applied sciences

Review

Intelligent HVAC Control in Residential Buildings: A Review of Advanced Techniques and AI Applications

Ricardo Felez and Jesus Felez *

Mechanical Engineering Department, Universidad Politécnica de Madrid, 28006 Madrid, Spain; ricardo.felez.val@alumnos.upm.es
* Correspondence: jesus.felez@upm.es

Featured Application

This review provides actionable insights for engineers, researchers, and building managers seeking to implement intelligent HVAC control in residential settings. Applications include energy-efficient heating and cooling, integration with renewable energy sources, and adaptive comfort management using AI-driven strategies such as MPC, DRL, and neural network-based prediction. The findings support practical deployment of hybrid control systems to reduce energy consumption, enhance occupant comfort, and enable grid-interactive smart homes.

Abstract

Increasing energy demand, decarbonization commitments, and growing expectations for thermal comfort are driving the need for more adaptive and efficient climate control in residential buildings. This review synthesizes contemporary intelligent HVAC control strategies, including model-predictive control (MPC), deep reinforcement learning (DRL), data-driven forecasting, and hybrid approaches. Following PRISMA guidelines, a set of studies published between 2010 and 2025 was systematically screened and analyzed to identify the dominant methodological trends, data requirements, implementation architectures, and evaluation practices reported in the literature. This review highlights how these methods differ in modeling assumptions, computational complexity, robustness to uncertainty, and suitability for residential environments characterized by stochastic occupancy and heterogeneous building stock. In addition, we examine enabling technologies such as sensing infrastructures, pricing signals, and embedded computation, as well as barriers to real-world deployment, including data availability, interpretability, and integration with existing building systems. The findings provide a consolidated framework for understanding the capabilities and limitations of intelligent HVAC control and outline research gaps that remain for achieving scalable, user-centered, and energy-efficient operation in residential buildings.

Keywords: intelligent HVAC systems; model predictive control (MPC); deep reinforcement learning (DRL); neural networks; residential energy efficiency; thermal comfort optimization; smart building automation; occupancy-aware control; hybrid control strategies; AI-driven HVAC

Academic Editors: Franco Cotana, Veronica Lucia Castaldo and Ilaria Pigliautile

Received: 5 January 2026
Revised: 1 February 2026
Accepted: 8 February 2026
Published: 18 February 2026

1. Introduction

Residential heating is a major contributor to greenhouse gas emissions in Europe, where cold winters and high comfort standards drive substantial energy demand. Across the continent, most homes rely on fossil fuel combustion for heating, with natural gas, oil, and coal accounting for most of the energy used in household heating systems. This widespread dependence on fossil fuels makes residential heating one of the leading sources of carbon dioxide (CO_2) and other greenhouse gas emissions, directly affecting air quality and accelerating climate change. Recent estimates indicate that the residential sector accounts for nearly 30% of Europe's total energy consumption, with heating representing the largest share of this demand. Residential heating alone produces approximately 12% of the European Union's total CO_2 emissions [1], a figure that varies with seasonal temperatures and energy sources. Countries with colder climates, particularly in Northern and Eastern Europe, often exhibit even higher emissions due to extended heating periods and lower average temperatures.

In the United States, buildings account for roughly 39% of all primary energy consumption and 74% of electricity use [2]. Thermal end uses—including space conditioning, water heating, and refrigeration—represent about 50% of building energy demand and are projected to increase in the coming years. This growing demand underscores the urgent need for innovative solutions to reduce energy consumption and mitigate environmental impacts in residential buildings worldwide. On the other hand, studies estimate that replacing fossil fuel heating systems with heat pumps, which include both aerothermal and geothermal technologies, could reduce household emissions by up to 70–80%. By 2022, approximately 20 million heat pumps had been installed across Europe [3], helping to displace emissions-intensive natural gas and oil heating systems. According to the European Heat Pump Association (EHPA), these installations are already saving about 50 million tons of CO_2 annually [4]. As part of the European Green Deal, the EU aims to accelerate this transition, with projections suggesting that heat pumps could supply up to 40% of Europe's residential heating demand by 2030 [5]. This large-scale shift is critical for achieving the EU's net-zero emissions target by 2050, as heat pumps not only reduce direct emissions but also operate more efficiently when paired with renewable electricity sources.

In this context, intelligent HVAC (Heating, Ventilation, and Air Conditioning) systems for residential buildings have become a cornerstone of modern energy-efficient architecture, driven by the growing need for sustainability, occupant comfort, and operational optimization. These systems integrate advanced control techniques—such as model predictive control, fuzzy logic, and adaptive algorithms—to dynamically regulate temperature, humidity, and air quality while minimizing energy consumption. Their applications extend beyond basic climate control, encompassing smart zoning, demand-response strategies, and integration with renewable energy sources, which collectively enhance system resilience and cost-effectiveness. Artificial intelligence methods, including machine learning and neural networks, play a pivotal role in predictive maintenance, fault detection, and real-time optimization, enabling HVAC systems to learn from historical data and adapt to changing environmental conditions. By leveraging these technologies, intelligent HVAC solutions not only reduce carbon footprints but also improve user experience through personalized comfort settings, marking a significant advancement in residential building automation.

A key challenge in residential HVAC control is the uncertainty associated with occupant behavior. Unlike commercial buildings with more predictable schedules, residential occupancy is highly variable and shaped by irregular routines and diverse comfort preferences. This variability affects internal heat gains and comfort demands, making occupancy one of the most influential and unpredictable disturbances in HVAC operation. In this review, robust handling of occupancy uncertainty refers to the ability of a control strategy to maintain comfort and efficiency despite fluctuations in presence, timing, and activity patterns. Recognizing this challenge is essential for evaluating how MPC, DRL, and hybrid approaches perform under real residential conditions.

In addition to highlighting the growing importance of energy efficiency and thermal comfort in residential buildings, this review establishes a clear framework for comparing advanced HVAC control strategies. Specifically, the analysis focuses on three key characteristics: (i) flexibility, defined as the ability of a control method to adapt to dynamic operating conditions, user behavior, and weather variability; (ii) predictability, referring to the accuracy and reliability of performance forecasting under diverse scenarios; and (iii) computational complexity, which encompasses the algorithmic sophistication and resource requirements for real-time implementation. These criteria provide a consistent basis for evaluating the strengths and limitations of model-predictive control, deep reinforcement learning, and neural network-based approaches throughout this review.

This review synthesizes findings from 97 sources examining intelligent HVAC systems for residential buildings, with emphasis on control techniques, applications, and artificial intelligence methods. The included studies span simulation-based investigations, field trials, laboratory experiments, and reviews, covering diverse geographic locations and building types.

The primary objective of this review is to provide a comprehensive synthesis of advanced HVAC control strategies for residential buildings, with a particular emphasis on AI-driven approaches such as Model Predictive Control, Deep Reinforcement Learning, and neural network-based methods. By analyzing recent developments, comparing performance metrics, and identifying research gaps, this paper aims to support the design and implementation of intelligent, energy-efficient, and occupant-centric HVAC systems.

The remainder of this paper is organized as follows: Section 2 presents the methodology adopted for this review, including selection criteria and classification framework. Section 3 provides an overview of conventional HVAC control strategies and their limitations. Section 4 discusses advanced control techniques, focusing on AI-based methods and hybrid approaches. Section 5 examines practical applications, integration with IoT and smart home technologies, and case studies from recent literature. Section 6 identifies key challenges, research gaps, and future directions for intelligent HVAC systems. Finally, Section 7 concludes the paper by summarizing the main findings and implications for residential building energy management.

2. Methods

This review was conducted following the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines [6]. The methodology was designed to ensure transparency, reproducibility, and rigor in identifying, selecting, and synthesizing relevant studies on intelligent HVAC control strategies for residential buildings.

2.1. Search Strategy

A comprehensive literature search was performed across major scientific databases, including Scopus, Web of Science, IEEE Xplore, and ScienceDirect, covering publications from January 2010 to December 2025. The search terms combined keywords related to HVAC systems, intelligent control, and artificial intelligence, using Boolean operators to refine results. The primary query included: ("HVAC" OR "Heating Ventilation Air Conditioning") AND ("intelligent control" OR "smart control" OR "AI" OR "machine learning" OR "model predictive control" OR "reinforcement learning") AND ("residential" OR "smart home"). Additional sources were identified through backward and forward citation tracking to ensure comprehensive coverage.

To determine whether each study met the inclusion criteria, we applied a two-stage screening process following PRISMA guidelines. First, titles and abstracts were screened, and then full-text articles were assessed for eligibility. Two independent reviewers, the paper authors, conducted the screening of all records and full-text reports. Disagreements were resolved through discussion and, when necessary, by consulting a third external reviewer to reach consensus.

2.2. Inclusion and Exclusion Criteria

Studies were selected based on the following inclusion criteria:

(i). Focus on residential buildings or mixed-use buildings with residential components;
(ii). Address intelligent HVAC control strategies, including AI-based methods (e.g., MPC, DRL, neural networks);
(iii). Present quantitative or qualitative performance metrics, such as energy savings, thermal comfort, or computational efficiency;
(iv). Published in peer-reviewed journals or conference proceedings between 2010 and 2025.
(v). Written in English.

The exclusion criteria were as follows:

1. Studies focusing exclusively on commercial or industrial buildings;
2. Papers addressing hardware design without control strategy analysis;
3. Articles lacking sufficient methodological detail or performance evaluation;
4. Non-peer-reviewed sources (e.g., blogs, white papers), conference abstracts, patents, and database records without a DOI (Digital Object Identifier).

A standardized checklist was used to evaluate inclusion criteria, which focused on the following: (i) residential building context, (ii) intelligent HVAC control strategies (including AI-based methods), and (iii) reporting of performance metrics. Each criterion was weighted equally (0.33), and studies scoring above 0.7 were included in the final synthesis. Automation tools were employed to assist in data extraction and classification. Specifically, the Elicit AI [7] assistant was used to extract predefined data fields from eligible studies, following detailed instructions for each column (e.g., building context, control method, performance results). However, all final decisions regarding inclusion were made by human reviewers, the two papers' authors, to ensure accuracy and consistency.

2.3. Study Selection Process

The initial search yielded 1429 records. Before screening, all retrieved records were imported into a reference management tool (Zotero) and automatically checked for duplicates using the software's built-in detection algorithm. Potential duplicates were then manually verified by the reviewers to ensure accuracy. After this process, only unique records were retained for title and abstract screening.

After removing duplicates, 375 studies remained for screening. Titles and abstracts were reviewed to exclude irrelevant papers, resulting in 163 papers for full-text assessment. Following the application of inclusion/exclusion criteria, 97 studies were retained for final synthesis. The selection process is illustrated in the PRISMA flow diagram (Figure 1).

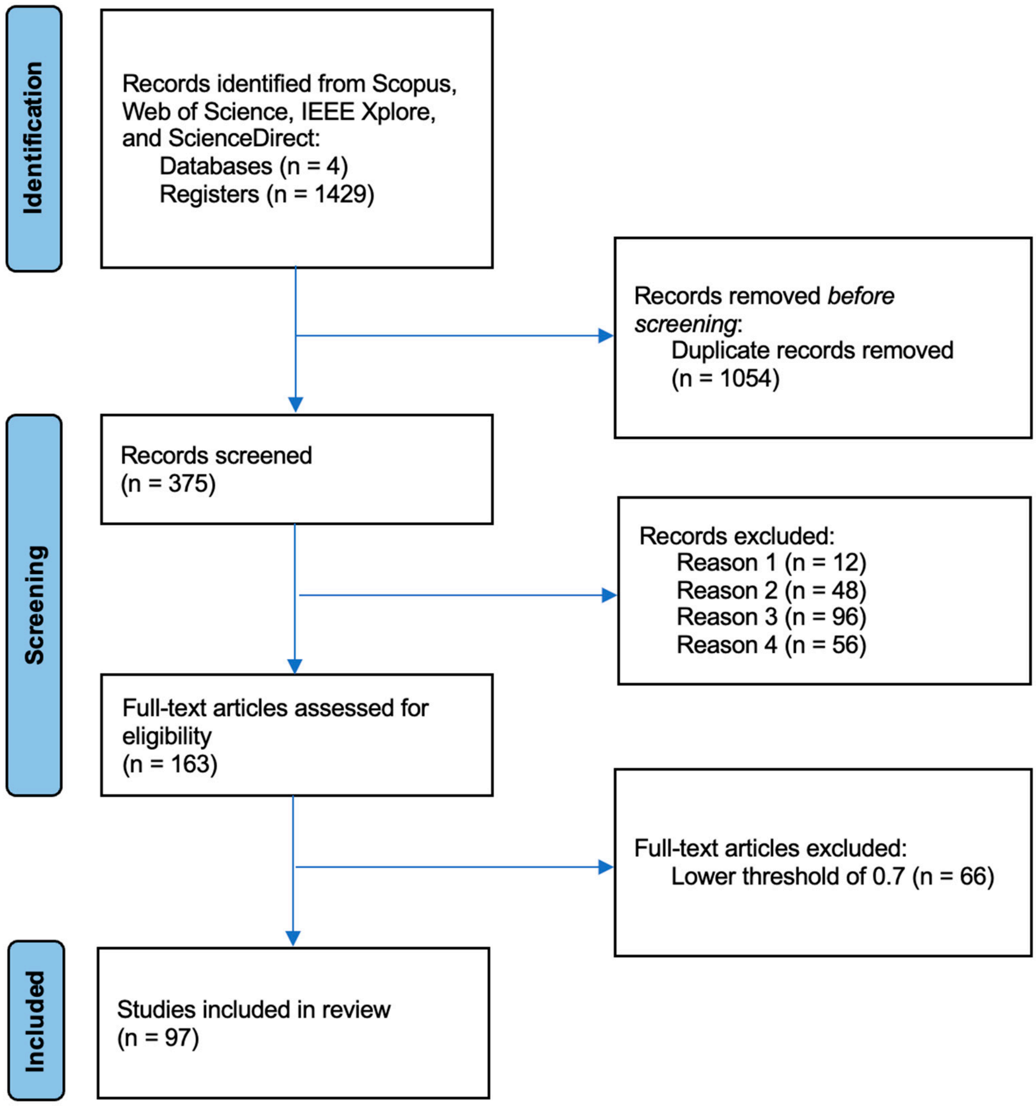

Figure 1. PRISMA Flow Diagram.

2.4. Data Extraction and Classification

Data from the selected studies were extracted using a standardized template, capturing the following:

- Control technique (e.g., MPC, DRL, fuzzy logic, hybrid approaches);
- Application context (e.g., thermal comfort optimization, energy efficiency, fault detection);
- Performance indicators (e.g., energy savings %, comfort index, computational cost).
- Integration aspects (e.g., IoT, renewable energy sources).

The studies were classified into three main categories:

1. Conventional control strategies and their limitations;
2. Advanced AI-based methods for HVAC optimization;
3. Hybrid and integrated approaches combining predictive and adaptive algorithms.

Based on the Elicit assistant [7], a large language model was used to extract each data column listed below from each paper. The model was provided with the extraction instructions shown for each column. This language model was employed to extract each data column listed below from the papers. Extraction instructions for each column were provided to the model as shown in Table 1.

Table 1. Screening criteria and instructions for screening.

Screening Criterion	Instructions for Screening
Building Context:	Extract details about the residential building setting, including the following: • Type of residential building (single-family house, apartment, villa, etc.) • Building size/characteristics (number of units, floors, square footage if provided) • Geographic location and climate zone • Building age or construction type if relevant • Any specific building features that affect HVAC performance
Intelligent Control Method:	Identify the specific intelligent control technique(s) used, including the following: • Type of control approach (MPC, rule-based, fuzzy logic, neural networks, etc.) • Specific Artificial Intelligence/machine Learning (AI/ML) algorithms employed: Long Short-Term Memory (LSTM), reinforcement learning (RL), genetic algorithms, etc.) • Control framework or methodology name • Whether it is predictive, reactive, or adaptive control • Integration with existing HVAC systems vs. new installation
HVAC Application:	Extract what HVAC functions and systems are being controlled, including the following: • Specific HVAC components controlled (air handling units, thermostats, dampers, fans, etc.) • Functions optimized (heating, cooling, ventilation, air quality, humidity) • Type of HVAC system (central air, heat pumps, mini-splits, mixed-mode, etc.) • Scope of control (individual rooms, zones, whole building) • Integration with other building systems (lighting, occupancy detection, etc.)
System Architecture:	Describe the technical implementation, including the following: • Hardware components (sensors, controllers, communication devices) • Software platform or framework used • Data collection and processing approach • Communication protocols and connectivity • Real-time vs. offline processing • Computational requirements and deployment location (edge, cloud, local)
Input Data:	Extract what data the intelligent system requires, including the following: • Environmental data (temperature, humidity, outdoor weather, etc.) • Occupancy information (presence, count, patterns, schedules) • Energy consumption data and utility pricing • Building physics data (thermal models, heat transfer coefficients) • Historical performance data • User preferences or comfort settings
Performance Results:	Extract quantitative and qualitative outcomes, including the following: • Energy savings (percentage reduction, kWh saved, cost savings) • Comfort improvements (temperature control accuracy, user satisfaction) • Operational benefits (reduced runtime, maintenance insights) • Environmental impact (carbon footprint reduction) • Comparison to baseline or conventional control methods • Statistical significance and confidence intervals were provided
Study Design:	Identify the research methodology, including the following: • Type of study (field trial, simulation, laboratory experiment, case study) • Duration of evaluation period • Sample size (number of buildings, households, or units) • Experimental design (before/after, controlled trial, comparative study) • Data collection methods and measurement tools • Validation approach for results

Table 1. *Cont.*

Screening Criterion	Instructions for Screening
Implementation Challenges:	Extract any barriers, limitations, or practical considerations, including the following: • Technical challenges during deployment • Cost considerations (capital, installation, operational) • User acceptance and behavioral factors • Scalability issues • Maintenance and reliability concerns • Regulatory or code compliance issues • Recommendations for future implementation

The risk of bias in the included studies was evaluated using a standardized approach to ensure methodological rigor. The initial assessment was conducted with the assistance of the Elicit AI tool, which extracted relevant information and applied the Joanna Briggs Institute (JBI) Critical Appraisal Checklist for experimental and quasi-experimental designs. This checklist examines key domains such as selection bias, measurement bias, and reporting bias. Each item was rated as "Yes," "No," or "Maybe," and an overall judgment of low, moderate, or high risk of bias was assigned based on the aggregate score.

All outputs generated by Elicit were independently reviewed and validated by the two authors of this paper to ensure accuracy and consistency. Discrepancies were resolved through discussion until consensus was reached. No fully automated decisions were made; human oversight was maintained throughout the process. Given the diversity of study designs, building contexts, and reported performance metrics, a meta-analysis was not feasible. Instead, we employed a narrative synthesis approach to systematically organize and interpret the findings. Studies were grouped according to the following:

(i). Type of intelligent HVAC control strategy;
(ii). Application and Key Features;
(iii). Reported outcomes (energy savings, thermal comfort, cost reduction).

Effect measures for each outcome were standardized where possible. For energy performance, results were expressed as a percentage reduction in energy consumption relative to baseline systems. Thermal comfort was summarized using mean deviation from target temperature (°C) and percentage of time within comfort range, while economic outcomes were reported as mean difference in operational cost. When studies provided multiple indicators, these were normalized and presented as relative improvement (%) to facilitate cross-study comparison.

Heterogeneity was addressed qualitatively by comparing trends across subgroups rather than pooling data statistically. We examined variations in performance based on climate zone, building type, and control strategy complexity. Subgroup analyses were conducted narratively to identify patterns and contextual factors influencing outcomes. Studies assessed as having a high risk of bias were included in the synthesis but flagged in tables and figures. Their influence on overall conclusions was minimized by prioritizing evidence from studies rated as low or moderate risk of bias. No formal sensitivity analysis was performed due to the absence of quantitative pooling.

Findings were summarized in comparative tables. This visualization aimed to highlight performance trends and identify gaps in the current literature. The included studies demonstrate substantial methodological diversity. Simulation-based studies constitute the majority, with several field trials providing real-world validation. The geographic distribution spans North America, Europe, Asia, Australia, and the Middle East, representing diverse climatic conditions from tropical to severe cold regions. Building types predominantly include single-family houses and apartments, though several studies examined commercial or institutional buildings for comparative purposes.

3. Overview of Conventional HVAC Control Strategies and Limitations

Conventional HVAC control strategies have historically relied on simple, rule-based mechanisms and proportional-integral-derivative (PID) controllers to regulate indoor temperature and air quality. These approaches are widely implemented due to their simplicity, low computational requirements, and ease of integration with existing systems [8]. Rule-based control typically operates on predefined setpoints and schedules, adjusting heating or cooling output based on fixed thresholds. Similarly, PID controllers modulate system performance by minimizing the error between the desired and actual temperature through proportional, integral, and derivative actions [9].

While these methods have proven effective in maintaining basic thermal comfort, they exhibit significant limitations in dynamic and complex environments. One major drawback is their inability to adapt to changing conditions, such as fluctuating occupancy, variable outdoor temperatures, and intermittent renewable energy supply [10]. Because conventional controllers rely on static parameters, they cannot anticipate future states or optimize performance under uncertainty. This often results in energy inefficiencies, frequent cycling of equipment, and suboptimal comfort levels.

Another limitation lies in the lack of integration with modern building automation systems and IoT technologies. Traditional controllers are generally designed for isolated operation, without considering interactions between HVAC systems and other building components such as lighting, shading, or energy storage [11]. Consequently, opportunities for demand-response strategies and grid-interactive optimization remain largely untapped.

Furthermore, conventional control strategies do not incorporate predictive capabilities or learning mechanisms. They are reactive by nature, responding only after deviations occur, which can lead to delayed adjustments and increased energy consumption during peak loads [12]. In contrast, advanced approaches such as Model Predictive Control (MPC) and machine learning-based algorithms can forecast future conditions and proactively adjust system parameters, significantly improving efficiency and occupant comfort.

In summary, while conventional HVAC control strategies remain prevalent due to their simplicity and cost-effectiveness, their inherent limitations—lack of adaptability, predictive capability, and integration—underscore the need for intelligent, AI-driven solutions. These shortcomings have motivated extensive research into advanced control techniques, which are discussed in detail in the following section.

4. Advanced Control Techniques for Intelligent HVAC Systems

The limitations of conventional HVAC control strategies have motivated the development of advanced techniques that leverage predictive algorithms, optimization methods, and artificial intelligence (AI) to improve energy efficiency and occupant comfort in residential buildings. These approaches enable HVAC systems to anticipate future conditions, adapt to dynamic environments, and optimize performance under uncertainty. The main families identified in this review are Model Predictive Control (MPC), Deep Reinforcement Learning (DRL), and neural network–based algorithms, often combined into hybrid and hierarchical frameworks that exploit complementary strengths, as well as fuzzy logic and rule-based intelligent control [12–16].

This section synthesizes the main findings from the reviewed studies regarding these advanced control families, focusing on their typical performance, implementation characteristics, and suitability for residential applications characterized by stochastic occupancy, heterogeneous building stock, and emerging grid-interactive requirements. To keep the main text focused on the most relevant comparative insights, the detailed descriptive analysis of individual control strategies has been moved to Supplementary Material—Extended Descriptive Analysis of Intelligent HVAC Control Strategies.

4.1. Model Predictive Control Approaches

MPC is the most established intelligent control strategy identified in this review, appearing in approximately 40% of the selected studies [13]. Its main advantage lies in the explicit use of a dynamic model of the building and system to forecast future states and to optimize control actions over a finite horizon subject to thermal comfort and actuator constraints [10,11,13]. Across residential and small-commercial applications, MPC-based controllers typically deliver energy savings in the range of 15–20% compared to conventional rule-based or PI/PID controllers, with some field studies reporting savings up to 70% during heating seasons when combined with thermal storage and time-of-use tariffs [13,17,18]. In addition, various economic MPC (EMPC) formulations explicitly optimize operating cost under time-varying electricity prices and feed-in tariffs, achieving 7–13.5% reductions in electricity costs while maintaining comfort [19–23].

The reviewed literature shows a wide variety of MPC formulations. Centralized MPC coordinates multiple components such as heat pumps, baseboards, photovoltaic (PV) systems, and batteries, yielding cost reductions of around 13.5% and notable peak demand reduction [17,19,20]. EMPC focuses on tariff-driven cost optimization and demand flexibility [21,22], while distributed MPC decomposes the optimization problem into local sub-problems associated with zones or subsystems to improve scalability [22,24,25]. Robust and stochastic MPC formulations address uncertainty in weather, occupancy, and model parameters, preserving comfort and stability at the expense of increased computational effort [13,26,27].

A clear trend is the integration of data-driven models into MPC. Neural networks (ANN, LSTM, and CNN–BiLSTM), Gaussian processes, and grey-box models are used as surrogate models of building thermal dynamics or load profiles, reducing modelling effort and improving prediction accuracy [21,28–31]. ANN-enhanced EMPC has achieved around 7% electricity cost reduction with indoor temperature RMSE values of 0.24–0.27 °C and optimization times under 30 min, which is acceptable for hourly control updates [21]. Mixed-integer MPC combined with LSTM prediction for variable-speed heat pumps has yielded 9–22% electricity cost savings and up to 22% reductions in CO_2 emissions, leveraging compressor speed modulation and building thermal mass [30]. Other works use symbolic regression, subspace identification, and DNN-based optimization to simplify models and tune MPC parameters [32–37].

Despite these benefits, MPC deployment in residential buildings faces several barriers. Detailed model development and calibration require expertise and can be costly in existing homes with limited documentation [11,21,33]. Nonlinear or large-scale formulations may be computationally demanding, particularly for multi-zone configurations and long prediction horizons [22,26,38]. Integration with legacy building management systems (BMS) also requires appropriate communication interfaces (e.g., BACnet, MQTT) and modular architectures to ensure interoperability [34,39,40]. Current research trends, therefore, point towards hybrid AI–MPC schemes, reduced-order or surrogate models, and hierarchical architectures that preserve the constraint-handling capabilities of MPC while improving adaptability and computational efficiency [13,14,41].

Table 2 includes an overview of the main MPC Control approaches.

Table 2. MPC control approaches and performance.

Control Approach	AI/ML Integration	Optimization Method	Key Studies
Centralized MPC	None	Multistep feedback	[19,20]
Economic MPC	ANN (Artificial Neural Network) prediction	Dynamic Programming	[21]
Distributed MPC	PSO (Particle Swarm Optimization), NOMAD (Nonlinear Optimization by Mesh Adaptive Direct Search)	Nonlinear global optimization	[22,24,25]
Robust MPC	None	Uncertainty set-based	[26]
Non-linear MPC	None	Reference governor	[42,43]
Heuristic MPC	Grey-box archetypes	Scenario clustering	[44]
ALSTM-Fast MPC	Attention-based LSTM (Long Short-Term Memory)	Deep learning optimization	[28,29]
Mixed-Integer MPC	LSTM prediction	MIP (Mixed-Integer Programming)	[30]
Model-Based Predictive Control	RBFNN (Radial Basis Function Neural Network), MOGA (Multi-Objective Genetic Algorithm)	Multi-objective optimization	[32,33]
Symbolic Regression MPC	Symbolic regression	Online optimization	[34]
DNN-Optimized MPC	Cheetah algorithm + DNN (Deep Neural Network)	Quadratic programming	[35]
Subspace Identification MPC	Compared to ANN	System identification	[36]
Genetic Algorithm MPC	Genetic algorithms	Dynamic horizon selection	[37]
Standard MPC	None	Various solvers	[38]
DRL-Enhanced MPC	LSTM, hybrid AI-physics	Various	[14]
Field-Deployed MPC	None	Online optimization	[17]

4.2. Deep Reinforcement Learning Approaches

DRL has emerged as a leading model-free alternative to MPC for intelligent HVAC control, particularly attractive when explicit building models are difficult or costly to obtain. In DRL, an agent learns a control policy through interaction with the environment, optimizing a cumulative reward that penalizes both energy consumption and comfort violations [41,45,46]. Representative algorithms include value-based methods such as DQN and DDQN, actor–critic methods such as DDPG and Soft Actor–Critic (SAC), policy-gradient methods like PPO, and hierarchical or multi-agent variants tailored to multi-zone and grid-interactive buildings [41,47–53].

Across the surveyed literature, DRL controllers often achieve energy savings between 15% and 35% compared to rule-based baselines, with cost reductions exceeding 50% in some multi-agent, occupant-centric configurations [45,47,51,54]. A DDPG-based strategy for multi-zone residential HVAC has demonstrated 15% cost reduction relative to DQN, with 79–98% reductions in comfort violations and near-zero average temperature deviation from setpoints under optimal conditions [47]. Hybrid SVR–DNN–DDPG configurations further improve thermal comfort prediction and reduce comfort violations by more than 60% compared to DQN, while providing additional energy savings [48]. Occupancy-driven DRL schemes combining LSTM with DDQN have reduced temperature and CO_2 violations by roughly 10–17% and shortened HVAC runtime by 10.9–30.4%, illustrating the value of integrating occupancy inference into the control loop [50].

https://doi.org/10.3390/app16042006

Advanced DRL architectures exploit hierarchical and multi-agent structures to handle complex objectives and distributed energy resources (DERs). Hierarchical DRL combining TD3 and PPO has achieved around 3.95% reductions in energy use and 8.37% improvements in temperature compliance while enhancing adaptability to changing occupancy and weather conditions [41]. Multi-agent hierarchical DRL using SAC and ensemble learning has delivered additional peak power and energy reductions (approximately 2.8% and 3.9%, respectively) while coordinating HVAC with DERs such as PV and batteries [53].

However, DRL presents important challenges for residential deployment. Training is sample-intensive and computationally demanding, often requiring millions of interactions in high-fidelity simulation environments before deployment [45,46,52,55,56]. Online learning in occupied homes is generally impractical due to unsafe exploration and poor comfort during early training phases [49,52,55]. DRL performance is highly sensitive to reward design, hyperparameter tuning, and exploration strategies, and standard algorithms do not natively enforce comfort or equipment constraints [56,57]. Recent work mitigates these issues through pre-training on digital twins, imitation learning from MPC, safety layers that filter actions, and conservative fine-tuning after deployment, but these practices add complexity and remain an active research area [45,49,55,57].

Comparative studies indicate that DRL can outperform MPC in highly dynamic, grid-interactive, or poorly modelled environments, while MPC retains advantages in predictability, stability, and interpretability [27,45,54]. Overall, DRL offers high flexibility and strong performance in complex scenarios, but its practical adoption requires careful attention to training methodology, safety, interpretability, and integration with existing building management systems [52,55,57].

Table 3 includes an overview of the different DRL Approaches.

Table 3. DRL control approaches and performance.

Algorithm	Algorithm Type	Application	Key Features	Key Studies
Deep Q-Network (DQN)	Value-based DRL	Multi-zone HVAC control	Discrete action spaces, experience replay	[49]
Deep Deterministic Policy Gradient (DDPG)	Actor-Critic DRL	Thermal comfort control	Continuous action spaces, deterministic policy	[47,48]
SVR (Support Vector Regression)—DNN with DDPG	Hybrid DRL approach	Multi-zone thermal comfort	Support Vector Regression + Deep Neural Network + DDPG	[48]
Double Deep Q-Network (DDQN)	Enhanced value-based DRL	Occupancy-driven control	Enhanced feature extraction, occupancy integration	[50]
Hierarchical Deep Reinforcement Learning (HDRL)	Multi-level DRL	Year-round HVAC optimization	TD3 and PPO algorithms, hybrid action spaces	[41]
Multi-Agent Deep Reinforcement Learning (MADRL)	Multi-agent DRL	Multi-zone building control	MADDPG (Multi-Agent Deep Deterministic Policy Gradient) algorithm, distributed control	[51,53]
Soft Actor-Critic (SAC)	Actor-Critic with entropy regularization	Hierarchical control with energy resources	Ensemble learning integration, continuous action space	[53]
Proximal Policy Optimization (PPO)	Policy gradient method	Human-in-the-loop HVAC control	66% faster convergence, real-time feedback integration	[58]

Table 3. *Cont.*

Algorithm	Algorithm Type	Application	Key Features	Key Studies
State-Space Approximate Dynamic Programming (SS-ADP)	Dynamic programming with neural networks	Smart home energy management	LSTM-RNN integration, reduced battery cycles	[59,60]
Occupancy-Driven LSTM-DDQN (OccD-LSTM-DDQN)	Hybrid LSTM + DDQN	Occupancy-aware HVAC control	CO_2-based occupancy estimation, privacy-preserving	[50]

4.3. Neural Network Architectures

Neural networks (NNs) play a central role in intelligent HVAC control, mainly as predictive and surrogate models rather than as standalone controllers in recent literature. NNs are widely used for forecasting indoor temperature, heating and cooling loads, occupancy patterns, CO_2 concentrations, and electricity demand [21,28,30,31,61–65]. Their ability to capture nonlinear relationships and temporal dependencies makes them particularly suitable for complex building dynamics where first-principles modelling is cumbersome or incomplete.

Among the surveyed architectures, Long Short-Term Memory (LSTM) networks are the most widely used for temperature prediction and energy forecasting tasks, achieving very high accuracy with MAE values around 0.05 °C and R^2 values above 0.99 for multistep indoor temperature forecasting [31]. Fusion models that combine CNNs with BiLSTM layers, sometimes optimized via metaheuristic algorithms such as the Dung Beetle Optimizer (DBO), further improve performance. For example, CNN–BiLSTM architectures using Gramian Angular Fields for feature extraction have reduced MAE, MAPE, and RMSE by more than 60% compared to standalone LSTM and BiLSTM models, albeit at the cost of longer training times [31]. Attention-based LSTM (ALSTM) networks have achieved MAPE values around 3.4% for short-term energy forecasting and have been integrated with fast MPC schemes (ALSTM–FMPC) to improve responsiveness to setpoint changes and thermal dynamics [28,61].

Other relevant NN architectures include radial basis function neural networks (RBFNN) for predictive control [29], GRU models optimized with metaheuristics for residential load forecasting [64], deep autoencoders for Koopman-based system identification with very low modelling error [66], and wavelet neural networks with improved ant colony optimization (I-ACO–WNN) for accurate HVAC load prediction [65]. These models are increasingly deployed on edge devices or FPGAs using frameworks such as HLS4ML, achieving inference times on the order of milliseconds or less and enabling real-time integration into control loops [63].

NNs are frequently embedded within MPC, DRL, and hybrid controllers. ANN- or LSTM-based surrogates replace physics-based models inside MPC to improve prediction accuracy and reduce modelling effort [21,30]. In DRL, deep networks approximate value functions and policies and are sometimes combined with auxiliary predictors such as SVR or DNN-based comfort models [48]. Federated learning approaches further use gradient-boosted regressors and decentralized training to personalize comfort models while preserving privacy [62].

Key methodological issues for NN-based HVAC models include generalization across buildings and seasons, dataset quality and length (ideally covering full annual cycles), non-stationarity in residential usage patterns, and computational constraints for embedded deployment [16,31,61,64]. Adaptive and incremental learning, model compression (pruning, quantization), and federated learning have been proposed to address these challenges while preserving privacy and scalability [16,62,63].

Table 4 includes an overview of the different neural network approaches.

Table 4. Overview of Neural network architectures.

Architecture	Primary Function	Input Features	Key Studies
LSTM	Temperature prediction and energy forecasting	Historical temperature, humidity	[31,61]
CNN and Fusion Models	Multi-step prediction	Gramian angular fields	[31]
ALSTM (Attention-based Long Short-Term Memory)	Energy forecasting	Operational data, weather	[28,61]
RBFNN (Radial Basis Function Neural Network)	Predictive control	Temperature, occupancy	[32]
GRU (Gated Recurrent Unit)	RL enhancement	Outdoor conditions, pricing	[64]
Deep Autoencoder	System identification	Thermal comfort indices	[66]
Wavelet Neural Network	Load prediction	Building parameters	[65]

4.4. Hybrid and Metaheuristic Approaches

Hybrid and hierarchical control strategies combine model-based optimization (e.g., MPC) with data-driven components (NNs, DRL, metaheuristics, and fuzzy logic) to exploit complementary strengths and mitigate individual weaknesses [14,21,35,41,48,67–71]. The reviewed literature shows that such controllers often achieve superior performance compared with single-method strategies, particularly in buildings with large variability or multiple interacting subsystems.

Neural-network-enhanced MPC is one of the most successful hybrid patterns. ANN–MPC and LSTM–MPC schemes have demonstrated around 7–22% reductions in electricity cost and up to 22% carbon emission reduction, while maintaining accurate temperature control and feasible optimization times [21,30]. Hybrid deep learning controllers combining SVR, DNN, and DDPG have improved thermal comfort prediction by more than 20%, reduced comfort violations by over 60%, and yielded additional energy and cost savings compared to pure DRL baselines [48]. ALSTM–FMPC integrates attention mechanisms with LSTM forecasting and fast MPC, improving dynamic adaptability to changing occupancy and weather conditions [28].

Metaheuristic optimization algorithms such as the Cheetah Optimization Algorithm (COA), Bald Eagle Search (BES), and improved ant colony optimization (I-ACO) are frequently used to tune NN architectures, select MPC horizons, or optimize operating schedules [37,65,70]. These methods have delivered cost savings of 34–57% in residential energy management under time-of-use tariffs and improved grid efficiency by up to 26% in building-to-grid integration scenarios [35,70]. Hybrid XGB–DQN models integrate gradient-boosted trees with DRL to capture occupant behavior and adjust HVAC operation, achieving substantial reductions in AC-related energy use and improvements in comfort duration [68].

Hierarchical and multi-agent control structures extend hybridization to the system architecture level. Upper layers typically perform slow, system-level economic or flexibility optimization using MPC or DRL, while lower layers implement fast zone-level comfort control using simpler controllers or local optimization [39,41,53,72–74]. This decomposition improves scalability and allows integration with legacy HVAC equipment, but introduces coordination and communication overhead that must be carefully managed to avoid instability.

Overall, hybrid and hierarchical approaches tend to outperform single-method strategies in terms of combined energy, comfort, and robustness metrics, especially in complex multi-zone or grid-interactive residential applications [14,16,41,48,53]. Their main drawback is increased design and tuning complexity, as well as higher requirements for sensing, computing, and data quality [16,22,35,70].

Table 5 includes an overview of the key hybrid and metaheuristic approaches.

Table 5. Summary of principal hybrid and metaheuristic methods applied in intelligent HVAC control.

Hybrid Approach	Component Methods	Application	Key Studies
ANN-MPC	Artificial neural networks + MPC	Demand flexibility	[21,30,67,75]
SVR-DNN-DDPG	Support vector regression + DNN + DDPG	Thermal comfort	[48]
ALSTM-FMPC	LSTM networks + fast MPC optimization	Occupant behavior	[28]
XGB—DQN (XGBoost + Deep Q-Network)	XGBoost + Deep Q-Network	Occupant behavior	[68]
ISPC (Intelligent Supervisory Predictive Control)	ANN + PSO	Supervisory control	[69]
Metaheuristic Optimization Algorithms	Cheetah Optimization Algorithm (COA) Bald Eagle Search (BES) algorithm Rule-based predictive control	Building-to-Grid (B2G) integration Residential building optimization	[35,70,71]
I-ACO-WNN	Improved ant colony + wavelet NN	Load forecasting	[65]
Fuzzy Logic Integration	ANFIS (Adaptive Neuro-Fuzzy Inference System)	HVAC performance prediction	[76–78]
HDRL (TD3 + PPO), Hierarchical DRL (SAC + Ensemble)	Multi-agent and hierarchical DRL frameworks	Dynamic occupancy and weather adaptability	[41,53]
Federated Learning and Privacy-Preserving Approaches	Federated Learning + Gradient Boosted Regressors	Personalized federated learning with occupant feedback and decentralized processing for privacy.	[62]

4.5. Fuzzy Logic and Rule-Based Systems

Fuzzy logic and advanced rule-based controllers provide an important bridge between classical control and fully data-driven methods. Fuzzy systems encode expert knowledge in the form of linguistic rules and membership functions, enabling intuitive handling of uncertainties and nonlinearities without requiring detailed mathematical models [16,79,80]. In the context of HVAC control, fuzzy controllers typically map variables such as indoor temperature, occupancy, and air quality into control actions with a high degree of interpretability.

The Human Building Interaction for Thermal Comfort (HBI–TC) framework is one of the most extensively documented fuzzy logic approaches [77,81]. It combines fuzzy predictive models with participatory sensing and zone-level control, enabling users to express comfort preferences while the system learns their profiles. Reported benefits include a 39% reduction in daily average airflow and significant improvements in user satisfaction, from 4.7 to 8.4 on a 10-point scale [77]. Adaptive Neuro-Fuzzy Inference System (ANFIS) models have achieved ventilation efficiency increases from 75% to 93%

and heat load efficiency from 79% to 97%, with low RMSE values indicating high predictive accuracy [76].

Systematic reviews of modelling techniques for HVAC control confirm that fuzzy logic is a mature approach, often delivering energy savings in the range of 20–40% compared to static setpoint or conventional rule-based control while maintaining low computational cost and high transparency [16,79]. However, pure fuzzy systems generally underperform MPC or DRL in highly dynamic, multi-objective scenarios and are increasingly used as components within hybrid architectures—for instance, as supervisory layers that define comfort envelopes, manage mode switching, or provide safe operating regions for learning-based controllers [14,76–78].

Implementation challenges include retrofitting complexity in existing HVAC systems, the need for carefully designed user interfaces to avoid feedback fatigue, and issues of scalability in open-plan or highly heterogeneous environments [77]. Nevertheless, fuzzy logic remains particularly suitable for retrofit applications and contexts where interpretability and user acceptance are critical [16,82].

4.6. Cross-Sectional Comparison

A cross-sectional comparison of MPC, DRL, NN-based prediction, hybrid frameworks, and fuzzy controllers highlights that no single control paradigm is universally optimal. MPC offers high predictability and robust constraint handling when accurate models are available, but its scalability is limited by model development and real-time optimization complexity, especially in large multi-zone systems [13,27,74]. DRL provides very high flexibility and strong performance in complex, uncertain environments, but requires extensive training data, careful reward design, and additional safety mechanisms, and its behavior can be difficult to interpret [45,47,55,57]. NNs excel as scalable predictors and surrogate models, enabling accurate forecasts that support other control strategies, but they depend on high-quality, representative datasets and may suffer from generalization issues if not regularly updated [16,31,61,64]. Fuzzy logic and rule-based systems deliver interpretable, robust performance with moderate energy savings and are well-suited for user-centric applications and retrofits, though they may not fully exploit the flexibility of emerging energy systems [16,77,79].

Hybrid and hierarchical approaches generally achieve the highest technical capabilities by combining these methods but also entail the greatest implementation complexity [14,21,35,41,53]. Their scalability depends on modular architectures, standardized communication protocols, and adequate sensing and computing infrastructure [39,40,63]. Several transversal trends emerge across studies. First, accurate and privacy-preserving occupancy information consistently appears as a key enabler for high performance, regardless of the control method, with occupancy-aware strategies achieving significantly better comfort–energy trade-offs [16,50,83–85]. Second, grid-interactive operation—including participation in demand response, time-of-use pricing, and coordination with PV and battery systems—is becoming a central objective, particularly for DRL and hybrid controllers [19,22,41,44,45,53,54]. Third, there is a clear movement toward modular, multi-layer architectures that integrate sensing, prediction, optimization, and user interaction into cohesive frameworks compatible with existing building infrastructures [14,40,41,85].

Finally, the gap between simulation-based performance and field results remains a critical challenge. Simulations often report high savings under idealized conditions, while real deployments reveal issues related to sensor reliability, network delays, occupant overrides, and changing building use that reduce achievable benefits [17,49,73,86]. Standardized benchmarks, shared datasets, and long-term field studies are therefore essential

https://doi.org/10.3390/app16042006

to evaluate and compare intelligent HVAC control strategies under realistic residential conditions [14,16,47,55].

A comparative analysis of the five major control categories is included in Table 6.

Table 6. Cross-sectional comparison of flexibility, predictability, computational complexity, and key limitations.

Approach	Flexibility	Predictability	Implementation Complexity	Key Limitations
MPC	Moderate	High	High (real-time optimization) [74]	Requires accurate models; scalability issues
DRL	Very High	Moderate	Very High (training + inference) [45]	High sample complexity; resource-intensive
Neural Networks	High	Moderate–High	Moderate (distributed computation) [72]	Communication overhead; coordination complexity
Hierarchical	High	High	Moderate–High (layered coordination) [87]	Inter-layer communication; design complexity
Hybrid	Very High	High	Very High [46]	Integration challenges; tuning complexity

4.7. Summary of Main Insights

Overall, advanced control techniques for intelligent residential HVAC systems offer substantial potential for improving energy efficiency, thermal comfort, and grid responsiveness, but their relative advantages are highly context-dependent. MPC provides a mature, well-understood framework with strong constraint handling and transparent operation, particularly appropriate when reliable models and forecasts are available [11,13,17]. DRL delivers high adaptability and strong performance in complex, uncertain scenarios, especially in occupant-centric and grid-interactive applications, though at the cost of significant training and implementation complexity [45,47,51,55]. Neural network architectures enable accurate prediction and surrogate modelling, supporting both MPC and DRL and facilitating deployment in data-rich environments [21,28,30,31,61–64]. Fuzzy logic and rule-based systems contribute robustness and interpretability and remain attractive for retrofits and user-centric control [65,76,77].

Hybrid and hierarchical control strategies, which integrate these paradigms, emerge as particularly promising for residential applications. They combine MPC's predictability and safety, DRL's flexibility, NNs' predictive intelligence, and fuzzy logic's interpretability into multi-layer architectures capable of handling diverse objectives and constraints [14,21,35,41,48,53]. Realizing this potential in practice will depend on advances in data infrastructure, scalable computing (especially edge and FPGA-based solutions), standardized communication protocols, and user-centered design, as well as robust validation through field deployments in heterogeneous residential contexts [14,16,39,40,63,86]. In this sense, intelligent HVAC control in residential buildings is moving toward flexible, adaptive, and multi-layer ecosystems that integrate technical performance with user acceptance and long-term resilience.

5. Comparative Analysis

This section provides a comparative evaluation of the intelligent HVAC control strategies discussed in previous sections, focusing on their performance in terms of energy efficiency, thermal comfort, computational complexity, and implementation feasibility in residential buildings.

5.1. HVAC Applications and System Integration

5.1.1. Functions Optimized

The reviewed studies address multiple HVAC functions with varying emphasis across different control objectives. Table 7 summarizes the main functions optimized, the number of studies focusing on each, primary performance metrics, and typical savings ranges reported in the literature.

Table 7. Comparative analysis of the principal models and performance metrics.

Function	Number of Studies	Primary Metrics	Typical Savings Range
Heating optimization	45+	Energy consumption [21], cost [19],	20–70%
Cooling optimization	40+	Peak demand [34], runtime [48]	4.9–30.2%
Thermal comfort	50+	PMV [42], temperature deviation [48]	79–98% violation reduction
Ventilation/IAQ (Indoor Air Quality)	15+	CO_2 levels [88], airflow [82]	12.2–17.4% violation reduction
Demand response	20+	Peak reduction [21], flexibility [44]	47–95% flexibility
Renewable integration	25+	Self-consumption [22], grid independence [89]	11–61% solar fraction

Multi-zone control presents particular challenges due to the complexity of thermal dynamics and interaction between zones. Multi-zone residential HVAC systems require consideration of natural air flow between floors and varying occupancy patterns across zones [65]. DDPG-based control achieves 98% reduction in comfort violations compared to rule-based strategies in multi-zone applications [47].

5.1.2. System Types and Components

Intelligent control has been applied across diverse HVAC system configurations. Table 8 summarizes the main system types, control components, integration levels, and representative studies.

Table 8. Overview of the main system types and their defining features.

System Type	Control Components	Integration Level	Key Studies
Central air systems	Air handling units, thermostats	Whole building	[17]
Heat pump systems	Variable-speed compressors	Zone-level	[30]
Multi-split systems	Indoor fan units, baseboards	Zone-level	[19]
VAV (Variable Air Volume) systems	Dampers, fans	Zone and system level	[82]
Radiant systems	Floor heating, water-air pumps	Room-level	[44]
Mixed-mode	HVAC + natural ventilation	Whole building	[90]

Variable-speed heat pumps demonstrate particular promise for intelligent control, with MPC achieving 9–22% energy cost reduction and up to 22% carbon emission reduction compared to conventional control policies [30]. The ability to modulate compressor speed enables finer control granularity than traditional on-off systems.

5.1.3. Integration with Building Systems

Advanced control strategies increasingly integrate HVAC with other building systems for holistic optimization. Table 9 summarizes the main integration types, combined systems, benefits, and representative implementation examples.

Table 9. Integration of intelligent HVAC with building systems.

Integration Type	Systems Combined	Benefits	Implementation Examples
Energy management	HVAC + PV + battery	11.6% cost reduction	[22]
Occupancy-aware	HVAC + occupancy sensors	5–13.3% savings	[83]
Grid-interactive	HVAC + demand response	44% cost reduction	[44]
Thermal storage	HVAC + building mass	61% solar fraction	[89]
EV integration	HVAC + EV battery	10% expense reduction	[20]

Building thermal mass serves as a form of thermal energy storage, enabling load shifting and increased renewable self-consumption. By strategically overheating buildings during periods of renewable availability, solar fractions can increase from 11% to 61% in single-family houses with heat pump systems [89].

5.2. Performance Outcomes

5.2.1. Energy Savings

Energy savings vary substantially across studies, influenced by baseline conditions, control method sophistication, and building characteristics. Table 10 summarizes typical energy savings, peak reduction, and cost savings reported for different control categories.

Table 10. Energy savings depending on the control category.

Control Category	Typical Energy Savings	Peak Reduction	Cost Savings	References
MPC vs. conventional	15–20%	10–30%	7–13.5%	[39]
DRL vs. rule-based	15%	Not reported	51.09%	[47]
Intelligent heating	21–26%	Not reported	34–57%	[91]
Occupancy-aware	5–13.3%	Not reported	24.4–27.2%	[83]
AI-assisted (general)	21.8–44.4%	25.8–35.1%	Variable	[16]

Field trial results demonstrate MPC achieving 20% energy savings during transition seasons and up to 70% during heating seasons compared to rule-based schedules, with peak power reduction exceeding 10% [17]. These results were obtained in a mid-size commercial building with direct digital control system integration.

The comprehensive review by [16] reports average energy savings between 21.81% and 44.36% across AI-based building control systems from 1993 to 2020. However, performance varies significantly based on building characteristics, with the highest savings observed in older buildings and households with high vacancy times [91].

5.2.2. Thermal Comfort Performance

Comfort improvements represent a critical outcome metric, often in tension with energy reduction objectives. Table 11 summarizes thermal comfort performance indicators and associated control methods.

Table 11. Thermal comfort performance and control methods.

Metric	Improvement Range	Control Method	References
PMV (Predicted Mean Vote) index control	Maintained at zero	Non-linear MPC	[42]
Temperature deviation	±0.3 °C from setpoint	VAV control	[88]
Comfort violation reduction	64–98%	DDPG vs. DQN/rule-based	[47,48]
Thermal comfort duration	24% increase	XGB-DQN	[51]
User satisfaction	4.7 to 8.4 (10-point scale)	Fuzzy predictive	[81]
Comfort improvement (AI)	21.7–85.8%	Various AI techniques	[16]

The trade-off between energy efficiency and comfort is explicitly managed through penalty factors and multi-objective optimization. MPC frameworks incorporating comfort constraints achieve temperature tracking errors less than 1 °F (0.56 °C) while maintaining energy savings [17]. The SVR-DNN model achieves 20.5% improvement in thermal comfort prediction compared to the standalone DNN, enabling more precise control targeting [48].

5.2.3. Operational and Environmental Benefits

Beyond energy and comfort, intelligent HVAC systems provide operational advantages and environmental benefits. Table 12 summarizes key outcomes.

Table 12. Operational and environmental outcomes of intelligent HVAC control.

Benefit Category	Quantified Outcome	Control Method	References
Reduced HVAC runtime	10.9–30.4% reduction	OccD-LSTM-DDQN	[50]
Carbon emission reduction	Up to 22%	Mixed-integer MPC	[30]
CO_2 emission reduction	198 gCO_2/day	Centralized MPC	[19]
Grid efficiency improvement	26%	DNN-MPC	[35]
Peak generator reduction	12%	DNN-MPC	[35]
Battery cycle reduction	Reduced deep cycles	SS-ADP	[59]

Generative AI-based control demonstrates exceptional real-world performance, achieving up to 47.92% energy reduction and 26.36% comfort improvement in an operational office setting compared to baseline operation [86]. Regression analysis confirmed robustness against confounding variables, including outdoor conditions and occupancy levels.

5.3. Input Data Requirements

Effective intelligent HVAC control depends on diverse data streams spanning environmental conditions, occupancy patterns, and system performance. Accurate and timely data acquisition is critical for predictive algorithms, optimization routines, and adaptive control strategies.

5.3.1. Environmental and Building Data

Weather forecasting integration enables proactive control strategies by anticipating outdoor conditions and adjusting HVAC operation accordingly. Ten-year datasets support solar radiation forecasting for MPC applications [21], while real-time weather API integration provides current condition updates [61]. Prediction accuracy typically degrades with forecast horizon, with optimal performance achieved within 4-h prediction windows [31]. Table 13 summarizes key environmental and building data requirements.

Table 13. Environmental and building data for intelligent HVAC control.

Data Type	Typical Sources	Update Frequency	Critical for	References
Indoor temperature	Thermostats, wireless sensors	Real-time	All control methods	[40]
Outdoor weather	Weather stations, forecasts	Hourly	Predictive control	[22]
Humidity	DHT sensors, hygrometers	Real-time	Comfort calculation	[61]
Solar radiation	Weather data, sensors	Hourly	Load prediction	[22]
Building thermal model	Calibration data	Static/periodic	MPC	[44]

5.3.2. Occupancy Detection and Its Impact on Residential HVAC Control

Occupancy is one of the most influential and least predictable disturbances in residential HVAC control. Unlike commercial buildings with standardized schedules, residential presence patterns depend on irregular routines, spontaneous activities, and heterogeneous comfort preferences. This variability directly affects internal heat gains, ventilation needs, and thermal comfort thresholds, making robust treatment of occupancy uncertainty essential for evaluating the performance of MPC, DRL, and hybrid control strategies under real residential conditions.

In this context, accurate and privacy-preserving occupancy detection becomes a key enabler for effective control. Empirical studies show that occupancy-aware strategies can achieve 19–45% energy savings compared to conventional rule-based HVAC operation. Table 14 summarizes the main detection techniques reported in the literature, including their accuracy, privacy implications, and implementation complexity.

Table 14. Occupancy detection methods: accuracy, privacy, and implementation complexity.

Detection Method	Accuracy	Privacy Concerns	Implementation Complexity	References
CO_2 concentration	70% precision, 60% recall	Low	Low	[88]
PIR sensors	Variable	Low	Low	[84]
Geo-fencing	Requires smartphone	Moderate	Moderate	[91]
Wi-Fi/Bluetooth	Variable	Moderate	Moderate	[84]
Cameras	High	High	High	[84]
Chair movement sensors	Room-specific	Low	High	[84]

Recent work shows that change-point analysis of CO_2 fluctuations provides a privacy-preserving alternative for estimating occupancy, achieving precision and recall levels of approximately 0.7 and 0.6 [50]. Although these values introduce uncertainty, the method enables timely ventilation during high-concentration periods while limiting unnecessary HVAC operation during vacant intervals.

5.3.3. Energy and Pricing Data

Dynamic electricity pricing enables economic optimization beyond simple energy minimization. Table 15 summarizes common pricing schemes, data requirements, optimization horizons, and typical savings.

Integration of energy storage systems (battery, thermal mass) requires coordination between HVAC control and energy management. MPC-based home energy management systems simultaneously optimize zone-based heating and energy flow among PV systems, batteries, and grid connections [20].

Table 15. Electricity pricing schemes: data requirements, optimization horizons, and typical savings.

Pricing Structure	Data Requirements	Optimization Horizon	Typical Savings	References
Time-of-use (ToU)	Rate schedules	24 h	13.5%	[19]
Real-time pricing	Market data	15–60 min	42.3%	[92]
Feed-in tariffs	Export rates	Variable	31%	[19]
Demand charges	Peak tracking	Monthly	2.79%	[53]

5.4. System Architecture and Implementation

5.4.1. Hardware Infrastructure

Implementation requirements vary significantly across control methodologies. Self-powered wireless sensors address deployment challenges in existing buildings, using energy harvesting from photovoltaic cells or thermoelectric generators to eliminate battery replacement requirements [32]. IEEE 802.15.4 transceivers enable low-power communication for distributed sensor networks. Table 16 summarizes hardware, computing, and communication requirements.

Table 16. Hardware, computing, and communications requirements for intelligent HVAC implementations.

Component Category	Typical Hardware	Cost Implications	Deployment Examples	References
Sensors	Temperature, humidity, CO_2	Low-moderate	Wireless networks	[40]
Controllers	Microcontrollers, edge devices	Moderate	Arduino, Raspberry Pi	[21]
Communication	Wi-Fi, LoRa, MQTT	Low	Building networks	[40]
Computing	Edge/cloud servers	Moderate-high	Local or cloud	[35]
Actuators	Valves, dampers, VFDs	High	Existing HVAC	[17]

5.4.2. Software Platforms and Frameworks

Real-time implementation requirements vary by control method. MPC optimization typically requires computational times under 30 min for practical building applications [21], while DRL inference can execute in milliseconds after offline training. The HLS4ML framework enables deployment of optimized neural network models on FPGAs, achieving substantial gains in hardware efficiency and inference speed for resource-constrained environments [63]. Table 17 summarizes software platforms.

Table 17. Software platforms for intelligent HVAC control: applications, tools, and computational requirements.

Platform	Application	Language/Tools	Computational Requirements	References
TRNSYS 18-MATLAB 2017b	Building simulation + MPC	MATLAB 2017b	High (Intel i7, 8GB RAM)	[19]
EnergyPlus 3.0	Building energy modeling	Python 3.10, BCVTB 1.6.0	Moderate-high	[93]
TensorFlow 2.16	Deep learning	Python 3.9	GPU recommended	[47]
PyTorch 2.0	Neural networks	Python 3.13	GPU (RTX 3090)	[50]
HLS4ML 1.2 0	FPGA deployment	Vivado 2025.1, TensorFlow 2.18	Low-latency edge	[63]
BOPTEST 0.9.0	RL testing	Python 3.12, OpenAI Gym 0.26.2	Moderate	[94]

5.4.3. Communication and Integration

Integration with existing building management systems represents a significant implementation consideration. BACnet protocol enables communication with standard building automation systems [34], while MQTT provides lightweight messaging for IoT sensor networks [40]. Cloud-based architectures enable remote monitoring and control but introduce latency and reliability considerations. Table 18 summarizes communication protocols.

Table 18. Communications protocols for intelligent HVAC integration: applications, latency, and integration complexity.

Protocol	Application	Latency	Integration Complexity	References
BACnet	Building automation	Low	High	[34]
Modbus	Sensor data	Low	Moderate	[88]
MQTT	IoT messaging	Low	Low	[40]
REST API	Cloud integration	Variable	Low-moderate	[80]
OPC	Industrial systems	Low	High	[17]

6. Synthesis of Findings

6.1. Explaining Variation in Energy Savings

The wide range of reported energy savings (5% to 70%) reflects systematic differences in study conditions rather than methodological inconsistency.

Studies in heating-dominated climates with older, poorly insulated buildings consistently report higher savings. Field trials during heating seasons demonstrate 70% savings [17], while transition season savings average 20%. Old buildings with high vacancy times show the greatest improvement potential, with median savings of 21–26% compared to simple on-off control [91]. Poorly-insulated buildings achieve 13% profit from intelligent control versus 26% for well-insulated buildings [95], suggesting that building envelope quality moderates but does not eliminate benefits.

Studies comparing against rule-based control consistently report higher savings than those comparing against conventional MPC. DDPG achieves 98% comfort violation reduction versus rule-based control but only 79% versus DQN [47]. Similarly, AI-assisted approaches show 51.09% cost reduction versus rule-based methods but only 4.34% versus single-agent DRL [51].

Longer prediction horizons enable greater savings but face diminishing returns beyond certain thresholds. The optimal horizon depends on building thermal mass and occupancy patterns [37]. Temperature prediction accuracy degrades with prediction horizon, with optimal performance within 4 h windows [63], while medium-term predictions benefit from cumulative training strategies, achieving the lowest RMSE of 34.9 kW [36].

6.2. MPC Versus Deep Reinforcement Learning

Both MPC and DRL demonstrate effectiveness, but their relative advantages depend on the application context.

MPC excels when accurate building models are available, and interpretability is important. MPC explicitly handles constraints [13] and provides predictable behavior through optimization-based decision making. Commercial implementation has been demonstrated with proven energy savings [17].

DRL eliminates the need for explicit building models, learning directly from interaction with the environment [47]. This model-free characteristic proves valuable when building physics are complex or poorly characterized. DRL demonstrates high generalization and adaptability to unseen environments, suggesting practical advantages for diverse building stocks.

Recent studies increasingly combine MPC with machine learning, using neural networks for prediction within MPC frameworks or employing DRL to optimize MPC parameters [14]. The ALSTM-Fast MPC system demonstrates adaptability to changing thermal dynamics while maintaining the constraint-handling capabilities of MPC [28].

6.3. Field Trial Versus Simulation Findings

Field trial results typically show lower but more reliable savings than simulation studies. Pre-trained DRL models achieve approximately 30% cost reduction in simulation, but up to 21% in real deployment [49]. This gap reflects unmodeled disturbances, sensor noise, and actuator limitations present in real buildings. Field trials also reveal practical challenges, including control delays, software malfunctions [49], and user override behaviors [44] that simulations typically neglect.

However, field trials demonstrate that intelligent control can succeed in operational settings. The "Office-in-the-Loop" system achieved 47.92% energy savings in a real office environment [86], while MPC implementation in a commercial building maintained temperature tracking errors below 1 °F throughout occupied periods [17].

6.4. Addressing Occupancy Uncertainty

Studies handling occupancy uncertainty demonstrate more robust performance across varied conditions. Monte Carlo-based uncertainty analysis provides robust estimates of MPC performance against randomness in EV arrival and departure schedules [20]. Probabilistic occupancy prediction enables aggressive demand response strategies by reducing the overestimation of productivity deterioration [96]. Change-point analysis for occupancy estimation achieves sufficient accuracy (70% precision, 60% recall) to enable meaningful energy savings while preserving privacy [50].

The federated learning approach addresses privacy concerns while enabling personalized comfort models, with real-time model adaptation at the client level [62]. This decentralized architecture reduces data transmission dependency while improving predictive accuracy through incremental client-side updates.

6.5. Implementation Challenges and Practical Considerations

6.5.1. Technical Barriers

Online learning for DRL applications is impractical due to long learning periods and poor comfort control during training. Pre-training on building models prior to deployment offers a solution, though developing accurate models for every house is not cost-effective [49]. Transfer learning and model-free approaches partially address this limitation [97]. Table 19 summarizes the main technical barriers.

Table 19. Implementation challenges and mitigation strategies for intelligent HVAC control.

Challenge	Description	Mitigation Strategies	References
Model development complexity	Building models requires significant expertise and calibration	Grey-box approaches, data-driven modelling	[21]
Computational requirements	Nonlinear optimization requires substantial resources	Distributed MPC, model simplification	[22]
Data quality and availability	High-quality real-world data is lacking	Synthetic data generation, transfer learning	[16]

Table 19. *Cont.*

Challenge	Description	Mitigation Strategies	References
Integration with existing systems	Legacy BMS (Building Management System) systems present compatibility challenges	Standard protocols (BACnet, MQTT)	[40]
Algorithm training time	DRL requires extensive training episodes	Pre-training, transfer learning	[47]

6.5.2. Cost Drivers and Economic Implications

Although the primary focus of this review is on technical control methods, economic feasibility consistently emerges in the literature as a key determinant for residential adoption. For this reason, a brief synthesis of cost categories and payback considerations is included to contextualize the practical viability of intelligent HVAC control.

The economic feasibility of intelligent HVAC control depends on several cost components and context-specific economic drivers. Rather than providing explicit calculations, this subsection synthesizes how existing studies characterize hardware, software, integration, and maintenance costs, and how these elements influence payback expectations in real deployments. Table 20 summarizes the main cost categories commonly discussed in the literature.

Table 20. Cost categories, key drivers, and payback considerations for intelligent HVAC implementations.

Cost Category	Typical Range	Key Drivers	Payback Considerations	References
Hardware (sensors)	Low-moderate	Quantity, wireless capability	Often <2 years	[32]
Software development	High	Customization requirements	Building-specific	[21,85]
Model development	High	Building uniqueness	One-time per building	[49,83]
Integration	Moderate-high	Existing infrastructure	Varies significantly	[39,40]
Maintenance	Low-moderate	Algorithm complexity	Ongoing	[62,63]

Importantly, the “Payback Considerations” column compiles qualitative insights reported across publications rather than numerical payback calculations produced in this review. These entries reflect recurrent findings related to sensor costs, customization requirements, model-development complexity, compatibility with existing infrastructure, and ongoing maintenance demands.

Previous economic analyses indicate that, in some scenarios, the financial savings enabled by MPC may not fully offset total implementation costs. Evidence of positive economic performance mainly comes from specific case studies rather than a broad consensus. For instance, the large-scale retrofitting project reported by [98] in a 560-villa community, an estimated payback period of 6.43 years and expected annual returns of USD 4 million (USD 7150 per villa). While this example illustrates the potential benefits of aggregated deployment, such results should be interpreted as context-dependent rather than universally representative.

6.5.3. Scalability and Deployment

Scalability challenges arise from building uniqueness and computational complexity.

Centralized control architectures may suffer from scalability issues, motivating the transition to decentralized and distributed settings [26]. Computational time increases with the number of devices [99] and zones [48], requiring careful algorithm design for large-scale applications.

https://doi.org/10.3390/app16042006

The modular architecture approach enables scalability and adaptability to different building configurations [39]. Edge computing with FPGAs provides low-latency processing for time-critical applications [63], while cloud-based architectures enable central management of distributed buildings [88].

6.5.4. User Acceptance and Behavioral Factors

User acceptance significantly influences the real-world performance of intelligent HVAC systems.

Participatory sensing approaches that engage occupants in comfort feedback demonstrate improved satisfaction (4.7 to 8.4 on a 10-point scale). However, poorly designed interfaces risk user fatigue and data quality degradation [77]. Human-in-the-loop frameworks balance automated optimization with occupant preferences through real-time feedback integration [58].

Occupants often override automated control decisions, particularly when comfort expectations are not met [44]. Transparent communication of system behavior and explicit accommodation of user preferences help maintain acceptance [86].

6.5.5. Regulatory and Standards Compliance

Intelligent HVAC systems must operate within established regulatory frameworks. Thermal comfort standards (e.g., [100]) provide constraints for optimization [96], while building codes define minimum ventilation rates and indoor air quality requirements [88]. Grid-interactive operation may also be influenced by guidelines and technical reports that outline interoperability requirements and demand response practices—such as the U.S. DOE technical report on grid-connected buildings [101]—although these documents are not formal interconnection standards. Privacy regulations increasingly affect occupancy sensing approaches, motivating privacy-preserving techniques, including federated learning [62] and CO_2-based estimation methods [50]. Data security for cloud-connected systems requires appropriate encryption and access controls [102].

6.5.6. Detailed Challenges and Implementation Considerations

The reviewed literature highlights that the successful deployment of intelligent HVAC control strategies in residential buildings depends on several interrelated methodological and infrastructural factors. The choice of control method is strongly conditioned by the availability and quality of building models and by the desired balance between interpretability and adaptability. MPC generally performs best when accurate thermal models are available and when transparency in decision-making is required, as demonstrated in [13]. In contrast, DRL becomes advantageous in buildings with uncertain or highly heterogeneous dynamics, particularly when avoiding the cost and effort associated with detailed model development [97]. Hybrid architectures combining data-driven prediction with MPC optimization have also been reported to offer robust performance by leveraging the strengths of both paradigms [21].

Data infrastructure is another critical foundation for effective intelligent HVAC operation. Studies consistently identify temperature measurements, weather forecasts, and equipment state data as essential inputs for ensuring reliable predictions and stable control [31]. The inclusion of occupancy data becomes particularly important in residential settings, where stochastic presence patterns introduce significant variability in thermal loads [84]. For control strategies incorporating economic objectives, access to real-time or time-of-use energy pricing information is required to optimize costs [19]. The literature also emphasizes that long-term performance depends on appropriate sensor selection, calibration, and maintenance practices to avoid degradation of data quality [63].

System architecture considerations are also decisive for practical implementation. Modular system designs enable incremental deployment, component replacement, and scalability, which are essential for long-term maintainability [39]. For latency-sensitive tasks—such as high-frequency DRL-based control or real-time fault detection—several studies highlight the advantages of edge computing solutions to reduce delays and reliance on cloud connectivity [63]. Interoperability issues are frequently addressed through the use of standardized communication protocols and structured APIs that ensure integration with existing building management systems and IoT ecosystems [13].

Finally, empirical validation is consistently described as a multi-stage process in advanced control research. Simulation-based testing and co-simulation frameworks remain the predominant first step for assessing control stability, safety, and potential savings [49]. Field pilots in representative homes or testbeds provide essential evidence on occupant interaction, sensor reliability, and real-world constraints [86]. Performance assessment typically requires comparison against clearly defined baseline conditions using metrics such as energy consumption, thermal comfort violations, and cost indicators [17]. Many studies further emphasize the need for periodic model updates or DRL retraining to adapt to changes in building dynamics, user behavior, and environmental conditions [62].

An additional consideration when evaluating AI-enabled HVAC control is the environmental cost associated with large-scale data processing and model training. Although AI-based controllers can substantially improve energy efficiency, their deployment requires collecting high-resolution environmental, occupancy, and equipment data, which increases sensing and communication demands. Training deep learning or reinforcement learning models—particularly when using cloud-based resources—can entail non-negligible computational energy consumption. As intelligent control becomes more widespread, the embodied environmental footprint of data acquisition, storage, and processing should be evaluated alongside operational energy savings. This highlights the importance of lightweight models, edge-computing architectures, and data-efficient learning strategies that reduce both computational burden and environmental impact.

6.5.7. System Resilience Under Blackout Conditions

Although intelligent HVAC control frameworks generally assume stable grid conditions and continuous availability of sensing, computation, and communication resources, real-world deployments must also consider the implications of large-scale power outages or blackout events. During prolonged loss of electrical supply, active HVAC control becomes unavailable, and indoor thermal evolution is governed primarily by the passive characteristics of the building envelope—such as insulation level, thermal mass, and airtightness [103]. These attributes largely determine how quickly indoor temperatures drift away from comfort thresholds in the absence of heating or cooling, and thus shape occupant safety during extreme conditions.

Recent studies on building resilience emphasize the role of emergency fallback modes, including the preservation of minimal control capability through battery-backed microcontrollers, manual override options, and simplified safe-state configurations to ensure ventilation for health-critical needs [104,105]. Although such strategies fall outside the typical scope of MPC or DRL-based HVAC control, they highlight the need for integrating resilience considerations into next-generation intelligent residential systems [106].

Moreover, the literature on grid-interactive buildings suggests that coupling intelligent HVAC control with home energy storage, rooftop PV systems, or microgrid islanding capabilities can significantly improve thermal survivability during blackouts by providing limited but targeted energy support [107,108]. However, these configurations require additional power-electronics interfaces and are not yet standard in residential installations.

Overall, intelligent HVAC systems cannot maintain normal operation during a prolonged blackout, but the combination of robust passive architectural characteristics, minimal emergency control capabilities, and hybrid on-site energy resources can substantially enhance resilience. Future work is expected to explore the integration of resilient control strategies, adaptive safe-state logic, and distributed energy resources within residential intelligent HVAC architectures.

7. Discussion

The findings of this review highlight the substantial progress made in intelligent HVAC control systems over the past decade and reveal a landscape of rapid technological evolution in intelligent HVAC control, marked by notable advances but also by persistent limitations that hinder large-scale deployment in residential contexts. Across model-based, model-free, and hybrid approaches, one of the most consistent observations is that performance is highly dependent on the assumptions made during system modeling and on the specific characteristics of the deployment environment. For instance, model predictive control continues to demonstrate strong reliability and constraint handling in cases where detailed and accurate thermal models are available, but real homes often exhibit heterogeneous building envelopes, unpredictable occupancy patterns, and variable thermal dynamics that challenge these assumptions. In contrast, reinforcement learning approaches show impressive adaptability and the capacity to learn effective policies without explicit modeling, yet these same strengths come with risks such as unstable learning behavior, lack of transparency, and safety concerns during exploration phases. The coexistence of these strengths and weaknesses highlights the need for multi-layered control strategies rather than reliance on a single paradigm.

A second prominent theme concerns the discrepancy between simulation-based results and real-world performance. Most studies report significant energy savings and comfort improvements in digital environments, driven by controlled disturbance profiles, idealized sensor conditions, and perfectly known system dynamics. However, real-world deployments are frequently exposed to sensor drift, communication delays, occupant interventions, seasonal changes, and unmodelled heat gains that can sharply reduce controller effectiveness. While simulations allow rapid testing of novel techniques, they often fail to capture the subtle but impactful complexities of real buildings. This divide suggests that the field still lacks standardized validation protocols and long-term field studies that examine how AI-based controllers behave under evolving conditions. Without such evidence, many promising techniques remain unproven beyond controlled experimental setups.

A third area where significant divergence arises across the literature is in the treatment of occupancy. As shown across multiple studies, occupancy is not only the main driver of comfort requirements but also the primary determinant of heating and cooling loads. Yet, in residential environments, occupancy follows patterns that are much more stochastic and individualized than in commercial buildings. Approaches such as motion sensing, CO_2 inference, machine learning prediction, and schedule profiling have all been proposed, each offering different trade-offs in accuracy, intrusiveness, and privacy. Despite these efforts, high uncertainty in occupancy estimation remains one of the major sources of energy waste and comfort violations. Furthermore, many current algorithms implicitly assume a universal comfort model, disregarding the fact that thermal preferences can vary significantly between individuals, even within the same home. This suggests the need for systems that not only detect presence but also learn resident-specific comfort profiles, potentially integrating adaptive or participatory approaches.

The comparison of different control families reveals that no single approach clearly dominates across all performance metrics: energy efficiency, comfort, interpretability, computational cost, scalability, and robustness. MPC offers transparency and stability; DRL offers flexibility and adaptiveness; neural models offer prediction strength; metaheuristics provide global search capabilities; and hybrid approaches attempt to integrate these strengths. However, hybridization also introduces complexity that may hinder adoption, particularly in typical residential settings where computational resources are limited, and HVAC equipment is often aging or constrained by proprietary interfaces. These practical considerations point to an important gap between academic proposals and market-ready solutions.

Additionally, user acceptance and human–system interaction emerge as underexplored but critical components of successful deployment. Several studies indicate that occupants frequently override automated controllers due to perceived discomfort or lack of trust in the system's decisions. These behaviors, while rarely modelled, have real implications: each override event disrupts controller learning, may force the system into suboptimal regimes, and ultimately reduces energy savings. The problem is more serious when controllers operate as "black boxes" with limited explainability, reducing user confidence. This highlights the need for more transparent, interpretable, and user-centric designs; controllers that clearly communicate their rationale, adapt to feedback, and allow intuitive interaction.

Another dimension that warrants attention is the growing influence of grid-interactive buildings and distributed energy resources. As rooftop PV, storage solutions, and electric vehicles become increasingly common, HVAC systems are becoming important assets for grid flexibility and demand response. However, most current studies focus on HVAC control as an isolated system, without considering interactions with wider household energy flows. Future solutions will likely need to account for time-of-use tariffs, carbon intensity signals, flexible loads, and local generation, shifting HVAC control from a comfort–energy optimization problem to a broader energy ecosystem optimization.

Finally, resilience and climate adaptation represent emerging priorities for the field. Extreme weather events, heatwaves, cold snaps, and potential grid instability require HVAC systems that can maintain minimum comfort and safety under adverse conditions. Yet only a small number of studies consider resilience explicitly, despite its increasing relevance. Integrating predictive mechanisms, safe fallback modes, thermal mass exploitation, and microgrid coordination could substantially enhance the robustness of residential buildings.

Taken together, these insights suggest that the future of intelligent HVAC control must not only rely on algorithmic sophistication but must also embrace deeper integration across sensing, prediction, optimization, user interaction, and energy systems. The field is moving from isolated optimization efforts toward holistic, multi-objective frameworks that align occupant needs, computational constraints, and societal goals such as decarbonization and resilience. The challenge ahead lies not only in advancing the technology but also in ensuring that these systems are scalable, trustworthy, and capable of operating effectively in the messy, ever-changing reality of residential environments.

8. Future Directions

Looking ahead, the evolution of intelligent HVAC control in residential buildings is likely to be shaped by a combination of technological, environmental, and human-centered factors. A central opportunity lies in the development of increasingly sophisticated hybrid control architectures that merge the strengths of model-based and model-free approaches. While MPC offers stability, explainability, and constraint handling, DRL and advanced neural architectures excel in adaptability and pattern recognition. Future research should therefore focus on systematically integrating these paradigms, potentially through meta-controllers capable of dynamically selecting or weighting strategies based on real-time building conditions, computational constraints, or user preferences. Such systems would enable controllers to optimize multiple objectives simultaneously—energy efficiency, thermal comfort, carbon footprint, and operational cost—while ensuring robust performance across diverse building types and climates.

Another critical direction concerns the refinement of occupancy detection and forecasting. Residential buildings exhibit highly irregular and individualized occupancy patterns, making traditional sensing strategies insufficient for reliable comfort–energy trade-offs. The next generation of HVAC systems will likely rely on multi-modal, privacy-preserving sensing frameworks that combine low-cost sensors (CO_2, motion, temperature), ambient intelligence (smart plugs, appliance signatures), and non-intrusive indicators such as Wi-Fi or BLE signal patterns. Complementing these approaches, advances in federated learning, edge-based inference, and differential privacy could enable high-quality occupancy models without compromising user data. Over time, such systems may progress beyond detection to learning personalized behavioral routines, enabling HVAC control that adapts to the specific lifestyle of each household.

A third avenue for impactful progress relates to bridging the persistent gap between simulation-based and real-world performance. Although computational environments have become increasingly sophisticated, they still fail to capture many of the nuanced uncertainties present in real buildings, such as sensor drift, actuator wear, network delays, and non-standard user behavior. Future research should therefore prioritize long-term field deployments across heterogeneous building stocks and climate zones, coupled with standardized benchmarking datasets that allow meaningful comparison between algorithms. To support this, new fault-tolerant and self-healing control frameworks will be needed: systems capable of gracefully handling missing data, model drift, hardware failures, or sudden environmental changes without requiring technician intervention.

Parallel to algorithmic advances, edge computing and lightweight AI will play an increasingly central role in enabling widespread residential adoption. Many households operate with limited computational resources, making large neural models or complex MPC solvers impractical for real-time execution. Promising areas include model compression (pruning, quantization), transformer-lite architectures, FPGA-optimized inference, and hierarchical scheduling that offloads heavy computations to the cloud only when necessary. Future systems may also incorporate energy-aware AI, where the computational footprint of the control algorithm itself is considered as part of the global optimization problem, ensuring that efficiency gains are not offset by processing overhead.

Integration with renewable energy and grid-interactive systems represents another major frontier. As homes increasingly adopt rooftop photovoltaics, storage, EV charging, and participate in demand-response programs, HVAC systems will need to function not as isolated controllers but as components of an integrated energy ecosystem. This shift will require predictive frameworks capable of coordinating thermal loads with solar generation forecasts, dynamic electricity pricing, carbon intensity signals, and grid flexibility requirements. By aligning HVAC operation with broader energy flows, such systems can play a significant role in reducing peak demand and supporting decarbonization goals.

Climate resilience also emerges as a pressing priority. With global increases in extreme weather events, buildings will require HVAC systems capable of maintaining acceptable indoor conditions during heatwaves, cold snaps, and potential grid disruptions. Future research should explore predictive resilience strategies that leverage building thermal mass, passive cooling, microgrid islanding, and minimal battery support to extend the period during which occupants remain safe during outages. Resilient HVAC control may evolve into multi-layered systems that incorporate early-warning signals, dynamic setpoint relaxation, and emergency fallback modes.

Finally, the long-term success of intelligent HVAC systems hinges on adopting a more human-centered philosophy. Despite technological sophistication, many systems

https://doi.org/10.3390/app16042006

fail due to user mistrust, overrides, or poor interpretability. Future work should prioritize interfaces that clearly communicate system decisions, enable intuitive interaction, and provide actionable explanations of expected energy or comfort outcomes. Advances in explainable AI could make DRL-based controllers more transparent, while participatory approaches may involve occupants more directly in preference shaping, comfort profiling, or feedback loops. Ultimately, the most effective systems will be those that seamlessly integrate technical performance with user acceptance and long-term engagement.

9. Conclusions

This review demonstrates that intelligent HVAC control in residential buildings sits at the intersection of advanced computational methods, increasingly heterogeneous occupant behavior, and evolving energy-system constraints. This review synthesizes findings from 97 sources examining intelligent HVAC systems for residential buildings, with particular emphasis on control techniques, applications, and artificial intelligence methods. MPC emerges as the most prevalent approach, appearing in approximately 40% of studies and demonstrating 15–20% energy savings with peak demand reductions of 10–30% compared to conventional control. Deep reinforcement learning, particularly DDPG, has emerged as a leading model-free alternative, achieving 15% energy cost reduction and 79–98% comfort violation reduction compared to rule-based strategies. Neural network architectures, including LSTM, CNN-BiLSTM, and attention mechanisms, serve critical roles in load prediction and thermal comfort modelling, with fusion models achieving prediction accuracy improvements of 66–85% over single-model approaches. Comprehensive AI-based control systems demonstrate average energy savings between 21.81% and 44.36% with comfort improvements between 21.67% and 85.77%.

The comparative analysis of model-predictive control, deep reinforcement learning, and hybrid AI-based approaches indicates that their performance is highly context dependent. MPC remains advantageous when reliable thermal models and strict constraint handling are required, whereas DRL excels in settings characterized by uncertainty, irregular occupancy, or limited modeling resources. Hybrid architectures, which combine physics-based and data-driven strengths, increasingly stand out as promising candidates for achieving balanced adaptability, interpretability, and robustness.

A central insight from the reviewed literature is the significant gap that still exists between algorithmic innovation and actual deployable solutions. Many studies rely heavily on simulation or controlled laboratory experiments, which limits the generalizability of reported results and obscures critical issues such as sensor degradation, missing data, network disruptions, or the impact of user interaction on system operation. Residential environments impose challenges that are not captured by idealized simulation setups, including device interoperability, limited computational resources, and the need for unobtrusive sensing infrastructures. Bridging the gap between research and real-world deployment will require more longitudinal field studies that test intelligent control strategies under realistic household conditions.

Human-centered considerations also emerge as essential for the widespread adoption of intelligent HVAC control in residential buildings. Most existing studies focus on energy or comfort optimization without addressing occupant acceptance, transparency, or trust. Interpretability remains a critical barrier: while rule-based and MPC strategies offer transparent decision-making processes, DRL and other data-driven approaches often operate as black boxes, raising concerns among users. Techniques such as federated learning and CO_2-based estimation offer promising alternatives, but their adoption remains limited and fragmented within the literature.

https://doi.org/10.3390/app16042006

System-level constraints reinforce the need for a holistic perspective. Intelligent HVAC solutions do not operate in isolation but interact with dynamic electricity pricing, local renewable generation, home energy storage, and evolving regulatory frameworks. Most existing studies overlook this interconnectedness, simplifying the control problem to comfort–energy trade-offs without incorporating carbon intensity, resilience, or the economic implications of grid integration. A more comprehensive research agenda is needed to address these multidimensional objectives simultaneously and to align residential HVAC control with broader sustainability and decarbonization goals.

Another aspect insufficiently covered in existing work concerns resilience to extreme events, especially blackout scenarios. Intelligent HVAC control systems generally assume stable power and communication availability, yet prolonged outages can severely compromise system operability. During such events, the building envelope's passive thermal characteristics primarily determine indoor comfort. Recent studies highlight the need for emergency fallback modes, microgrid support, battery-backed critical loads, and adaptive safe-state logic to enhance continuity of service. Integrating these resilience-oriented features into next-generation intelligent HVAC systems represents an important direction for future research.

Looking forward, progress will require advancements not only in algorithmic sophistication but also in data infrastructures, communication standards, and multi-objective design frameworks. The development of shared benchmark datasets, hybrid modeling methodologies, interpretable AI, and resilient control architectures will be essential for enabling scalable and robust deployment. Large-scale field demonstrations in diverse residential settings will also help validate the practicality and generalizability of emerging techniques. Ultimately, intelligent HVAC control holds significant potential to transform residential energy use and comfort; however, fully realizing this potential will depend on coordinated progress across sensing, computation, human factors, interoperability, and regulatory considerations. By synthesizing dispersed findings and identifying cross-cutting challenges, this review provides a comprehensive foundation for future research aimed at developing adaptive, resilient, and sustainable climate-control solutions for residential buildings.

Finally, although this review focuses on residential buildings, many of the intelligent control strategies discussed—including MPC, DRL, neural-network–based prediction, and hybrid frameworks—hold significant potential for non-residential buildings as well. Commercial and institutional buildings typically feature larger HVAC systems, higher occupant densities, and more complex operational schedules, making them strong candidates for advanced predictive and adaptive control. Extending these approaches to non-residential contexts may unlock even greater energy savings and flexibility due to larger thermal storage capacity and more sophisticated building management infrastructures.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/app16042006/s1.

Author Contributions: Conceptualization, R.F. and J.F.; methodology, R.F. and J.F.; formal analysis, R.F. and J.F.; investigation, R.F. and J.F.; validation, R.F.; writing—original draft preparation, R.F.; writing—review and editing, R.F. and J.F.; supervision, J.F. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Data used in this article can be found in 10.5281/zenodo.18171660.

Conflicts of Interest: The 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:

ACO	Ant Colony Optimization
AI	Artificial Intelligence
ALSTM	Attention-based Long Short-Term Memory
ALSTM-FMPC	Attention-based Fast MPC
ANFIS	Adaptive Neuro-Fuzzy Inference System
ANN	Artificial Neural Network
ASHRAE	American Society of Heating, Refrigerating, and Air-Conditioning Engineers
B2G	Building-to-Grid
BACnet	Building Automation and Control Networks
BES	Bald Eagle Search
BiLSTM	Bidirectional Long Short-Term Memory
BMS	Building Management System
CL	Cooling Load
CNN	Convolutional Neural Network
CO_2	Carbon Dioxide
COA	Cheetah Optimization Algorithm
DBO	Dung Beetle Optimizer
DDPG	Deep Deterministic Policy Gradient
DDQN	Dueling Deep Q-Network
DNN	Deep Neural Network
DOI	Digital Object Identifier
DP	Dynamic Programming
DQN	Deep Q-Network
DRL	Deep Reinforcement Learning
EEA	European Environment Agency
EHPA	European Heat Pump Association
EMPC	Economic Model Predictive Control
EMS	Energy Management System
EV	Electric Vehicle
FPGA	Field-Programmable Gate Array
GRU	Gated Recurrent Unit
GTO	Gorilla Troop Optimizer
HBI-TC	Human Building Interaction for Thermal Comfort
HDRL	Hierarchical Deep Reinforcement Learning
HL	Heating Load
HLS4ML	High-Level Synthesis for Machine Learning
HVAC	Heating, Ventilation, and Air Conditioning
I-ACO	Improved Ant Colony Optimization
IAQ	Indoor Air Quality
IoT	Internet of Things
ISPC	Intelligent Supervisory Predictive Control
LSTM	Long Short-Term Memory
MADDPG	Multi-Agent Deep Deterministic Policy Gradient
MADRL	Multi-Agent Deep Reinforcement Learning
MAE	Mean Absolute Error
MAPE	Mean Absolute Percentage Error
MIP	Mixed-Integer Programming
ML	Machine Learning
MOGA	Multi-Objective Genetic Algorithm
MPC	Model Predictive Control

MQTT	Message Queuing Telemetry Transport
NOMAD	Nonlinear Optimization by Mesh Adaptive Direct Search
OccD-LSTM-DDQN	Occupancy-Driven LSTM–DDQN
OPC	Open Platform Communications
PMV	Predicted Mean Vote
PPO	Proximal Policy Optimization
PRISMA	Preferred Reporting Items for Systematic reviews and Meta-Analyses
PSO	Particle Swarm Optimization
PV	Photovoltaic
R^2	Coefficient of Determination
RBFNN	Radial Basis Function Neural Network
REST API	Representational State Transfer API
RL	Reinforcement Learning
RMSE	Root Mean Square Error
SAC	Soft Actor-Critic
SS-ADP	State-Space Approximate Dynamic Programming
SVM	Support Vector Machine
SVR	Support Vector Regression
TD3	Twin Delayed DDPG
ToU	Time-of-Use (pricing)
TRNSYS	TRaNsient SYstems Simulation Program
VAV	Variable Air Volume
WNN	Wavelet Neural Network
XGB	eXtreme Gradient Boosting (XGBoost)
XGB-DQN	XGBoost + Deep Q-Network

References

1. Eurostats. Energy Consumption in Households. Available online: https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Energy_consumption_in_households (accessed on 30 December 2025).
2. US Department of Energy. Thermal Energy Storage Technologies Subprogram Area. Available online: https://www.energy.gov/eere/buildings/thermal-energy-storage (accessed on 30 December 2025).
3. European Environment Agency. Decarbonising Heating and Cooling—A Climate Imperative. Available online: https://www.eea.europa.eu/en/analysis/publications/decarbonising-heating-and-cooling (accessed on 30 December 2025).
4. Nowak, T.; Westring, P. European Heat Pump Market and Statistics Report 2023. 2023. Available online: https://www.ehpa.org/wp-content/uploads/2023/06/EHPA_market_report_2023_Executive-Summary.pdf (accessed on 30 December 2025).
5. European Heat Pump Association. REPowerEU: Heat Pump Strategy Required to Help Sector Deliver. 2024. Available online: https://ehpa.org/news-and-resources/press-releases/repowereu-heat-pump-strategy-required-to-help-sector-deliver-2/ (accessed on 30 December 2025).
6. Page, M.J.; McKenzie, J.E.; Bossuyt, P.M.; Boutron, I.; Hoffmann, T.C.; Mulrow, C.D.; Shamseer, L.; Tetzlaff, J.M.; Akl, E.A.; Brennan, S.E.; et al. The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ* **2021**, *372*, n71. [CrossRef]
7. Whitfield, S.; Hofmann, M.A. Elicit: AI literature review research assistant. *Public Serv. Q.* **2023**, *19*, 201–207. [CrossRef]
8. Wang, S.; Xu, X. Simplified building model for transient thermal performance estimation using GA-based parameter identification. *Int. J. Therm. Sci.* **2006**, *45*, 419–432. [CrossRef]
9. Åström, K.J.; Hägglund, T.; Åström, K.J. *PID Controllers*, 2nd ed.; International Society for Measurement and Control: Research Triangle Park, NC, USA, 1995.
10. Ma, Y.; Borrelli, F.; Hencey, B.; Coffey, B.; Bengea, S.; Haves, P. Model Predictive Control for the Operation of Building Cooling Systems. *IEEE Trans. Control Syst. Technol.* **2011**, *20*, 796–803. [CrossRef]
11. Afram, A.; Janabi-Sharifi, F. Theory and applications of HVAC control systems—A review of model predictive control (MPC). *Build. Environ.* **2014**, *72*, 343–355. [CrossRef]
12. Afram, A.; Janabi-Sharifi, F. Review of modeling methods for HVAC systems. *Appl. Therm. Eng.* **2014**, *67*, 507–519. [CrossRef]

13. Serale, G.; Fiorentini, M.; Capozzoli, A.; Bernardini, D.; Bemporad, A. Model Predictive Control (MPC) for Enhancing Building and HVAC System Energy Efficiency: Problem Formulation, Applications and Opportunities. *Energies* **2018**, *11*, 631. [CrossRef]
14. Gassar, A.A.A.; Jafar, R. Artificial Intelligence-Enabled Heating entilation, and Air Conditioning Systems Toward Zero-Emission Buildings: A Systematic Review of Applications, Challenges, and Future Directions. *Appl. Sci.* **2025**, *15*, 10497. [CrossRef]
15. Xin, X.; Zhang, Z.; Zhou, Y.; Liu, Y.; Wang, D.; Nan, S. A comprehensive review of predictive control strategies in heating entilation, and air-conditioning (HVAC): Model-free vs. model. *J. Build. Eng.* **2024**, *94*, 110013. [CrossRef]
16. Merabet, G.H.; 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]
17. Bengea, S.C.; Kelman, A.D.; Borrelli, F.; Taylor, R.; Narayanan, S. Implementation of model predictive control for an HVAC system in a mid-size commercial building. *HVACR Res.* **2014**, *20*, 121–135. [CrossRef]
18. Felez, R.; Felez, J. Advanced Energy Management for Residential Buildings Optimizing Costs and Efficiency Through Thermal Energy Storage and Predictive Control. *Appl. Sci.* **2025**, *15*, 880. [CrossRef]
19. Seal, S.; Boulet, B.; Dehkordi, V.R. Centralized model predictive control strategy for thermal comfort and residential energy management. *Energy* **2020**, *212*, 118456. [CrossRef]
20. Seal, S.; Dehkordi, B.B.R.; Bouffard, F.; Joos, G. Centralized MPC for Home Energy Management with EV as Mobile Energy Storage Unit. *IEEE Trans. Sustain. Energy* **2023**, *14*, 1425–1435. [CrossRef]
21. Finck, C.; Li, R.; Zeiler, W. Economic model predictive control for demand flexibility of a residential building. *Energy* **2019**, *176*, 365–379. [CrossRef]
22. Kuboth, S.; Heberle, F.; König-Haagen, A.; Brüggemann, D. Economic model predictive control of combined thermal and electric residential building energy systems. *Appl. Energy* **2019**, *240*, 372–385. [CrossRef]
23. Philipp, Z.; Alexander, E.; Lutz, G.; Veit, H.; Dominique, S.; Timm, F. A Comparison of Economic MPC Formulations for Thermal Building Control. In *Proceedings of the 2019 IEEE PES Innovative Smart Grid Technologies Europe (ISGT-Europe)*; IEEE: Piscataway, NJ, USA, 2019.
24. Narayanan, M.; Mengedoht, G.; Commerell, W. Evaluation of SOFC-CHP's ability to integrate thermal and electrical energy system decentrally in a single-family house with model predictive controller. *Sustain. Energy Technol. Assess.* **2021**, *48*, 101643. [CrossRef]
25. Scherer, H.F.; Pasamontes, M.; Guzmán, J.L.; Álvarez, J.D.; Camponogara, E.; Normey-Rico, J.E. Efficient building energy management using distributed model predictive control. *J. Process Control.* **2014**, *24*, 740–749. [CrossRef]
26. Carli, R.; Cavone, G.; Pippia, T.; De Schutter, B.; Dotoli, M. Robust Optimal Control for Demand Side Management of Multi-Carrier Microgrids. *IEEE Trans. Autom. Sci. Eng.* **2022**, *19*, 1338–1351. [CrossRef]
27. Killian, M.; Kozek, M. Ten questions concerning model predictive control for energy efficient buildings. *Build. Environ.* **2016**, *105*, 403–412. [CrossRef]
28. Safari, A.; Kharrati, H.; Rahimi, A. A hybrid attention-based long short-term memory fast model for thermal regulation of smart residential buildings. *IET Smart Cities* **2024**, *6*, 361–371. [CrossRef]
29. Azzi, A.; Abid, M.; Hanif, A.; Bensag, H.; Tabaa, M.; Hachimi, H.; Youssfi, M. A Hybrid Control Strategy Combining Reinforcement Learning and MPC-LSTM for Energy Management in Building. *Energies* **2025**, *18*, 4783. [CrossRef]
30. Lee, Z.; Gupta, K.; Kircher, K.J.; Zhang, K.M. Mixed-integer model predictive control of variable-speed heat pumps. *Energy Build.* **2019**, *198*, 75–83. [CrossRef]
31. Zheng, G.; Jia, R.; Yi, W.; Yue, X. Multi-step fusion model for predicting indoor temperature in residential buildings based on attention mechanism and neural network. *J. Build. Eng.* **2025**, *102*, 112057. [CrossRef]
32. Ruano, A.; Silva, S.; Duarte, H.; Ferreira, P.M. Wireless Sensors and IoT Platform for Intelligent HVAC Control. *Appl. Sci.* **2018**, *8*, 370. [CrossRef]
33. Luzi, M.; Vaccarini, M.; Lemma, M. A tuning methodology of Model Predictive Control design for energy efficient building thermal control. *J. Build. Eng.* **2019**, *21*, 28–36. [CrossRef]
34. Zhao, D.; Watari, D.; Ozawa, Y.; Taniguchi, I.; Suzuki, T.; Shimoda, Y.; Onoye, T. Data-driven online energy management framework for HVAC systems: An experimental study. *Appl. Energy* **2023**, *352*, 121921. [CrossRef]
35. Babosalam, S.; Kargar, S.M.; Moazzami, M.; Zanjani, S.M. Occupancy-aware energy optimization in building-to-grid systems using deep neural networks and model predictive control. *J. Build. Eng.* **2025**, *112*, 113714. [CrossRef]
36. Jiang, W.; Wang, P.; Ma, X.; Liu, Y. Development of a grey-box heat load prediction model by subspace identification method for heating building. *Build. Environ.* **2025**, *280*, 113119. [CrossRef]
37. Laguna, G.; Mor, G.; Lazzari, F.; Gabaldon, E.; Erfani, A.; Saelens, D.; Cipriano, J. Dynamic horizon selection methodology for model predictive control in buildings. *Energy Rep.* **2022**, *8*, 10193–10202. [CrossRef]
38. Drgona, J.; Kvasnica, M. Comparison of MPC strategies for building control. In *Proceedings of the 2013 International Conference on Process Control (PC), Strbske Pleso, Slovakia, 18–21 June 2013*; IEEE: Piscataway, NJ, USA, 2013; pp. 401–406. [CrossRef]

39. Banjac, A.; Novak, H.; Vašak, M. Implementation of model predictive indoor climate control for hierarchical building energy management. *Control Eng. Pract.* **2023**, *136*, 105536. [CrossRef]
40. Markiewicz, M.; Skała, A.; Grela, J.; Janusz, S.; Stasiak, T.; Latoń, D.; Bielecki, A.; Bańczyk, K. The Architecture for Testing Central Heating Control Algorithms with Feedback from Wireless Temperature Sensors. *Energies* **2023**, *16*, 5584. [CrossRef]
41. Liao, C.; Miyata, S.; Qu, M.; Akashi, Y. Year-round operational optimization of HVAC systems using hierarchical deep reinforcement learning for enhancing indoor air quality and reducing energy consumption. *Appl. Energy* **2025**, *390*, 125816. [CrossRef]
42. Castilla, M.; Álvarez, J.D.; Normey-Rico, J.E.; Rodríguez, F. Thermal comfort control using a non-linear MPC strategy: A real case of study in a bioclimatic building. *J. Process Control.* **2014**, *24*, 703–713. [CrossRef]
43. Castilla, M.; Alvarez, J.D.; Normey-Rico, J.E.; Rodriguez, F. A nonlinear model based predictive control strategy to maintain thermal comfort inside a bioclimatic building. In *Proceedings of the 2012 20th Mediterranean Conference on Control & Automation (MED), Barcelona, Spain, 3–6 July 2012*; IEEE: Piscataway, NJ, USA, 2012; pp. 665–671. [CrossRef]
44. Morovat, N.; Athienitis, A.K.; Candanedo, J.A.; Nouanegue, H.F. Heuristic model predictive control implementation to activate energy flexibility in a fully electric school building. *Energy* **2024**, *296*, 131126. [CrossRef]
45. Wei, T.; Wang, Y.; Zhu, Q. Deep Reinforcement Learning for Building HVAC Control. In *DAC '17: Proceedings of the 54th Annual Design Automation Conference 2017*; ACM: New York, NY, USA, 2017. [CrossRef]
46. 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]
47. 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]
48. Ding, Z.-K.; Fu, Q.-M.; Chen, J.-P.; Wu, H.-J.; Lu, Y.; Hu, F.-Y. Energy-efficient control of thermal comfort in multi-zone residential HVAC via reinforcement learning. *Connect. Sci.* **2022**, *34*, 2364–2394. [CrossRef]
49. Kurte, K.; Munk, J.; Kotevska, O.; Amasyali, K.; Smith, R.; McKee, E.; Du, Y.; Cui, B.; Kuruganti, T.; Zandi, H. Evaluating the Adaptability of Reinforcement Learning Based HVAC Control for Residential Houses. *Sustainability* **2020**, *12*, 7727. [CrossRef]
50. Hou, F.; Cheng, J.C.P.; Ma, J.; Kwok, H.H.L.; Huang, C.; Wu, Z. Occupancy-driven HVAC control optimization via LSTM and deep reinforcement learning for enhanced indoor air quality, thermal comfort and energy efficiency. *Build. Environ.* **2025**, *284*, 113501. [CrossRef]
51. Liu, X.; Wu, Y.; Wu, H. Enhancing HVAC energy management through multi-zone occupant-centric approach: A multi-agent deep reinforcement learning solution. *Energy Build.* **2023**, *303*, 113770. [CrossRef]
52. Al Sayed, K.; Boodi, A.; Broujeny, R.S.; Beddiar, K. Reinforcement learning for HVAC control in intelligent buildings: A technical and conceptual review. *J. Build. Eng.* **2024**, *95*, 110085. [CrossRef]
53. Pei, Y.; Yao, Y.; Zhao, J.; Hao, J.; Ding, F.; Wang, J. Multi-Agent Hierarchical Deep Reinforcement Learning for HVAC Control with Flexible DERs. *IEEE Trans. Smart Grid* **2025**, *16*, 5589–5601. [CrossRef]
54. Zhan, S.; Lei, Y.; Chong, A. Comparing model predictive control and reinforcement learning for the optimal operation of building-PV-battery systems. *E3S Web Conf.* **2023**, *396*, 4018. [CrossRef]
55. Ajifowowe, I.; Chang, H.; Lee, C.S.; Chang, S. Prospects and challenges of reinforcement learning- based HVAC control. *J. Build. Eng.* **2024**, *98*, 111080. [CrossRef]
56. Azuatalam, D.; Lee, W.-L.; De Nijs, F.; Liebman, A. Reinforcement learning for whole-building HVAC control and demand response. *Energy AI* **2020**, *2*, 100020. [CrossRef]
57. Al Sayed, K.; Boodi, A.; Broujeny, R.S.; Beddiar, K. Reinforcement Learning for Optimal HVAC Control: From Theory to Real-World Applications. In *Proceedings of the IECON 2023—49th Annual Conference of the IEEE Industrial Electronics Society*; IEEE: Piscataway, NJ, USA, 2023; pp. 1–6. [CrossRef]
58. Liang, X.; De Nijs, F.; Say, B.; Wang, H. Human-in-the-Loop AI for HVAC Management Enhancing Comfort and Energy Efficiency. In *E-Energy '25: Proceedings of the 16th ACM International Conference on Future and Sustainable Energy Systems, Rotterdam, The Netherlands, 17–20 June 2025*; ACM: New York, NY, USA, 2025; pp. 359–370. [CrossRef]
59. Zhao, Y.; Genovese, P.V.; Li, Z. Intelligent Thermal Comfort Controlling System for Buildings Based on IoT and AI. *Future Internet* **2020**, *12*, 30. [CrossRef]
60. Jeddi, B.; Mishra, Y.; Ledwich, G. Differential Dynamic Programming Based Home Energy Management Scheduler. *IEEE Trans. Sustain. Energy* **2019**, *11*, 1427–1437. [CrossRef]
61. Sarker, G.C.; Hasan, M.M.; Hoque, M.R.; Alam, M.K.; Ahmed, M.H.U. Development of an IoT-driven deep learning framework for smart energy prediction in commercial buildings. *J. Build. Eng.* **2025**, *111*, 113315. [CrossRef]
62. Avci, A.B. Federated machine learning for indoor thermal comfort: Real-time zone-based personalization and HVAC control. *J. Build. Eng.* **2025**, *114*, 114476. [CrossRef]
63. Mshragi, M.; Petri, I. Fast machine learning for building management systems. *Artif. Intell. Rev.* **2025**, *58*, 211. [CrossRef]

64. Wang, B.; Wang, X.; Wang, N.; Javaheri, Z.; Moghadamnejad, N.; Abedi, M. Machine learning optimization model for reducing the electricity loads in residential energy forecasting. *Sustain. Comput. Inform. Syst.* **2023**, *38*, 100876. [CrossRef]
65. Huang, Y.; Li, C. Accurate heating entilation and air conditioning system load prediction for residential buildings using improved ant colony optimization and wavelet neural network. *J. Build. Eng.* **2021**, *35*, 101972. [CrossRef]
66. Wahba, N.; Rismanchi, B.; Pu, Y.; Aye, L. Efficient HVAC system identification using Koopman operator and machine learning for thermal comfort optimisation. *Build. Environ.* **2023**, *242*, 110567. [CrossRef]
67. Chua, K.J.; Chou, S.K.; Yang, W.M.; Yan, J. Achieving better energy-efficient air conditioning—A review of technologies and strategies. *Appl. Energy* **2013**, *104*, 87–104. [CrossRef]
68. Liu, X.; Gou, Z. Occupant-centric HVAC and window control: A reinforcement learning model for enhancing indoor thermal comfort and energy efficiency. *Build. Environ.* **2024**, *250*, 111197. [CrossRef]
69. Gonçalves, D.; Sheikhnejad, Y.; Oliveira, M.; Martins, N. One step forward toward smart city Utopia: Smart building energy management based on adaptive surrogate modelling. *Energy Build.* **2020**, *223*, 110146. [CrossRef]
70. Saini, S.Y.K.; Kumar, R.; Lamba, R.; Saxena, A. Optimal energy management system for residential buildings considering the time of use price with swarm intelligence algorithms. *J. Build. Eng.* **2022**, *59*, 105062. [CrossRef]
71. Alibabaei, N.; Fung, A.S.; Raahemifar, K.; Moghimi, A. Effects of intelligent strategy planning models on residential HVAC system energy demand and cost during the heating and cooling seasons. *Appl. Energy* **2017**, *185*, 29–43. [CrossRef]
72. Wilk, P.; Wang, N.; Li, J. Multi-Agent Reinforcement Learning for Smart Community Energy Management. *Energies* **2024**, *17*, 5211. [CrossRef]
73. Bae, Y.; Bhattacharya, S.; Cui, B.; Lee, S.; Li, Y.; Zhang, L.; Im, P.; Adetola, V.; Vrabie, D.; Leach, M.; et al. Sensor impacts on building and HVAC controls: A critical review for building energy performance. *Adv. Appl. Energy* **2021**, *4*, 100068. [CrossRef]
74. Oldewurtel, F.; Parisio, A.; Jones, C.N.; Gyalistras, D.; Gwerder, M.; Stauch, V.; Lehmann, B.; Morari, M. Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy Build.* **2012**, *45*, 15–27. [CrossRef]
75. Allahseh, D.; Böttner, J.; Al-Addous, M.; Lenz, V. Advancements in hybrid heating systems for residential applications. *Energy Explor. Exploit.* **2025**, *43*, 2221–2275. [CrossRef]
76. Khan, O.; Parvez, M.; Seraj, M.; Yahya, Z.; Devarajan, Y.; Nagappan, B. Optimising building heat load prediction using advanced control strategies and Artificial Intelligence for HVAC system. *Therm. Sci. Eng. Prog.* **2024**, *49*, 102484. [CrossRef]
77. Jazizadeh, F.; Ghahramani, A.; Becerik-Gerber, B.; Kichkaylo, T.; Orosz, M. User-led decentralized thermal comfort driven HVAC operations for improved efficiency in office buildings. *Energy Build.* **2014**, *70*, 398–410. [CrossRef]
78. Ghahramani, A.; Jazizadeh, F.; Becerik-Gerber, B. A knowledge based approach for selecting energy-aware and comfort-driven HVAC temperature set points. *Energy Build.* **2014**, *85*, 536–548. [CrossRef]
79. Afroz, Z.; Shafiullah, G.; Urmee, T.; Higgins, G. Modeling techniques used in building HVAC control systems: A review. *Renew. Sustain. Energy Rev.* **2018**, *83*, 64–84. [CrossRef]
80. Metzmacher, H.; Syndicus, M.; Warthmann, A.; van Treeck, C. Exploratory comparison of control algorithms and machine learning as regulators for a personalized climatization system. *Energy Build.* **2022**, *255*, 111653. [CrossRef]
81. Jazizadeh, F.; Joshi, V.; Battaglia, F. Adaptive and distributed operation of HVAC systems: Energy and comfort implications of active diffusers as new adaptation capacities. *Build. Environ.* **2020**, *186*, 107089. [CrossRef]
82. Pang, Z.; Chen, Y.; Zhang, J.; O'Neill, Z.; Cheng, H.; Dong, B. Nationwide HVAC energy-saving potential quantification for office buildings with occupant-centric controls in various climates. *Appl. Energy* **2020**, *279*, 115727. [CrossRef]
83. Turley, C.; Jacoby, M.; Pavlak, G.; Henze, G. Development and Evaluation of Occupancy-Aware HVAC Control for Residential Building Energy Efficiency and Occupant Comfort. *Energies* **2020**, *13*, 5396. [CrossRef]
84. Esrafilian-Najafabadi, M.; Haghighat, F. Occupancy-based HVAC control systems in buildings: A state-of-the-art review. *Build. Environ.* **2021**, *197*, 107810. [CrossRef]
85. Aguilar, J.; Garces-Jimenez, A.; Gallego-Salvador, N.; De Mesa, J.A.G.; Gomez-Pulido, J.M.; Garcia-Tejedor, A.J. Autonomic Management Architecture for Multi-HVAC Systems in Smart Buildings. *IEEE Access* **2019**, *7*, 123402–123415. [CrossRef]
86. Sawada, T.; Mizuno, M.; Hasegawa, T.; Yokoyama, K.; Kono, M. Office-in-the-Loop: An investigation into Agentic AI for advanced building HVAC control systems. *Data-Centric Eng.* **2025**, *6*, e31. [CrossRef]
87. Parisio, A.; Rikos, E.; Glielmo, L. A Model Predictive Control Approach to Microgrid Operation Optimization. *IEEE Trans. Control Syst. Technol.* **2014**, *22*, 1813–1827. [CrossRef]
88. Rahnama, S.; Hultmark, G.; Rupnik, K.; Vogler-Finck, P.; Afshari, A. Control logic for a novel HVAC system providing room-based indoor climate control in residential buildings. *J. Build. Eng.* **2022**, *65*, 105766. [CrossRef]
89. Thür, A.; Calabrese, T.; Streicher, W. Smart grid and PV driven ground heat pump as thermal battery in small buildings for optimized electricity consumption. *Sol. Energy* **2018**, *174*, 273–285. [CrossRef]
90. Do, H.; Cetin, K.S. Mixed-Mode Ventilation in HVAC System for Energy and Economic Benefits in Residential Buildings. *Energies* **2022**, *15*, 4429. [CrossRef]

91. Nägele, F.; Kasper, T.; Girod, B. Turning up the heat on obsolete thermostats: A simulation-based comparison of intelligent control approaches for residential heating systems. *Renew. Sustain. Energy Rev.* **2017**, *75*, 1254–1268. [CrossRef]
92. Li, H.; Wei, Q. A novel energy management method for multiple residential energy systems with energy exchange. *Neurocomputing* **2023**, *575*, 127185. [CrossRef]
93. Nguyen, A.T.; Pham, D.H.; Oo, B.L.; Santamouris, M.; Ahn, Y.; Lim, B.T.H. Modelling building HVAC control strategies using a deep reinforcement learning approach. *Energy Build.* **2024**, *310*, 114065. [CrossRef]
94. Gao, Y.; Shi, S.; Miyata, S.; Akashi, Y. Successful application of predictive information in deep reinforcement learning control: A case study based on an office building HVAC system. *Energy* **2024**, *291*, 130344. [CrossRef]
95. Louadj, K.; Demim, F.; Nemra, A.; Aidene, M.; Marthon, P.; Iddir, H. Optimization of Electricity Consumption in a building. In *Proceedings of the 2018 5th International Conference on Control, Decision and Information Technologies (CoDIT), Thessaloniki, Greece, 10–13 April 2018*; IEEE: Piscataway, NJ, USA, 2018; pp. 1–5. [CrossRef]
96. Lim, J.; Choi, W. Influence of a Better Prediction of Thermal Satisfaction for the Implementation of an HVAC-Based Demand Response Strategy. *Energies* **2022**, *15*, 3094. [CrossRef]
97. Michailidis, P.; Michailidis, I.; Vamvakas, D.; Kosmatopoulos, E. Model-Free HVAC Control in Buildings: A Review. *Energies* **2023**, *16*, 7124. [CrossRef]
98. Binu, B.S.; Dudhe, R.; More, B.; Mathew, E.; Kunhabdulla, F.; Khanam, M. Comprehensive retrofitting solution aligned with COP28 cross-cutting themes for energy efficient residential community: Toward net-zero. *Open House Int.* **2025**, *50*, 607–628. [CrossRef]
99. Zhao, Z.; Keerthisinghe, C. A Fast and Optimal Smart Home Energy Management System: State-Space Approximate Dynamic Programming. *IEEE Access* **2020**, *8*, 184151–184159. [CrossRef]
100. Magnuski, M.J. Review of ASHRAE symposium session on the application of adaptive control to real HVAC applications. *Int. J. Refrig.* **1988**, *11*, 11–15. [CrossRef]
101. Harris, C. Grid-Interactive Efficient Buildings Technical Report Series: Windows and Opaque Envelope. 2019. Available online: https://www1.eere.energy.gov/buildings/pdfs/75387.pdf (accessed on 28 December 2025).
102. Wang, L. Location method of transit station for rail transit based on variable step size any colony algorithm. In *Proceedings of the Sixth International Conference on Computer Information Science and Application Technology (CISAT 2023), Hangzhou, China, 26–28 May 2023*; Jia, S., Dong, H., Eds.; SPIE: Bellingham, WA, USA, 2023; p. 292. [CrossRef]
103. Hong, T.; Malik, J.; Krelling, A.; O'BRien, W.; Sun, K.; Lamberts, R.; Wei, M. Ten questions concerning thermal resilience of buildings and occupants for climate adaptation. *Build. Environ.* **2023**, *244*, 110806. [CrossRef]
104. Rehman, H.U.; Hamdy, M.; Hasan, A. Towards Extensive Definition and Planning of Energy Resilience in Buildings in Cold Climate. *Buildings* **2024**, *14*, 1453. [CrossRef]
105. Attia, S. Resilient cooling of buildings to protect against heatwaves and power outages. In *Adapting the Built Environment for Climate Change*; Elsevier: Amsterdam, The Netherlands, 2023; pp. 215–228. [CrossRef]
106. Mehraban, M.H.; Mirzabeigi, S.; Faraji, S.; Soltanian-Zadeh, S.; Sepasgozar, S.M.E. AI-Driven Prediction of Building Energy Performance and Thermal Resilience During Power Outages: A BIM-Simulation Machine Learning Workflow. *Buildings* **2025**, *15*, 3950. [CrossRef]
107. Mahfuz-Ur-Rahman, A.M.; Islam, M.R.; Muttaqi, K.M.; Sutanto, D. An Effective Energy Management with Advanced Converter and Control for a PV-Battery Storage Based Microgrid to Improve Energy Resiliency. *IEEE Trans. Ind. Appl.* **2021**, *57*, 6659–6668. [CrossRef]
108. Galvan, E.; Mandal, P.; Sang, Y. Networked microgrids with roof-top solar PV and battery energy storage to improve distribution grids resilience to natural disasters. *Int. J. Electr. Power Energy Syst.* **2020**, *123*, 106239. [CrossRef]

https://doi.org/10.3390/app16042006