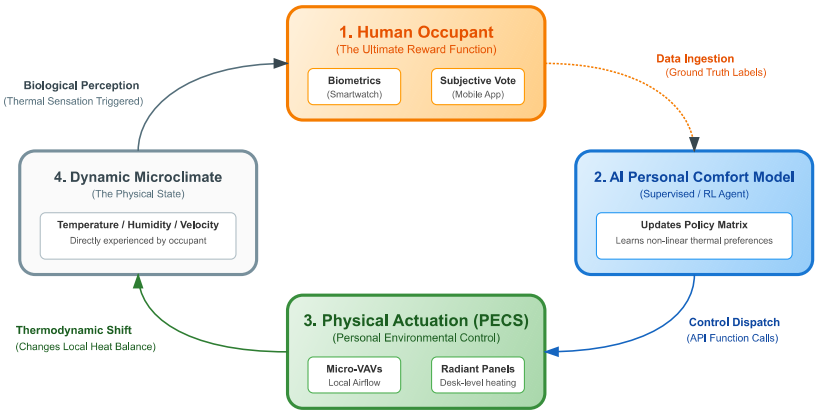


Figure 11. The closed-loop reinforcement learning cycle of Human-Centric AI. The cycle initiates with the Human Occupant providing Biometric and Subjective Vote data, serving as the ultimate ground truth label. This data updates the AI Personal Comfort Model, which learns non-linear preferences and updates its Policy Matrix. Through API dispatch, the model drives Physical Actuation via the Personal Environmental Control (PECS) layer (Micro-VAVs, Radiant Panels). This actuation dynamically alters the Dynamic Microclimate (Temperature/Humidity/Velocity) experienced by the occupant, continuously driving biological perception and closing the loop.

The main barrier in implementing PCMs is not in the algorithmic framework but in physical spatial resolution. A cognitive AI cannot offer personalized thermal comfort if the building only uses monolithic, centralized Air Handling Units (AHUs) to serve large open-plan offices. Therefore, the implementation of PCMs needs to be integrated with the decentralization of physical actuation, i.e., Personalized Environmental Control Systems (PECS). These include smart VAV diffusers, actively cooled desk chairs, and localized radiant space heaters. In shared spaces without PECS, the Agentic AI uses spatial weighting to calculate a multi-occupant consensus by aggregating the intersecting PCMs of the individuals currently in the space using computer vision or Wi-Fi triangulation [163]. Ultimately, Human-Centric AI redescribes the human occupant from an essentially passive receiver of the building's internal climate to an active component of the cyber-physical control system. By focusing on individual human biology rather than adhering to aggregate standards, cognitive buildings can often expand traditional temperature deadbands in

unoccupied or transient spaces to achieve greater energy savings while also increasing human productivity and satisfaction [164].

9. Conclusions

The transition towards decarbonized and intelligent environments marks a crucial engineering imperative of the modern world. Based on a systematic review of 144 studies, this paper synthesizes the current state-of-the-art and extracts concrete, actionable takeaways for the future of intelligent building management. This comprehensive review has mapped the fundamental evolution of Building Energy Management Systems (BEMS) from rigid, rule-based automation towards Generation 4 Cognitive Digital Twins (CDTs). These highly autonomous cyber-physical systems are distinguished by their capabilities for dynamic reasoning, continuous learning, and bi-directional actuation. A key finding of this literature analysis is that the realization of cognitive buildings relies fundamentally on translating raw, multi-modal IoT telemetry into machine-readable knowledge graphs using semantic interoperability schemas, such as the Brick Schema. Operating from this structured semantic base, advanced intelligence cores utilizing predictive Transformer architectures and model-free Deep Reinforcement Learning (DRL) agents have revolutionized the non-convex optimization of building energy. Furthermore, the integration of Agentic AI and Large Language Models (LLMs) via Retrieval-Augmented Generation (RAG) and Chain-of-Thought (CoT) prompting has emerged as a breakthrough paradigm, enabling systems to automatically map high-level semantic objectives to executable API controls.

The practical implications of these AIoT technologies for sustainable building management are profound and immediately actionable. By transitioning into Grid-Interactive Efficient Buildings (GEBs), cognitive environments can leverage Model Predictive Control (MPC) and DRL to execute sophisticated demand response strategies, peak shaving, and energy arbitrage. This dynamic operational capability aligns building energy consumption directly with renewable grid availability, storing energy surpluses and reducing demand during peak pricing periods. Concurrently, these continuous operations are safeguarded by unsupervised Fault Detection and Diagnostics (FDD) pipelines. Utilizing Autoencoders and semantic graphs, these diagnostic systems proactively detect mechanical faults and prevent catastrophic failures, shifting maintenance strategies from reactive to condition-based operational continuity, thereby maximizing energy efficiency and extending the lifespan of mechanical infrastructure.

Despite the impressive capabilities of these data-driven methods in simulated environments, this review identifies several critical research gaps that continue to obstruct widespread commercial deployment. Foremost among these is the “Sim-to-Real” trust gap. DRL agents struggle to transition safely from simulators to physical buildings due to a pervasive lack of high-quality, labeled historical data, coupled with the inherent risks of exploratory actions violating thermal comfort constraints or damaging physical equipment. Furthermore, the heterogeneous nature of legacy BMS protocols severely limits the scalable, plug-and-play deployment of AI models across diverse building topologies. Another significant hurdle is the “Green AI” paradox, wherein the computational carbon debt incurred by training massive deep learning models in centralized cloud data centers threatens to offset the operational energy savings they generate at the building level. Finally, the shift toward bi-directional, human-centric sensing introduces severe vulnerabilities regarding occupant privacy and the cybersecurity of critical building infrastructure, which must be resolved before broad adoption can occur.

To bridge these gaps and realize the ubiquitous adoption of CDTs, future research must pursue several concrete directions. The integration of thermodynamic laws into neural networks through Physics-Informed Machine Learning (PIML), specifically Physics-Informed

Neural Networks (PINNs), is essential to ensure AI control policies are mathematically sound and capable of safe extrapolation outside training distributions. To resolve the Green AI paradox and minimize operational latency, the research community must prioritize Edge intelligence and TinyML, focusing on model quantization and pruning to deploy algorithms directly onto resource-constrained microcontrollers. Addressing privacy and scalability mandates will require advancing distributed AI architectures, particularly Federated Learning (FL) and Transfer Learning (TL), allowing diverse building portfolios to collaboratively train global energy models without exposing raw, sensitive telemetry. Ultimately, future cognitive buildings must transition away from static, macroscopic comfort models toward dynamic Personal Comfort Models (PCMs). This evolution necessitates Human-in-the-Loop DRL systems that continuously couple human physiological and subjective feedback directly into the AI's reward function, supported by decentralized Personalized Environmental Control Systems (PECS). By balancing advanced physical actuators with precise, individualized human needs, the cognitive buildings of the future will achieve unprecedented decarbonization while substantially elevating human productivity, well-being, and comfort. Ultimately, the successful deployment of these paradigms will bridge the gap between theoretical AI models and practical facility management, providing an actionable roadmap for achieving global net-zero emission targets.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/buildings16132569/s1, Table S1: Excluded Full-Text Studies; Table S2: Characteristics of Included Studies (Thematic Summary).

Author Contributions: Y.T. (Yunzhi Tian): Development of the conceptualization of the "Cognitive Building" framework; Methodology design; Writing—original draft (drafting of Section 1: Introduction, Section 4: Cognitive Infrastructure and Section 9: Discussion & Synthesis); Overall Investigation of the research topic; Y.T. (Yuan Tian): Methodology execution (applied the PRISMA protocol in Section 2); Formal analysis (performed the scientometric study); Writing—original draft (prepared Section 3: Synthesized Summary of the Results of Existing Studies, and Section 5: Algorithmic Intelligence Core); Y.J.: Data curation (screening and extraction of literature); Investigation and Validation of the Literature; Writing—review & editing (prepared Section 6: Applications, and Section 7: Challenges); Provision of resources for the studies; V.M.: Funding acquisition; Validation of the proposed future roadmap (Section 8) and technical findings; Investigation through critical review and final manuscript refinement. All authors have read and agreed to the published version of the manuscript.

Funding: 2025 Smart Course Project of Undergraduate Universities in Henan Province, China, "History of Chinese and Foreign Architecture".

Data Availability Statement: The original contributions presented in this study are included in the article/Supplementary Material. Further inquiries can be directed to the corresponding authors.

Conflicts of Interest: Author Yuan Tian was employed by the company Xi'an ChangDa Asset Management Co., Ltd. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AHU	Air Handling Unit
AI	Artificial Intelligence
AIoT	Artificial Intelligence of Things
ANN	Artificial Neural Network
BEMS	Building Energy Management Systems
BESS	Battery Energy Storage Systems
BIM	Building Information Modeling

CCPA	California Consumer Privacy Act
CCx	Continuous Commissioning
CDTs	Cognitive Digital Twins
CNN	Convolutional Neural Network
CoT	Chain-of-Thought
CPP	Carbon Payback Period
CPS	Cyber-Physical System
CSI	Channel State Information
DA	Domain Adaptation
DDPG	Deep Deterministic Policy Gradient
DERs	Distributed Energy Resources
DL	Deep Learning
DNN	Deep Neural Networks
DP	Differential Privacy
DQN	Deep Q-Network
DR	Demand Response
DRIP	Data Rich, Information Poor
DRL	Deep Reinforcement Learning
DT	Decision Tree/Digital Twin
EROI	Energy Return on Investment
FDD	Fault Detection and Diagnosis
FedAvg	Federated Averaging
FL	Federated Learning
FLOPs	Floating-Point Operations
GAN	Generative Adversarial Network
GDPR	General Data Protection Regulation
GEBs	Grid-Interactive Efficient Buildings
GNNs	Graph Neural Networks
HVAC	Heating, Ventilation, and Air Conditioning
IoT	Internet of Things
JSON	JavaScript Object Notation
LLM	Large Language Model
LoRaWAN	Long Range Wide Area Network
LSTM	Long Short-Term Memory
MCUs	Microcontrollers
ML	Machine Learning
MOO	Multi-Objective Optimization
MPC	Model Predictive Control
MQTT	Message Queuing Telemetry Transport
MRT	Mean Radiant Temperature
NLP	Natural Language Processing
NSGA	Non-dominated Sorting Genetic Algorithm
O&M	Operations and Maintenance
OWL	Web Ontology Language
PCMs	Personal Comfort Models
PDEs	Partial Differential Equations
PECS	Personalized Environmental Control Systems
pFL	Personalized Federated Learning
PID	Proportional–Integral–Derivative
PIML	Physics-Informed Machine Learning
PINNs	Physics-Informed Neural Networks
PIR	Passive Infrared
PMV	Predicted Mean Vote
PPD	Predicted Percentage of Dissatisfied

PPO	Proximal Policy Optimization
PRISMA	Preferred Reporting Items for Systematic Reviews and Meta-Analyses
QA	Quality Assessment
RAG	Retrieval-Augmented Generation
RBC	Rule-Based Control
RDF	Resource Description Framework
RL	Reinforcement Learning
RNNs	Recurrent Neural Networks
RQ	Research Question
RSSI	Received Signal Strength Indicator
SAC	Soft Actor–Critic
SAREF	Smart Applications REFerence ontology
SCADA	Supervisory Control and Data Acquisition
SLR	Systematic Literature Review
SMPC	Secure Multi-Party Computation
SVM	Support Vector Machines
TFT	Temporal Fusion Transformer
TinyML	Tiny Machine Learning
TL	Transfer Learning
UDA	Unsupervised Domain Adaptation
VAV	Variable Air Volume
WSNs	Wireless Sensor Networks

References

1. Tirelli, D.; Besana, D. Moving toward Net Zero Carbon Buildings to Face Global Warming: A Narrative Review. *Buildings* **2023**, *13*, 684. [CrossRef]
2. Du, Q.; Wan, Z.; Yang, M.; Wang, X.; Bai, L. Dynamic Integrated Simulation of Carbon Emission Reduction Potential in China's Building Sector. *Sustain. Cities Soc.* **2024**, *116*, 105944. [CrossRef]
3. Li, J.; Zhang, Q.; Etienne, X.L. Optimal Carbon Emission Reduction Path of the Building Sector: Evidence from China. *Sci. Total Environ.* **2024**, *919*, 170553. [CrossRef] [PubMed]
4. Zhou, Y.; Liu, J. Advances in Emerging Digital Technologies for Energy Efficiency and Energy Integration in Smart Cities. *Energy Build.* **2024**, *315*, 114289. [CrossRef]
5. Tang, H.; Wang, S. Energy Flexibility Quantification of Grid-Responsive Buildings: Energy Flexibility Index and Assessment of Their Effectiveness for Applications. *Energy* **2021**, *221*, 119756. [CrossRef]
6. Hannan, M.A.; Faisal, M.; Ker, P.J.; Mun, L.H.; Parvin, K.; Mahlia, T.M.I.; Blaabjerg, F. A Review of Internet of Energy Based Building Energy Management Systems: Issues and Recommendations. *IEEE Access* **2018**, *6*, 38997–39014. [CrossRef]
7. Poyyamozhi, M.; Murugesan, B.; Rajamanickam, N.; Shorfuzzaman, M.; Aboelmagd, Y. IoT—A Promising Solution to Energy Management in Smart Buildings: A Systematic Review, Applications, Barriers, and Future Scope. *Buildings* **2024**, *14*, 3446. [CrossRef]
8. Salama, A.K.; Abdellatif, M.M. AIoT-Based Smart Home Energy Management System. In *Proceedings of the 2022 IEEE Global Conference on Artificial Intelligence and Internet of Things (GCAIoT), Alamein New City, Egypt, 18–21 December 2022*; IEEE: New York, NY, USA, 2022; pp. 177–181.
9. Péan, T.Q.; Salom, J.; Costa-Castelló, R. Review of Control Strategies for Improving the Energy Flexibility Provided by Heat Pump Systems in Buildings. *J. Process Control* **2019**, *74*, 35–49. [CrossRef]
10. Mariano-Hernández, D.; Hernández-Callejo, L.; Zorita-Lamadrid, A.; Duque-Pérez, O.; Santos García, F. A Review of Strategies for Building Energy Management System: Model Predictive Control, Demand Side Management, Optimization, and Fault Detect & Diagnosis. *J. Build. Eng.* **2021**, *33*, 101692. [CrossRef]
11. Ardakanian, O.; Bhattacharya, A.; Culler, D. Non-Intrusive Occupancy Monitoring for Energy Conservation in Commercial Buildings. *Energy Build.* **2018**, *179*, 311–323. [CrossRef]
12. Ali, A.I.; Zorlu Partal, S. Development and Performance Analysis of a ZigBee and LoRa-Based Smart Building Sensor Network. *Front. Energy Res.* **2022**, *10*, 933743. [CrossRef]
13. Abdelkarim, M.; Al-Hubaishi, M. A Novel Fuzzy AHP-TOPSIS Model for IoT Protocol Selection in Smart Air Quality Monitoring Systems. In *Proceedings of the 2025 5th International Conference on Pervasive Computing and Social Networking (ICPCSN), Salem, India, 14–16 May 2025*; IEEE: New York, NY, USA, 2025; pp. 194–201.

14. Bernus, P.; Noran, O. Data Rich—But Information Poor. In *Collaboration in a Data-Rich World*; Camarinha-Matos, L.M., Afsarmanesh, H., Fornasiero, R., Eds.; IFIP Advances in Information and Communication Technology; Springer International Publishing: Cham, Switzerland, 2017; Volume 506, pp. 206–214.
15. Guo, F.; Ham, S.W.; Kim, D.; Moon, H.J. Deep Reinforcement Learning Control for Co-Optimizing Energy Consumption, Thermal Comfort, and Indoor Air Quality in an Office Building. *Appl. Energy* **2025**, *377*, 124467. [CrossRef]
16. Touzani, S.; Prakash, A.K.; Wang, Z.; Agarwal, S.; Pritoni, M.; Kiran, M.; Brown, R.; Granderson, J. Controlling Distributed Energy Resources via Deep Reinforcement Learning for Load Flexibility and Energy Efficiency. *Appl. Energy* **2021**, *304*, 117733. [CrossRef]
17. Yu, L.; Qin, S.; Zhang, M.; Shen, C.; Jiang, T.; Guan, X. A Review of Deep Reinforcement Learning for Smart Building Energy Management. *IEEE Internet Things J.* **2021**, *8*, 12046–12063. [CrossRef]
18. Ali, D.M.T.E.; Motuzienė, V.; Džiugaitė-Tumėnienė, R. AI-Driven Innovations in Building Energy Management Systems: A Review of Potential Applications and Energy Savings. *Energies* **2024**, *17*, 4277. [CrossRef]
19. Manmatharasan, P.; Bitsuamlak, G.; Grolinger, K. AI-Driven Design Optimization for Sustainable Buildings: A Systematic Review. *Energy Build.* **2025**, *332*, 115440. [CrossRef]
20. Halhoul Merabet, G.; Essaaidi, M.; Ben Haddou, M.; Qolomany, B.; Qadir, J.; Anan, M.; Al-Fuqaha, A.; Abid, M.R.; Benhaddou, D. Intelligent Building Control Systems for Thermal Comfort and Energy-Efficiency: A Systematic Review of Artificial Intelligence-Assisted Techniques. *Renew. Sustain. Energy Rev.* **2021**, *144*, 110969. [CrossRef]
21. Farzaneh, H.; Malehmirchegini, L.; Bejan, A.; Afolabi, T.; Mulumba, A.; Daka, P.P. Artificial Intelligence Evolution in Smart Buildings for Energy Efficiency. *Appl. Sci.* **2021**, *11*, 763. [CrossRef]
22. Iluyomade, T.D.; Okwandu, A.C. Smart Buildings and Sustainable Design: Leveraging AI for Energy Optimization in the Built Environment. *Int. J. Sci. Res. Arch.* **2024**, *12*, 2448–2456. [CrossRef]
23. Li, J.; Herdem, M.S.; Nathwani, J.; Wen, J.Z. Methods and Applications for Artificial Intelligence, Big Data, Internet of Things, and Blockchain in Smart Energy Management. *Energy AI* **2023**, *11*, 100208. [CrossRef]
24. Muniandi, B.; Maurya, P.K.; Bhavani, C.; Kulkarni, S.; Yellu, R.R.; Chauhan, N. AI-Driven Energy Management Systems for Smart Buildings. *Power Syst. Technol.* **2024**, *48*, 322–337. [CrossRef]
25. Andrews, J.; Kowsika, M.; Vakil, A.; Li, J. A Motion Induced Passive Infrared (PIR) Sensor for Stationary Human Occupancy Detection. In *Proceedings of the 2020 IEEE/ION Position, Location and Navigation Symposium (PLANS), Portland, OR, USA, 20–23 April 2020*; IEEE: New York, NY, USA, 2020; pp. 1295–1304.
26. Yang, Z.; Zhou, Z.; Liu, Y. From RSSI to CSI: Indoor Localization via Channel Response. *ACM Comput. Surv.* **2013**, *46*, 1–32. [CrossRef]
27. Huang, H.-Y.; Chen, Y.-J. Fusion of Received Signal Strength and Channel State Information for Wireless Human Sensing. In *Proceedings of the 2023 International Conference on Consumer Electronics—Taiwan (ICCE-Taiwan), PingTung, Taiwan, 17–19 July 2023*; IEEE: New York, NY, USA, 2023; pp. 617–618.
28. Natarajan, A.; Krishnasamy, V.; Singh, M. A Machine Learning Approach to Passive Human Motion Detection Using WiFi Measurements From Commodity IoT Devices. *IEEE Trans. Instrum. Meas.* **2023**, *72*, 1–10. [CrossRef]
29. Jung, W.; Jazizadeh, F. Comparative Assessment of HVAC Control Strategies Using Personal Thermal Comfort and Sensitivity Models. *Build. Environ.* **2019**, *158*, 104–119. [CrossRef]
30. Özbey, M.F.; Lotfi, B.; Turhan, C. Estimation of the Mean Radiant Temperature in Office Buildings Using an Artificial Neural Network Developed in a Phyton Environment. *Adv. Build. Energy Res.* **2025**, *19*, 300–315. [CrossRef]
31. Lin, X.; Prabowo, A.; Razzak, I.; Xue, H.; Amos, M.; Behrens, S.; Salim, F.D. A Gap in Time: The Challenge of Processing Heterogeneous IoT Data in Digitalized Buildings. *IEEE Pervasive Comput.* **2025**, *24*, 19–31. [CrossRef]
32. Plageras, A.P.; Psannis, K.E.; Stergiou, C.; Wang, H.; Gupta, B.B. Efficient IoT-Based Sensor BIG Data Collection–Processing and Analysis in Smart Buildings. *Future Gener. Comput. Syst.* **2018**, *82*, 349–357. [CrossRef]
33. Pritoni, M.; Paine, D.; Fierro, G.; Mosiman, C.; Poplawski, M.; Saha, A.; Bender, J.; Granderson, J. Metadata Schemas and Ontologies for Building Energy Applications: A Critical Review and Use Case Analysis. *Energies* **2021**, *14*, 2024. [CrossRef]
34. Delgoshaei, P.; Heidarinejad, M.; Austin, M.A. A Semantic Approach for Building System Operations: Knowledge Representation and Reasoning. *Sustainability* **2022**, *14*, 5810. [CrossRef]
35. Mizutani, I.; Ramanathan, G.; Mayer, S. Semantic Data Integration with DevOps to Support Engineering Process of Intelligent Building Automation Systems. In *Proceedings of the 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*; Association for Computing Machinery: New York, NY, USA, 2021.
36. Balaji, B.; Bhattacharya, A.; Fierro, G.; Gao, J.; Gluck, J.; Hong, D.; Johansen, A.; Koh, J.; Ploennigs, J.; Agarwal, Y.; et al. Brick: Metadata Schema for Portable Smart Building Applications. *Appl. Energy* **2018**, *226*, 1273–1292. [CrossRef]
37. Wang, M.; Lilis, G.N.; Katsigarakis, K.; Mavrokapnidis, D.; Korolija, I.; Rovas, D. Automatic HVAC Topology Generation Using BIM Geometry Checking and Knowledge Graph Technologies. In Proceedings of the 2024 European Conference on Computing in Construction, Chania, Crete, Greece, 15–17 July 14 2024.

38. Guldner, A.; Hoffmann, M.; Lohr, C.; Machhamer, R.; Malburg, L.; Morgen, M.; Rodermund, S.C.; Schäfer, F.; Schaupeter, L.; Schneider, J.; et al. A Framework for AI-Based Self-Adaptive Cyber-Physical Process Systems. *It–Inf. Technol.* **2023**, *65*, 113–128. [CrossRef]
39. Chiosa, R.; Piscitelli, M.S.; Pritoni, M.; Capozzoli, A. A Portable Application Framework for Energy Management and Information Systems (EMIS) Solutions Using Brick Semantic Schema. *Energy Build.* **2024**, *323*, 114802. [CrossRef]
40. Delgoshaei, P.; Heidarinejad, M.; Austin, M.A. Semantic Inference-Based Control Strategies for Building HVAC Systems Using Modelica-Based Physical Models. *Procedia Eng.* **2017**, *205*, 1975–1982. [CrossRef]
41. De Barrena, T.F.; Fernandes, A.; Ferrando, J.L.; García, A.; Landaluce, H.; Angulo, I. Adaptive High Frequency Data Streaming for Soft Real-Time Industrial AI: A Scalable Microservices Based Architecture with Dynamic Downsampling. *Array* **2025**, *27*, 100488. [CrossRef]
42. Lin, J.; Zhu, L.; Chen, W.-M.; Wang, W.-C.; Han, S. Tiny Machine Learning: Progress and Futures [Feature]. *IEEE Circuits Syst. Mag.* **2023**, *23*, 8–34. [CrossRef]
43. Ravaglia, L.; Rusci, M.; Nadalini, D.; Capotondi, A.; Conti, F.; Benini, L. A TinyML Platform for On-Device Continual Learning With Quantized Latent Replays. *IEEE J. Emerg. Sel. Top. Circuits Syst.* **2021**, *11*, 789–802. [CrossRef]
44. Heydari, S.; Mahmoud, Q.H. Tiny Machine Learning and On-Device Inference: A Survey of Applications, Challenges, and Future Directions. *Sensors* **2025**, *25*, 3191. [CrossRef] [PubMed]
45. Dash, S. Green AI: Enhancing Sustainability and Energy Efficiency in AI-Integrated Enterprise Systems. *IEEE Access* **2025**, *13*, 21216–21228. [CrossRef]
46. Alharbi, H.A.; Aldossary, M. Energy-Efficient Edge-Fog-Cloud Architecture for IoT-Based Smart Agriculture Environment. *IEEE Access* **2021**, *9*, 110480–110492. [CrossRef]
47. Romanssini, M.; De Aguirre, P.C.C.; Girardi, A.G. IoT Sensor Node with Edge Computing for Efficient Vibration Monitoring of Industrial Motors. In *Proceedings of the 2025 IEEE 16th Latin America Symposium on Circuits and Systems (LASCAS), Bento Gonçalves, Brazil, 25–28 February 2025*; IEEE: New York, NY, USA, 2025; pp. 1–5.
48. Yuan, J.; Xiao, H.; Shen, Z.; Zhang, T.; Jin, J. ELECT: Energy-Efficient Intelligent Edge–Cloud Collaboration for Remote IoT Services. *Future Gener. Comput. Syst.* **2023**, *147*, 179–194. [CrossRef]
49. Madhumathi, C.S.; Ravi, R. Energy-Efficient IoT Sensor Networks Using LoRaWAN and Edge Intelligence. *Int. Innov. Res. J. Eng. Technol.* **2025**, *10*, 12–21. [CrossRef]
50. Dutta, L.; Bharali, S. TinyML Meets IoT: A Comprehensive Survey. *Internet Things* **2021**, *16*, 100461. [CrossRef]
51. Oprea, S.-V.; Bâra, A. An Edge-Fog-Cloud Computing Architecture for IoT and Smart Metering Data. *Peer-To-Peer Netw. Appl.* **2023**, *16*, 818–845. [CrossRef]
52. Daraghmi, Y.-A.; Daraghmi, E.Y.; Daraghma, R.; Fouchal, H.; Ayaida, M. Edge–Fog–Cloud Computing Hierarchy for Improving Performance and Security of NB-IoT-Based Health Monitoring Systems. *Sensors* **2022**, *22*, 8646. [CrossRef] [PubMed]
53. Van Houdt, G.; Mosquera, C.; Nápoles, G. A Review on the Long Short-Term Memory Model. *Artif. Intell. Rev.* **2020**, *53*, 5929–5955. [CrossRef]
54. Dai, Z.-Q.; Li, J.; Cao, Y.-J.; Zhang, Y.-X. SALSTM: Segmented Self-Attention Long Short-Term Memory for Long-Term Forecasting. *J. Supercomput.* **2025**, *81*, 115. [CrossRef]
55. Moveh, S.; Merchán-Cruz, E.A.; Abuhussain, M.; Dodo, Y.A.; Alhumaid, S.; Alhamami, A.H. Deep Learning Framework Using Transformer Networks for Multi Building Energy Consumption Prediction in Smart Cities. *Energies* **2025**, *18*, 1468. [CrossRef]
56. Zheng, P.; Zhou, H.; Liu, J.; Nakanishi, Y. Interpretable Building Energy Consumption Forecasting Using Spectral Clustering Algorithm and Temporal Fusion Transformers Architecture. *Appl. Energy* **2023**, *349*, 121607. [CrossRef]
57. Li, L.; Su, X.; Bi, X.; Lu, Y.; Sun, X. A Novel Transformer-Based Network Forecasting Method for Building Cooling Loads. *Energy Build.* **2023**, *296*, 113409. [CrossRef]
58. López Santos, M.; García-Santiago, X.; Echevarría Camarero, F.; Blázquez Gil, G.; Carrasco Ortega, P. Application of Temporal Fusion Transformer for Day-Ahead PV Power Forecasting. *Energies* **2022**, *15*, 5232. [CrossRef]
59. Lim, B.; Arık, S.Ö.; Loeff, N.; Pfister, T. Temporal Fusion Transformers for Interpretable Multi-Horizon Time Series Forecasting. *Int. J. Forecast.* **2021**, *37*, 1748–1764. [CrossRef]
60. Zhang, S.; Kong, C. Deep Learning and Multi-Objective Optimization for Real-Time Occupancy-Based Energy Control in Smart Buildings. *Sci. Rep.* **2025**, *15*, 39466. [CrossRef] [PubMed]
61. Ibrhim, H.; Khattab, S.; Elsayed, K.; Badr, A.; Nabil, E. A Formal Methods-Based Rule Verification Framework for End-User Programming in Campus Building Automation Systems. *Build. Environ.* **2020**, *181*, 106983. [CrossRef]
62. Shakya, A.K.; Pillai, G.; Chakrabarty, S. Reinforcement Learning Algorithms: A Brief Survey. *Expert Syst. Appl.* **2023**, *231*, 120495. [CrossRef]
63. Banerjee, C.; Chen, Z.; Noman, N. Improved Soft Actor-Critic: Mixing Prioritized Off-Policy Samples With On-Policy Experiences. *IEEE Trans. Neural Netw. Learn. Syst.* **2024**, *35*, 3121–3129. [CrossRef] [PubMed]

64. Sun, H.; Hu, Y.; Luo, J.; Guo, Q.; Zhao, J. Enhancing HVAC Control Systems Using a Steady Soft Actor–Critic Deep Reinforcement Learning Approach. *Buildings* **2025**, *15*, 644. [CrossRef]
65. Adibhesami, M.A.; Hassanzadeh, A. Optimizing HVAC Energy Efficiency in Low-Energy Buildings: A Comparative Analysis of Reinforcement Learning Control Strategies under Tehran Climate Conditions. *Data-Centric Eng.* **2025**, *6*, e40. [CrossRef]
66. Wang, H.; Chen, X.; Vital, N.; Duffy, E.; Razi, A. Energy Optimization for HVAC Systems in Multi-VAV Open Offices: A Deep Reinforcement Learning Approach. *Appl. Energy* **2024**, *356*, 122354. [CrossRef]
67. Yu, L.; Su, X.; Zhang, J. A Deep Reinforcement Learning Framework for Multi-Interval HVAC Control. In *Proceedings of the 2025 IEEE 22nd International Conference on Mobile Ad-Hoc and Smart Systems (MASS), Chicago, IL, USA, 6–8 October 2025*; IEEE: New York, NY, USA, 2025; pp. 567–573.
68. Jiang, Z.; Risbeck, M.J.; Ramamurti, V.; Murugesan, S.; Amores, J.; Zhang, C.; Lee, Y.M.; Drees, K.H. Building HVAC Control with Reinforcement Learning for Reduction of Energy Cost and Demand Charge. *Energy Build.* **2021**, *239*, 110833. [CrossRef]
69. Ding, X.; Cerpa, A.; Du, W. Multi-Zone HVAC Control With Model-Based Deep Reinforcement Learning. *IEEE Trans. Autom. Sci. Eng.* **2025**, *22*, 4408–4426. [CrossRef]
70. Du, Y.; Zandi, H.; Kotevska, O.; Kurte, K.; Munk, J.; Amasyali, K.; Mckee, E.; Li, F. Intelligent Multi-Zone Residential HVAC Control Strategy Based on Deep Reinforcement Learning. *Appl. Energy* **2021**, *281*, 116117. [CrossRef]
71. Biemann, M.; Scheller, F.; Liu, X.; Huang, L. Experimental Evaluation of Model-Free Reinforcement Learning Algorithms for Continuous HVAC Control. *Appl. Energy* **2021**, *298*, 117164. [CrossRef]
72. Silvestri, A.; Coraci, D.; Brandi, S.; Capozzoli, A.; Borkowski, E.; Köhler, J.; Wu, D.; Zeilinger, M.N.; Schlueter, A. Real Building Implementation of a Deep Reinforcement Learning Controller to Enhance Energy Efficiency and Indoor Temperature Control. *Appl. Energy* **2024**, *368*, 123447. [CrossRef]
73. Boutahri, Y.; Tilioua, A. Reinforcement Learning for HVAC Control and Energy Efficiency in Residential Buildings with BOPTEST Simulations and Real-Case Validation. *Discov. Comput.* **2025**, *28*, 45. [CrossRef]
74. Zhang, H.; Seal, S.; Wu, D.; Boulet, B.; Bouffard, F.; Joos, G. Data-Driven Model Predictive and Reinforcement Learning Based Control for Building Energy Management: A Survey. *arXiv* **2021**, arXiv:2106.14450.
75. Wang, X.; Mahdavi, N.; Aryai, V.; Sethuvenkatraman, S.; West, S.R. From Simulation to Reality: A Study of Reinforcement Learning Control in Operational Building Environments. *ASME J. Eng. Sustain. Build. Cities* **2025**, *6*, 041006. [CrossRef]
76. Gu, S.; Yang, L.; Du, Y.; Chen, G.; Walter, F.; Wang, J.; Knoll, A. A Review of Safe Reinforcement Learning: Methods, Theories, and Applications. *IEEE Trans. Pattern Anal. Mach. Intell.* **2024**, *46*, 11216–11235. [CrossRef] [PubMed]
77. Chowa, S.S.; Alvi, R.; Rahman, S.S.; Rahman, M.A.; Raiaan, M.A.K.; Islam, M.R.; Hussain, M.; Azam, S. From Language to Action: A Review of Large Language Models as Autonomous Agents and Tool Users. *Artif. Intell. Rev.* **2026**, *59*, 71. [CrossRef]
78. Güitta-López, L.; Suriani, V.; Boal, J.; López-López, Á.J.; Nardi, D. Boosting Deep Reinforcement Learning with Semantic Knowledge for Robotic Manipulators. *Robotics* **2025**, *14*, 86. [CrossRef]
79. Brohi, S.; Mastoi, Q.; Jhanjhi, N.Z.; Pillai, T.R. A Research Landscape of Agentic AI and Large Language Models: Applications, Challenges and Future Directions. *Algorithms* **2025**, *18*, 499. [CrossRef]
80. Basu, K. Bridging Knowledge Gaps in LLMs via Function Calls. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, Boise, ID USA, 21–25 October 2024*; ACM: New York, NY, USA, 2024; pp. 5556–5557.
81. Zhang, L.; Ford, V.; Chen, Z.; Chen, J. Automatic Building Energy Model Development and Debugging Using Large Language Models Agentic Workflow. *Energy Build.* **2025**, *327*, 115116. [CrossRef]
82. Karlović, R.; Lorencin, I. Large Language Models as Retail Cart Assistants: A Prompt-Based Evaluation. *AHFE Open Access* **2025**, *198*, 529. [CrossRef]
83. Miao, J.; Thongprayoon, C.; Suppadungsuk, S.; Krisanapan, P.; Radhakrishnan, Y.; Cheungpasitporn, W. Chain of Thought Utilization in Large Language Models and Application in Nephrology. *Medicina* **2024**, *60*, 148. [CrossRef] [PubMed]
84. Turpin, M.; Michael, J.; Perez, E.; Bowman, S. Language Models Don' t Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting. In *Proceedings of the Advances in Neural Information Processing Systems*; Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2023; Volume 36, pp. 74952–74965.
85. Xu, X.; Zhang, D.; Liu, Q.; Lu, Q.; Zhu, L. Agentic RAG with Human-in-the-Retrieval. In *Proceedings of the 2025 IEEE 22nd International Conference on Software Architecture Companion (ICSA-C), Odense, Denmark, 31 March 2025–4 April 2025*; IEEE: New York, NY, USA, 2025; pp. 498–502.
86. Arroyo, J.; Manna, C.; Spiessens, F.; Helsen, L. Reinforced Model Predictive Control (RL-MPC) for Building Energy Management. *Appl. Energy* **2022**, *309*, 118346. [CrossRef]
87. Van De Berg, D.; Shah, N.; Del Rio-Chanona, E.A. Hierarchical Planning-Scheduling-Control—Optimality Surrogates and Derivative-Free Optimization. *Comput. Chem. Eng.* **2024**, *188*, 108726. [CrossRef]

88. Wang, Z.; Zhao, Y.; Zhang, C.; Ma, P.; Liu, X. A General Multi Agent-Based Distributed Framework for Optimal Control of Building HVAC Systems. *J. Build. Eng.* **2022**, *52*, 104498. [CrossRef]
89. Zhu, M.; Zhang, H.; Jiao, A.; Karniadakis, G.E.; Lu, L. Reliable Extrapolation of Deep Neural Operators Informed by Physics or Sparse Observations. *Comput. Methods Appl. Mech. Eng.* **2023**, *412*, 116064. [CrossRef]
90. Meng, C.; Griesemer, S.; Cao, D.; Seo, S.; Liu, Y. When Physics Meets Machine Learning: A Survey of Physics-Informed Machine Learning. *Mach. Learn. Comput. Sci. Eng.* **2025**, *1*, 20. [CrossRef]
91. Pan, L.; Li, G.; Zhu, T.; Liu, D.; Wang, Y.; Lu, Y. Physics-Informed Machine Learning in Design and Manufacturing: Status and Challenges. *J. Comput. Inf. Sci. Eng.* **2025**, *25*, 120804. [CrossRef]
92. Mowlavi, S.; Nabi, S. Optimal Control of PDEs Using Physics-Informed Neural Networks. *J. Comput. Phys.* **2023**, *473*, 111731. [CrossRef]
93. Plankovskyy, S.; Tsegelnyk, Y.; Shyshko, N.; Litvinchev, I.; Romanova, T.; Velarde Cantú, J.M. Review of Physics-Informed Neural Networks: Challenges in Loss Function Design and Geometric Integration. *Mathematics* **2025**, *13*, 3289. [CrossRef]
94. Jiang, Z.; Wang, X.; Dong, B. Physics-Informed Modularized Neural Network for Advanced Building Control by Deep Reinforcement Learning. *Adv. Appl. Energy* **2025**, *19*, 100237. [CrossRef]
95. Karniadakis, G.E.; Kevrekidis, I.G.; Lu, L.; Perdikaris, P.; Wang, S.; Yang, L. Physics-Informed Machine Learning. *Nat. Rev. Phys.* **2021**, *3*, 422–440. [CrossRef]
96. Chen, W.; Wang, Q.; Hesthaven, J.S.; Zhang, C. Physics-Informed Machine Learning for Reduced-Order Modeling of Nonlinear Problems. *J. Comput. Phys.* **2021**, *446*, 110666. [CrossRef]
97. Chen, Y.; Yang, Q.; Chen, Z.; Yan, C.; Zeng, S.; Dai, M. Physics-Informed Neural Networks for Building Thermal Modeling and Demand Response Control. *Build. Environ.* **2023**, *234*, 110149. [CrossRef]
98. Arowoiya, V.A.; Onososen, A.O.; Moehler, R.C.; Fang, Y. Influence of Thermal Comfort on Energy Consumption for Building Occupants: The Current State of the Art. *Buildings* **2024**, *14*, 1310. [CrossRef]
99. Yan, B.; Yang, W.; He, F.; Zeng, W. Occupant Behavior Impact in Buildings and the Artificial Intelligence-Based Techniques and Data-Driven Approach Solutions. *Renew. Sustain. Energy Rev.* **2023**, *184*, 113372. [CrossRef]
100. Franco, A.; Bartoli, C.; Conti, P.; Miserocchi, L.; Testi, D. Multi-Objective Optimization of HVAC Operation for Balancing Energy Use and Occupant Comfort in Educational Buildings. *Energies* **2021**, *14*, 2847. [CrossRef]
101. Tartarini, F.; Schiavon, S. Comparative Analysis of PMV Models Accuracy Implemented in the ISO 7730:2005 and ASHRAE 55:2023. *Build. Environ.* **2025**, *275*, 112766. [CrossRef]
102. Izhar; Wang, X.; Xu, W.; Tavakkoli, H.; Yuen, Z.; Shan, X.; Lee, Y.-K. Integrated Predicted Mean Vote Sensing System Using MEMS Multi-Sensors for Smart HVAC Systems. *IEEE Sens. J.* **2021**, *21*, 8400–8410. [CrossRef]
103. Omidvar, A.; Kim, J. Modification of Sweat Evaporative Heat Loss in the PMV/PPD Model to Improve Thermal Comfort Prediction in Warm Climates. *Build. Environ.* **2020**, *176*, 106868. [CrossRef]
104. Li, H.; Xu, J.; Zhao, Q.; Wang, S. Economic Model Predictive Control in Buildings Based on Piecewise Linear Approximation of Predicted Mean Vote Index. *IEEE Trans. Autom. Sci. Eng.* **2024**, *21*, 3384–3395. [CrossRef]
105. Cigler, J.; Prívara, S.; Váňa, Z.; Žáčeková, E.; Ferkl, L. Optimization of Predicted Mean Vote Index within Model Predictive Control Framework: Computationally Tractable Solution. *Energy Build.* **2012**, *52*, 39–49. [CrossRef]
106. Yang, X.; Chen, Z.; Huang, X.; Li, R.; Xu, S.; Yang, C. Robust Capacity Optimization Methods for Integrated Energy Systems Considering Demand Response and Thermal Comfort. *Energy* **2021**, *221*, 119727. [CrossRef]
107. Bayasgalan, A.; Park, Y.S.; Koh, S.B.; Son, S.-Y. Comprehensive Review of Building Energy Management Models: Grid-Interactive Efficient Building Perspective. *Energies* **2024**, *17*, 4794. [CrossRef]
108. Satchwell, A.; Piette, M.; Khandekar, A.; Granderson, J.; Frick, N.; Hledik, R.; Faruqui, A.; Lam, L.; Ross, S.; Cohen, J.; et al. *A National Roadmap for Grid-Interactive Efficient Buildings*; LBL Publications: Los Angeles, CA, USA, 2021; pp. 1–150. [CrossRef] [PubMed]
109. Silva, B.N.; Khan, M.; Han, K. Futuristic Sustainable Energy Management in Smart Environments: A Review of Peak Load Shaving and Demand Response Strategies, Challenges, and Opportunities. *Sustainability* **2020**, *12*, 5561. [CrossRef]
110. Vázquez-Canteli, J.R.; Nagy, Z. Reinforcement Learning for Demand Response: A Review of Algorithms and Modeling Techniques. *Appl. Energy* **2019**, *235*, 1072–1089. [CrossRef]
111. Ghafoori, M.; Abdallah, M.; Kim, S. Electricity Peak Shaving for Commercial Buildings Using Machine Learning and Vehicle to Building (V2B) System. *Appl. Energy* **2023**, *340*, 121052. [CrossRef]
112. Hossain, J.; Saeed, N.; Manojkumar, R.; Marzband, M.; Sedraoui, K.; Al-Turki, Y. Optimal Peak-Shaving for Dynamic Demand Response in Smart Malaysian Commercial Buildings Utilizing an Efficient PV-BES System. *Sustain. Cities Soc.* **2024**, *101*, 105107. [CrossRef]
113. Saleem, M.I.; Saha, S.; Izhar, U.; Ang, L. Optimized Energy Management of a Solar Battery Microgrid: An Economic Approach towards Voltage Stability. *J. Energy Storage* **2024**, *90*, 111876. [CrossRef]

114. Gilabert-Torres, C.; Rus-Casas, C.; Jiménez-Castillo, G.; Muñoz-Rodríguez, F.J. Photovoltaic Systems with Battery Storage: A Novel and Comprehensive Scheduling Method for High-Consumption Facilities. *J. Energy Storage* **2025**, *139*, 118858. [CrossRef]
115. Fernandez-Izquierdo, A.; Cimmino, A.; Patsonakis, C.; Tsolakis, A.C.; Garcia-Castro, R.; Ioannidis, D.; Tzovaras, D. OpenADR Ontology: Semantic Enrichment of Demand Response Strategies in Smart Grids. In *Proceedings of the 2020 International Conference on Smart Energy Systems and Technologies (SEST), Istanbul, Turkey, 7–9 September 2020*; IEEE: New York, NY, USA, 2020; pp. 1–6.
116. Ribó-Pérez, D.; Larrosa-López, L.; Pecondón-Tricas, D.; Alcázar-Ortega, M. A Critical Review of Demand Response Products as Resource for Ancillary Services: International Experience and Policy Recommendations. *Energies* **2021**, *14*, 846. [CrossRef]
117. Chen, Z.; O'Neill, Z.; Wen, J.; Pradhan, O.; Yang, T.; Lu, X.; Lin, G.; Miyata, S.; Lee, S.; Shen, C.; et al. A Review of Data-Driven Fault Detection and Diagnostics for Building HVAC Systems. *Appl. Energy* **2023**, *339*, 121030. [CrossRef]
118. Gholipour, Y.; Zare, M.; Sereshk, M.V.; Gholipour, Y. A Comprehensive Review of Maintenance Strategies: From Reactive to Proactive Approaches. *Cent. ASIA Cauc.* **2025**, *26*, 70–83. [CrossRef]
119. Li, T.; Sun, C.; Yan, R.; Chen, X. A Novel Unsupervised Graph Wavelet Autoencoder for Mechanical System Fault Detection. *J. Intell. Manuf.* **2025**, *36*, 5397–5414. [CrossRef]
120. Holly, S.; Heel, R.; Katic, D.; Schoeffl, L.; Stiftinger, A.; Holzner, P.; Kaufmann, T.; Haslhofer, B.; Schall, D.; Heitzinger, C.; et al. Autoencoder Based Anomaly Detection and Explained Fault Localization in Industrial Cooling Systems. *PHM Soc. Eur. Conf.* **2022**, *7*, 200–210. [CrossRef]
121. Shahnazari, H.; Mhaskar, P.; House, J.M.; Salsbury, T.I. Heating, Ventilation and Air Conditioning Systems: Fault Detection and Isolation and Safe Parking. *Comput. Chem. Eng.* **2018**, *108*, 139–151. [CrossRef]
122. Cai, B.; Huang, L.; Xie, M. Bayesian Networks in Fault Diagnosis. *IEEE Trans. Ind. Inform.* **2017**, *13*, 2227–2240. [CrossRef]
123. Chen, Y.; Wen, J.; Pradhan, O.; Lo, L.J.; Wu, T. Using Discrete Bayesian Networks for Diagnosing and Isolating Cross-Level Faults in HVAC Systems. *Appl. Energy* **2022**, *327*, 120050. [CrossRef]
124. Lai, C.; Baraldi, P.; Zio, E. Physics-Informed Deep Autoencoder for Fault Detection in New-Design Systems. *Mech. Syst. Signal Process.* **2024**, *215*, 111420. [CrossRef]
125. Bakar, A.A.; Yussof, S.; Ghapar, A.A.; Sameon, S.S.; Jørgensen, B.N. A Review of Privacy Concerns in Energy-Efficient Smart Buildings: Risks, Rights, and Regulations. *Energies* **2024**, *17*, 977. [CrossRef]
126. Karimireddy, S.P.; Kale, S.; Mohri, M.; Reddi, S.; Stich, S.; Suresh, A.T. SCAFFOLD: Stochastic Controlled Averaging for Federated Learning. In Proceedings of the 37th International Conference on Machine Learning, Online, 13–18 July 2020; Volume 119, pp. 5132–5143.
127. Sravanthi, K.; Valikar, V.; Kowsalya, S.S.N.; Ghouse, S.M.; Mutharasan, A.; Vajravelu, A. Federated Learning-Based Energy Management Systems for Privacy-Preserving Demand Forecasting in Smart Cities. In *Proceedings of the 2025 13th International Conference on Smart Grid (icSmartGrid), Glasgow, UK, 27–29 May 2025*; IEEE: New York, NY, USA, 2025; pp. 446–451.
128. Rizwan, A.; Khan, A.N.; Ahmad, R.; Khan, Q.W.; Kim, D.H. Personalized Hierarchical Heterogeneous Federated Learning for Thermal Comfort Prediction in Smart Buildings. *Eng. Appl. Artif. Intell.* **2025**, *139*, 109464. [CrossRef]
129. Varno, F.; Saghayi, M.; Rafiee Sevyeri, L.; Gupta, S.; Matwin, S.; Havaei, M. AdaBest: Minimizing Client Drift in Federated Learning via Adaptive Bias Estimation. In *Computer Vision—ECCV 2022*; Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T., Eds.; Lecture Notes in Computer Science; Springer Nature: Cham, Switzerland, 2022; Volume 13683, pp. 710–726.
130. Fereidooni, H.; Marchal, S.; Miettinen, M.; Mirhoseini, A.; Mollering, H.; Nguyen, T.D.; Rieger, P.; Sadeghi, A.-R.; Schneider, T.; Yalame, H.; et al. SAFELearn: Secure Aggregation for Private FEderated Learning. In *Proceedings of the 2021 IEEE Security and Privacy Workshops (SPW), San Francisco, CA, USA, 27 May 2021*; IEEE: New York, NY, USA, 2021; pp. 56–62.
131. Zangato, T.; Osmani, A.; Alizadeh, P. Data-Driven Policy Mapping for Safe RL-Based Energy Management Systems. *Energy Rep.* **2025**, *13*, 1888–1909. [CrossRef]
132. Bayram, F.; Ahmed, B.S. A Domain-Region Based Evaluation of ML Performance Robustness to Covariate Shift. *Neural Comput. Appl.* **2023**, *35*, 17555–17577. [CrossRef]
133. Coraci, D.; Brandi, S.; Hong, T.; Capozzoli, A. Online Transfer Learning Strategy for Enhancing the Scalability and Deployment of Deep Reinforcement Learning Control in Smart Buildings. *Appl. Energy* **2023**, *333*, 120598. [CrossRef]
134. Pinto, G.; Wang, Z.; Roy, A.; Hong, T.; Capozzoli, A. Transfer Learning for Smart Buildings: A Critical Review of Algorithms, Applications, and Future Perspectives. *Adv. Appl. Energy* **2022**, *5*, 100084. [CrossRef]
135. Lei, Y.; Fan, C.; He, H.; Xie, Y. Enabling Efficient Cross-Building HVAC Fault Inferences through Novel Unsupervised Domain Adaptation Methods. *Build. Environ.* **2025**, *272*, 112678. [CrossRef]
136. Park, C.; Kim, I.; Kim, W. Transfer Learning–Based Energy Consumption Prediction for Variable Refrigerant Flow System in Buildings. *Appl. Therm. Eng.* **2025**, *267*, 125811. [CrossRef]
137. Xing, Z.; Pan, Y.; Yang, Y.; Yuan, X.; Liang, Y.; Huang, Z. Transfer Learning Integrating Similarity Analysis for Short-Term and Long-Term Building Energy Consumption Prediction. *Appl. Energy* **2024**, *365*, 123276. [CrossRef]

138. Barbierato, E.; Gatti, A. Toward Green AI: A Methodological Survey of the Scientific Literature. *IEEE Access* **2024**, *12*, 23989–24013. [CrossRef]
139. Delanoë, P.; Tchuente, D.; Colin, G. Method and Evaluations of the Effective Gain of Artificial Intelligence Models for Reducing CO2 Emissions. *J. Environ. Manag.* **2023**, *331*, 117261. [CrossRef] [PubMed]
140. Paula, E.; Soni, J.; Upadhyay, H.; Lagos, L. Comparative Analysis of Model Compression Techniques for Achieving Carbon Efficient AI. *Sci. Rep.* **2025**, *15*, 23461. [CrossRef] [PubMed]
141. Upadhyay, G.; Prasad, C.S.; Akash, V.; Singh, A.K.; Monish Sai, R.; Ghosal, S.; Kar, S.; Shantala, S.H.; Thakur LalitKrushna, J.; Srividhya, S.; et al. Enhancing CNN Inference Time and Reducing Latency on Edge and Resource-Constrained Systems through Quantization. In Proceedings of the 2024 Second International Conference on Networks, Multimedia and Information Technology (NMITCON), Bengaluru, India, 9–10 August 2024; pp. 1–6.
142. Li, J.; Shao, H.; Deng, X.; Jiang, Y. Efficient Filter Pruning: Reducing Model Complexity through Redundancy Graph Decomposition. *Neurocomputing* **2024**, *599*, 128108. [CrossRef]
143. Singh, G. Carbon-Aware Resource Allocation: Dynamically Balancing Compute Loads with Renewable Energy Availability. *Eur. J. Comput. Sci. Inf. Technol.* **2025**, *13*, 30–42. [CrossRef]
144. Lee, H.; Yang, H. Cognitive Digital Twin for Multiobjective Production Scheduling. *J. Manuf. Sci. Eng.* **2025**, *147*, 091004. [CrossRef]
145. Ashtari Talkhestani, B.; Jung, T.; Lindemann, B.; Sahlab, N.; Jazdi, N.; Schloegl, W.; Weyrich, M. An Architecture of an Intelligent Digital Twin in a Cyber-Physical Production System. *Automatisierungstechnik* **2019**, *67*, 762–782. [CrossRef]
146. El Mokhtari, K.; Panushev, I.; McArthur, J.J. Development of a Cognitive Digital Twin for Building Management and Operations. *Front. Built Environ.* **2022**, *8*, 856873. [CrossRef]
147. Revolti, A.; Gualtieri, L.; Pauwels, P.; Dallasega, P. From Building Information Modeling to Construction Digital Twin: A Conceptual Framework. *Prod. Manuf. Res.* **2024**, *12*, 2387679. [CrossRef]
148. Deng, M.; Menassa, C.C.; Kamat, V.R. From BIM to Digital Twins: A Systematic Review of the Evolution of Intelligent Building Representations in the AEC-FM Industry. *J. Inf. Technol. Constr.* **2021**, *26*, 58–83. [CrossRef]
149. Baghalzadeh Shishehgarkhaneh, M.; Keivani, A.; Moehler, R.C.; Jelodari, N.; Roshdi Laleh, S. Internet of Things (IoT), Building Information Modeling (BIM), and Digital Twin (DT) in Construction Industry: A Review, Bibliometric, and Network Analysis. *Buildings* **2022**, *12*, 1503. [CrossRef]
150. Banfi, F.; Brumana, R.; Salvalai, G.; Previtali, M. Digital Twin and Cloud BIM-XR Platform Development: From Scan-to-BIM-to-DT Process to a 4D Multi-User Live App to Improve Building Comfort, Efficiency and Costs. *Energies* **2022**, *15*, 4497. [CrossRef]
151. Nguyen, T.D.; Adhikari, S. The Role of BIM in Integrating Digital Twin in Building Construction: A Literature Review. *Sustainability* **2023**, *15*, 10462. [CrossRef]
152. Pan, Y.; Zhang, L. A BIM-Data Mining Integrated Digital Twin Framework for Advanced Project Management. *Autom. Constr.* **2021**, *124*, 103564. [CrossRef]
153. Sepasgozar, S.; Khan, A.; Smith, K.; Romero, J.; Shen, X.; Shirowzhan, S.; Li, H.; Tahmasebinia, F. BIM and Digital Twin for Developing Convergence Technologies as Future of Digital Construction. *Buildings* **2023**, *13*, 441. [CrossRef]
154. Sepasgozar, S.M.E. Differentiating Digital Twin from Digital Shadow: Elucidating a Paradigm Shift to Expedite a Smart, Sustainable Built Environment. *Buildings* **2021**, *11*, 151. [CrossRef]
155. Boje, C.; Guerriero, A.; Kubicki, S.; Rezgui, Y. Towards a Semantic Construction Digital Twin: Directions for Future Research. *Autom. Constr.* **2020**, *114*, 103179. [CrossRef]
156. Zheng, X.; Lu, J.; Kiritsis, D. The Emergence of Cognitive Digital Twin: Vision, Challenges and Opportunities. *Int. J. Prod. Res.* **2022**, *60*, 7610–7632. [CrossRef]
157. Meschini, S.; Pellegrini, L.; Locatelli, M.; Accardo, D.; Tagliabue, L.C.; Di Giuda, G.M.; Avena, M. Toward Cognitive Digital Twins Using a BIM-GIS Asset Management System for a Diffused University. *Front. Built Environ.* **2022**, *8*, 959475. [CrossRef]
158. Martins Gnecco, V.; Pigliautile, I.; Pisello, A.L. Personal Comfort Models in Long-Term Monitoring Using Physiological Data from Wearable Sensors. *J. Phys. Conf. Ser.* **2023**, *2600*, 092005. [CrossRef]
159. Qavidel Fard, Z.; Zomorodian, Z.S.; Korsavi, S.S. Application of Machine Learning in Thermal Comfort Studies: A Review of Methods, Performance and Challenges. *Energy Build.* **2022**, *256*, 111771. [CrossRef]
160. Haghirad, M.; Heidari, S.; Hosseini, H. Advancing Personal Thermal Comfort Prediction: A Data-Driven Framework Integrating Environmental and Occupant Dynamics Using Machine Learning. *Build. Environ.* **2024**, *262*, 111799. [CrossRef]
161. Feng, Y.; Liu, S.; Wang, J.; Yang, J.; Jao, Y.-L.; Wang, N. Data-Driven Personal Thermal Comfort Prediction: A Literature Review. *Renew. Sustain. Energy Rev.* **2022**, *161*, 112357. [CrossRef]
162. Wu, J.; Huang, Z.; Hu, Z.; Lv, C. Toward Human-in-the-Loop AI: Enhancing Deep Reinforcement Learning via Real-Time Human Guidance for Autonomous Driving. *Engineering* **2023**, *21*, 75–91. [CrossRef]

163. Khovalyg, D.; Bivolarova, M.P.; Shinoda, J.; Al-Assaad, D.; Vellei, M.; Bandurski, K.; Chinazzo, G.; Kazanci, O.B.; Kim, J.; Kramer, T.; et al. Personalized Environmental Control Systems (PECS): Systematic Review of Benefits for Thermal Comfort, Air Quality, Health, and Human Performance. *Build. Environ.* **2025**, *286*, 113541. [CrossRef]
164. Kim, H.; Kang, H.; Choi, H.; Jung, D.; Hong, T. Human-Building Interaction for Indoor Environmental Control: Evolution of Technology and Future Prospects. *Autom. Constr.* **2023**, *152*, 104938. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.