



International Journal of Advance Research Publication and Reviews

Vol 02, Issue 08, pp 770-788, August 2025

Challenges in Reinforcement Learning for Control and Optimization of Real Buildings

Geku, Diton

*Department of Electrical and Electronic Engineering, Federal university otuoke
gekudd@fuotuoke.edu.ng*

ABSTRACT

One potential approach to improving and enhancing energy use in buildings is artificial intelligence, which is part of the larger digitalisation of society. Historically, rule-based controllers that required human tuning have been used to govern building service systems. Although they are easy to use and inexpensive, they may not be the best option and won't be able to adjust to new circumstances. Control and optimisation problems are ideal for the kind of machine learning known as reinforcement learning (RL). Aiming to maximise cumulative rewards, RL takes a trial-and-error approach to optimal control. Its potential for controlling and optimising construction applications has been demonstrated in simulated studies. Although there is a lot of literature on the topic, practical applications are more difficult to come by. After establishing the existence of AI-based market solutions, this study compiles the results of the applications' extensive testing on real-world deployments, including how the studies dealt with issues like adaptable solution security and the potential for solutions to be used on a larger scale. Pretraining and online deployment could be made possible with digital twin technologies. Furthermore, the study assesses the implementations' handling of ethical considerations and finds that RL-control systems do not have an ethical analysis in place. There are major gaps in our knowledge of the dependability and efficacy of the suggested RL solutions in complicated real-world contexts due to a lack of long-term research and an oversimplification that focusses on HVAC alone rather than the full building management system.

1. Introduction

Europe's digital transformation reflects the adoption and integration of digital technologies across governance, industry, and society. Among these, artificial intelligence (AI) has been highlighted as a key enabler, capable of supporting cost-effective and sustainable energy solutions. By April 2024, the European Union (EU) had already enacted 63 policies related to AI, with the region's AI market projected to reach USD 46.67 billion in 2024. Within the EU energy agenda, digitalisation is particularly vital because it enhances the integration of renewable energy into the grid, thereby supporting decarbonisation. The use of AI in energy systems is often combined with Internet of Things (IoT) devices, which provide real-time monitoring and control. Globally, IoT devices are forecasted to grow from 15.1 billion in 2023 to more than 29 billion by 2030.

Paiho et al. [7] described the "twin transition" as the merging of digitalisation with the European Green Deal (EGD) targets, aiming to create resilient, carbon-neutral, and sustainable environments. AI applications are increasingly being used throughout the lifecycle of buildings, especially in smart building management [8]. While the overall market maturity of AI in this field remains at the demonstration stage, its role in enabling intelligent building systems is considered highly significant [7]. For instance, Tuomela [9] reported that Finnish households achieved up to 30% winter energy savings after adopting smart home energy management systems.

Several studies have assessed AI's role in the construction sector, covering topics such as intelligent control for thermal comfort [10], machine learning in energy demand forecasting [11], challenges in adopting AI for efficiency [12], energy

optimisation throughout the building lifecycle [13], neural network-based energy prediction [14], and AI-driven smart building solutions [15]. Despite these advances, practical deployment in real buildings is rare, mainly due to the need for large, high-quality datasets. In most cases, building data is used for training models rather than applying them in real-world control scenarios [10].

Digital twin technology, closely connected to AI-based control, has attracted increasing attention. A digital twin links IoT-enabled physical infrastructure to a virtual model, where AI/ML analytics interpret and predict building behaviour. Much of the literature on digital twins consists of conceptual or review papers [16–18], while some focus on AI's role in shaping their development [19]. Theoretical models have been proposed for digital twins to improve maintenance [20], occupant comfort [21], construction automation [22], and industrial building operations [23]. Only a handful of studies have evaluated digital twins in operational buildings through real-world case studies [24–26].

Ethical questions surrounding AI are especially prominent in Nordic countries [27,28]. At the EU level, ethical frameworks are being formalised, and UNESCO has issued global recommendations. However, these guidelines remain general, without prescribing specific applications [32]. Within the building sector, how to embed ethical concerns such as transparency and human-centricity into AI deployment remains unclear. Although reinforcement learning (RL) and digital twin applications show strong promise, they lack tested, large-scale real-world implementations. This study therefore seeks to provide a foundation for applying AI-driven control and optimisation in actual buildings, focusing particularly on reinforcement learning for HVAC systems.

2. Reinforcement learning in HVAC applications

Heating, ventilation, and air conditioning (HVAC) systems have traditionally been managed with simple rule-based controllers. These rely on fixed schedules and reactive feedback loops [33–35]. While such approaches are inexpensive and easy to apply, they are inefficient in complex settings where nonlinear dynamics and long time delays exist. Adjustments are usually handled manually by specialists, as opposed to automated optimisation methods. However, these systems cannot adapt to new conditions and often require recalibration when significant changes occur, such as retrofitting or variations in occupant behaviour [33–35]. They also lack the ability to anticipate external factors like weather forecasts or occupancy patterns. Model predictive control (MPC) offers a more advanced option, capable of predicting behaviour several steps ahead [33,36]. Yet, despite its strengths, MPC depends on highly detailed building and HVAC models, which are expensive and time-intensive to design and maintain.

In recent years, reinforcement learning (RL) has emerged as a promising alternative for HVAC control [37–39]. Unlike MPC, RL does not require a complete model of the building system. Instead, it adapts to evolving conditions, such as renovations or equipment upgrades. RL, a branch of machine learning, works by trial and error, identifying a policy that maps environmental states to actions. An RL agent chooses actions based on its policy, interacts with the system, and updates its strategy according to rewards received.

In HVAC use, the environment includes the building, its occupants, and the HVAC infrastructure. Many buildings already employ management systems (BMSs) that provide environmental data [41]. From a practical standpoint, fewer sensors and actuators are desirable, as they reduce cost and improve reliability [42,43]. Optimisation is driven by a reward function, often designed to reduce energy consumption or cost. However, comfort for occupants must also be safeguarded, since energy efficiency cannot come at the expense of well-being. Indoor environmental quality (IEQ) encompasses visual, acoustic, thermal, and air quality dimensions, with the last two having the greatest impact on health [44].

Comfort is frequently measured through thermal conditions and indoor air quality [44]. IAQ involves pollutants from indoor and outdoor sources, which may be physical, chemical, or biological [45]. Humidity and airflow also affect well-being. For example, low relative humidity can facilitate the spread of influenza, while moderate humidity supports comfort [46]. Health-focused strategies to reduce harmful exposures are therefore vital [44]. Thermal comfort, meanwhile, is defined as a person's satisfaction with their thermal surroundings [47]. Temperature is a dominant factor, with warmth

linked to improved respiratory health, particularly in children, whereas cold exposure can drive higher energy use but also produce indirect health benefits [48,49].

Conventional tools for assessing comfort, such as Predicted Mean Vote (PMV) and Predicted Percentage of Dissatisfied (PPD), are designed for large populations using simplified assumptions about metabolism or clothing [50,51]. These models often lack the precision to tailor comfort to individuals or small groups, which poses both technical and ethical challenges in human-centered applications [51].

Deploying RL-based HVAC controllers in practice may follow two approaches: static or dynamic. In static deployment, a pretrained policy is applied directly without further learning. This is inexpensive, computationally light, and works on low-cost hardware. However, static systems may underperform if the building changes, and periodic retraining may be necessary [42,52]. Dynamic deployment, on the other hand, continuously updates the control policy during operation. While it demands more computational resources and carries a risk of instability, it allows the system to adapt in real time, improving performance over time and reducing dependence on perfect models [42,52]. The choice between these approaches depends on the requirements and constraints of the building system in question.

3. Material and methods

The study's methodology is laid out in this section. Literature reviews and other sources, such as company websites, were analysed as part of the project. Case analysis and the incorporation of various ethical standards were also carried out. The next paragraphs elaborate on these points. Both Google and well-known building control company websites were combed through in search of commercially available AI-based building control apps. Similar opportunities arose through other mediums, such corporate marketing, for some of the featured market applications. A lot of products on the market claim to be "smart," but nobody ever really explains what that term means. For our assessment, we looked only at solutions that made it crystal obvious on their website that they leverage AI or ML. While solutions specific to the Finnish market were the primary emphasis, those accessible in other Nordic nations were also considered when appropriate. We have been going over papers in Section 4.2 that discuss the use of reinforcement learning controllers in real-world HVAC systems. The review comprised documents that had been published by the time the work was carried out in late 2023 and early 2024. Since artificial intelligence (AI) as a whole would have been too generic, we opted to focus on reinforcement learning instead due to its promising applications in smart controls. Cases from the lab and the actual world are both included. The initial searches were conducted with ScienceDirect using the following keywords: "Reinforcement learning," "HVAC," "deployment," or "real building." Articles that make it obvious in the abstract or section titles that a deployment is taking place will be filtered out. Articles discovered through many other means were also incorporated. Researchers looked for studies that addressed occupant health because it is a key factor in bringing AI-based controls into actual buildings. Reviewing previous work and commercial applications of Digital Twin technology in actual built environments (including health considerations) has supplemented the studies on RL controllers in Section 4.3. Building a digital twin and getting it up and running can give a data-rich platform that helps with RL-based solution creation, deployment, and runtime functionality. The majority of the content has been sourced from Google Scholar and ScienceDirect.

In this article, we examined three publications that dealt with actual developing AI scenarios by integrating the ethical AI principles found in numerous references (see section 5.1) into an ethical framework. By first organising the principles with the help of the Copilot AI program and then, with the help of an expert researcher's insights, we were able to integrate several ethical principles. First, a focus on people; second, a commitment to their safety and security; third, an examination of the relationship between ethics and social impact; fourth, safeguards for personal information and privacy; fifth, openness and oversight from humans; and sixth, responsibility for decision-making and oversight. Using this ethical paradigm, we then examined three practical AI applications for building automation. Three articles were chosen from Section 4.2. One article, "Occupant-centered real-time control of indoor temperature using deep learning algorithms," represents personalised thermal comfort control with physiological measurements in a laboratory setting. The other two articles are about energy usage optimisation with just building data and a non-laboratory version of personalised thermal comfort control, respectively.

4. AI-based smart control solutions in real buildings

4.1 AI-based market solutions

The main promotional benefits consist of lower heating costs, improved indoor air quality, and reduced carbon dioxide output. A variety of management and enhancement methods include error identification features. Heimar Andersen and colleagues [53] found that when exploring industry views on the challenges and motivations for adopting automatic fault detection and diagnosis in buildings and HVAC systems, the detection of faults is not a sellable item; instead, clients are more interested in acquiring the benefits it provides, like energy efficiency. This could clarify the reason apps are promoted with an emphasis on management and enhancement instead of flaw identification. Typically, numerous and significant benefits are claimed on the platforms of the applications. Still, it is unclear whether the benefits can be directly linked to artificial intelligence or if they could also be achieved through other technological approaches. In addition, the explanations were remarkably unclear. The particular technologies related to artificial intelligence used were not revealed. A noticeable enthusiasm for AI-driven building management and enhancement solutions can be observed in the Nordic region, highlighted by a significant array of recognised instances.

4.2 Improved education in the advancement of construction services strategies

Our scholarly investigation produced 16 publications that specifically outlined the evaluation of deep reinforcement learning (DRL) methods in practical HVAC scenarios. Two of the investigations [36,42] utilised experimental arrangements lacking real inhabitants, yet featured physical areas equipped with authentic heating devices, whereas the remaining two involved participants within the experimental setting. Merely two had set up devices in a home environment, whereas the rest were located in workplaces or communal areas. In nearly half of the cases, the controller was utilised exclusively for a designated area of the structure.

The length of the field trials differed greatly, spanning from 2 days to 1 year. The concise illustrations mainly evaluated the pretrained RL-controller's ability for efficient management in an actual structure, whereas the prolonged scenarios offered additional understanding of the systems' flexibility and continuous development in a genuine environment. Our evaluation concentrated solely on HVAC enhancement, with the algorithms primarily designed to reduce energy usage or expenses related to heating and cooling, all while maintaining thermal comfort for residents. One paper discussed the quality of air inside buildings and energy consumption, another concentrated on enhancing residential hot water systems, and a third analysed electricity consumption alongside heating. Table 2 summarises the key difficulties linked to the recognised reinforcement learning-driven smart control methods in real-world building implementations, which will be analysed and discussed in the following sections.

4.2.1 Creating the reinforcement learning regulator

Four point two point one point one. Measurements and functions. Sierla and colleagues [76] classified the actions in reinforcement learning for HVAC into three unique categories: 1) control signal, where the action serves as a direct command for an actuator sent through the programmable logic controller (PLC) of the building automation; 2) setpoint adjustment, in which the action is defined as a setpoint within the control loop of the building automation; and 3) planning, which pertains to actions linked to the building energy management system, generally ai This could be linked to very basic management examples. Heidari and colleagues [42] observed that the benefit of a target signal is that, although the RL controller operates as an overseeing controller, simply determining the target for a traditional controller, the effectiveness stays stable and robust, enabling the incorporation of user-specified targets.

A particular situation did not clearly correspond with the categories suggested by Sierla et al. [76], since Wei et al. [72] sought to save energy by offering users guidance on the best times for arriving and leaving, along with recommended locations for presence. Their financial reserves relied on a strategy that permitted each area to reduce its heating, ventilation, and air conditioning as well as lighting consumption when unoccupied, employing a machine learning approach to subtly

encourage energy conservation through changes in tenant habits. This could be seen as a broadening of the organising section.

The category of action represents a vital choice in the design of a reinforcement learning controller; however, the timing of control and the precision of output signals could also face limitations dictated by the physical system in use. Based on the utilised reinforcement learning method, the regulator might produce either continuous or discrete results. Instead of exclusively linking output clarity to algorithm selection, publications also reference elements like actuator increment size [41] and the feasibility of more precise adjustments. For example, when residents cannot detect slight changes in indoor warmth, heightened precision might result in conflicting comfort responses [50]. The lowest limit for the controller timing is set by the measurement duration, which varies from 5 seconds to 5 minutes, while the control time step generally lies between 5 and 30 minutes (refer to Table 2). Heidari and colleagues [42] support a 30-minute regulation period because of the slow reaction time of the hydronic heating system, while Chen and associates [34] noted that their practical test setup functioned more swiftly than the simulation, suggesting that the utilised 15-minute control interval might be shortened even more. Alongside management periods, undue adjustments have been reduced via the incentive mechanism by applying sanctions for repetitive changes in regulator behaviours [35].

4.2.1.2. Motivational system and resident contentment.

In most of the analysed situations, thermal comfort is simplified to a direct relationship with indoor temperature, for example, Refs. [36, 69, 70]. This approach is simple to implement for both models and real structures. While various somewhat different roles have been proposed, they all possess a particular target temperature or temperature range; the larger the difference between the measured temperature and this target, the more substantial the consequence enforced by the reward mechanism. Furthermore, enhanced thermal comfort measurements have been assessed in the studies. PPD is employed in the preliminary training stage when real thermal comfort information is unavailable [52,74], for setting up the individual comfort framework [50], or as a standard reference point [73]. Jung and colleagues [71] assessed both actual and generated thermal sensation ratings, subsequently translating them into the PPD index to measure thermal comfort as a definitive figure.

In real-world scenarios, structures are often occupied by many people with different tastes; therefore, several research efforts have concentrated on enhancing the personal satisfaction of the residents. They utilise tailored thermal comfort frameworks created from information gathered directly from residents' feedback through surveys, biological monitors, or a mix of the two. As an example, Lei and colleagues [50] used a tailored comfort matrix, set up with a preliminary comfort questionnaire for every individual, which was applied in the pretraining stage. The framework was regularly refreshed during the implementation phase. The writers note that the comfort framework has two benefits compared to direct occupant input: it enhances the learning process by reducing potentially conflicting responses and allows for feedback gathering at a lower frequency than the control is calculated. The ease drawback in the incentive formula is defined by the typical unease of the present residents. However, programs dependent on ongoing questionnaires might face respondent exhaustion, leading to participants being unable to reply reliably, thus generating less precise models [50,51]. As an illustration. Zhang and colleagues [52] opt to update their approach by allowing residents to offer feedback only when changes are required.

Portable devices offer a different approach to obtaining immediate responses from users. As an example, Jung and colleagues [71] evaluated the participants' physical activity by employing acceleration data, in addition to heart rate and skin temperature collected from a wristband. Nonetheless, gathering information through wearable biological devices from every individual in non-residential structures would be excessively expensive, unfeasible, and would provoke worries regarding data confidentiality. Consequently, it is essential to develop an all-encompassing thermal comfort framework that considers personal differences [77].

Another issue regarding occupant satisfaction is that most studies concentrate exclusively on one factor affecting thermal comfort, usually indoor temperature. Nonetheless, for instance. Jung and colleagues [71] along with Qin and associates [70] acknowledge this as a constraint and argue that further factors, including relative humidity and air speed, should be

included to provide a more comprehensive insight into occupant comfort. An & Chen [69] also focused on reducing particulate air pollution, specifically indoor PM_{2.5} levels, while maintaining thermal comfort and conserving energy. This was achieved by simultaneously operating a window, an air cleaner, and a cooling system.

Moreover, the equilibrium between energy efficiency and resident satisfaction is another aspect to consider when creating a reward mechanism. The reasoning behind the selected values among rival reward expressions was frequently insufficiently explained; however, they were commonly described as flexible factors, for example. Citations. [43, 50, 70]. A number of writers also aligned rival incentive phrases to a similar magnitude utilising past information, for instance. Citations. [35,74]. Chen and colleagues [34] noted that the emphasis on thermal comfort in relation to energy use might differ during times of occupancy compared to when spaces are unoccupied.

4.2.2 Initial Training

An essential benefit of RL-driven management is its autonomy from the requirement of developing and maintaining complex frameworks. Permitting the system to start fresh and independently explore possible actions might lead to unhappy residents or even harm to the structure, and attaining acceptable control results could take a considerable duration. To ensure a smooth start, the documents mainly rely on the prior training of the RL agent in a simulated setting before it is launched in the real world.

A variety of methods aimed at minimising pretraining efforts have been recorded. Zhang and colleagues [52] as well as Lei and associates [50] utilised automated calibration of physical models to synchronise quantifiable data through the use of Genetic Algorithms (GAs) in conjunction with Bayesian calibration or Gaussian Process-driven Bayesian optimisation. For example, additional sources. [33, 35, 36, 70] employed analytical frameworks based on data.

An & Chen [69] demonstrated via simulations that a controller developed in one environment could be applied in spaces with under a 20% difference in key parameters relative to the original and still exceed the effectiveness of the rule-based controller. Qin and colleagues [70] noted that transfer learning could potentially be utilised in neural networks within reinforcement learning, enabling the implementation of a pre-trained agent from an alternate situation in a fresh target setting; nonetheless, they did not perform any experiments regarding this application.

Heidari and colleagues [42] suggested that the agent should be pretrained with broadly applicable insights, employing differences in temperature reaction durations and occupancy patterns from various similar regions, instead of focussing on just one setting. Introducing the representative to various forms might reduce overfitting, allowing for quicker adjustment when applied in one structure. They claim that fluctuations in occupancy information are more essential in the training dataset, since the characteristics of the building will show less variation than the occupancy information when utilised in a particular scenario. Since the aim is to foster adaptable conduct, the foundational characteristics do not have to correspond with the particular example. Chen and colleagues [34] implemented pretraining; however, instead of using a refined model to develop a new strategy, they guided the agent based on the earlier strategy via imitation learning. This guideline is then executed in the physical structure, where it is improved digitally via a method of investigation that gradually refines its efficiency.

4.2.3. Virtual execution

The majority of publications seem to straightforwardly apply the trained regulator in the real construction. They might overlook enhancements to governance strategies due to short implementation periods. Nonetheless, a limited number have examined the complexities of implementation and virtual education more thoroughly. Y. Lei and colleagues [50] divided the implementation stage into two sections. As long as the replay buffer is filled with real-time experiences, the control strategy remains refined daily within the simulated setting. Automated model adjustment has been implemented through Gaussian Process-driven Bayesian Optimisation to enhance the precision of the virtual setting upon receiving new data. During the following phase, the representative acquires knowledge straight from the storage of past observations. Naug and colleagues [35] divided the digital method into two parts: the implementation cycle and the re-education cycle.

The offline re-education procedure entails a model based on data revising the structure dimensions, succeeded by the re-education of the fresh control strategy to correspond with the revised model. The process of retraining begins with evaluating and examining online performance. The execution cycle simply implements the operational strategy without participating in discovery. This method reduces the likelihood of safety concerns emerging from investigation. Nonetheless, it is crucial to acknowledge that in-person training allows for a greater sampling frequency, leading to faster reacquisition. Nonetheless, adjusting with inadequate information could result in overfitting and severe memory loss; therefore, they utilised Elastic Weight Consolidation as a method of regularisation. Zhang and colleagues [52] utilised the deployment/relearning loop framework in conjunction with automated model adjustment methods, such as Bayesian calibration and genetic algorithms.

No matter the implementation approach, HVAC control systems utilising reinforcement learning need to be carefully crafted to tackle various real-world issues, such as missing or unrealistic measurement data, handling variable changes in building operations, and making sure that the actions of the controller do not interfere with critical building functions to secure approval from facility managers [35,41]. Luo and colleagues [51] highlighted the importance of automated processing of collected data, including filling in short interruptions in sensor measurements using retrospective methods. They implemented safety evaluations to evaluate the security of actions suggested by the RL controller before execution, automatically deactivating the controller and permitting the conventional rule-based controller to take over control [41].

4.2.4 Confirmation stage

A further difficulty present in practical situations is that assessing the effectiveness of previous and current control systems is complex, as environmental factors and occupancy fluctuate between evaluations carried out at various times. The arrangement of the evaluative assessments is the main focus. A common approach consists of analysing successive time periods. As an illustration, Naug and colleagues [35] compared the results over a span of two years using various controller configurations, whereas Svetozarevic and associates [33] investigated consecutive weeks. Qin and colleagues [70] evaluated the RL controller against a standard rule-based controller by determining a reference time period with roughly equivalent outdoor temperature and relative humidity. Nonetheless, it is important to keep in mind that effectiveness can also be affected by modifications in structural conduct. This can be accomplished solely through prolonged evaluation periods when all essential actions from the prior management strategy are already in place. A suggested approach entails switching daily between the differing regulatory strategies, permitting the rival policies to encounter largely similar weather conditions while still allowing both managers to tackle daily and weekly load variations. In this case, intervals between the separate policies ought to be omitted from the analysis to prevent skewing of the findings caused by the policy change [41]. In the final stages, adjustments must be made to reduce the effects of varying external conditions. Naug and colleagues [35] utilised dynamic time warping to ensure the evaluation of similar situations regarding external temperature and relative humidity. Svetozarevic and colleagues [33] adjusted the data based on Heating Degree Days, whereas Luo and others [41] classified outdoor temperature and building load intervals, making comparisons only when both factors fell within identical ranges.

On the other hand, some writers grounded their examination in a juxtaposition of the results achieved with a modern real-world controller and those from a modelled baseline version. An and Chen [69] evaluated the simulated baseline controller alongside the simulated RL controller against real-world measurements, whereas Zhang et al. [52] proposed a Gaussian process-based method to create a baseline utilising historical data for a direct comparison of savings achieved with the RL controller.

4.3 Creating virtual replicas to enable the incorporation of AI-based management systems into real structures

A variety of comprehensive case analyses have recorded the implementation and assessment of digital twin technology in real-world buildings, such as the educational campus in Cambridge, UK [78], military installations in Italy [79], and hospitals in China [80]. This research offers an extensive summary of the integration and management of IoT data within structures to develop information models grounded in Building Information Modelling (BIM). Nonetheless, this study offers limited insights into the use of AI for performing analytical tasks on the collected information. The implementation

of virtual replica technologies in real structures has been recorded for educational and corporate environments in the sources. Nonetheless, there is a lack of information and proof concerning the application of AI or the educational potential of digital twins in the documented cases.

The approach to classifying the functional energy efficiency of structures via digital replicas is explored in Reference [83]. By leveraging data collection from the Internet of Things, a virtual representation was used to analyse the functioning efficiency of a university facility in Nicosia, Cyprus; nonetheless, artificial intelligence analysis for assessing building performance metrics was not utilised in this study. Manfren et al. [24] created regression models using data at different time intervals to improve the effectiveness of smart thermostatic radiator valves (TRVs) and gas absorption heat pumps (GAHPs), while also offering insights into the overall performance of the building in a case study of Hale Court sheltered housing in Portsmouth, UK. Nonetheless, the phrase “digital twin” was used more as an abstract idea than as a functional instrument for building energy monitoring and control.

In order to tackle the limitations posed by technology in the incorporation and duplication of the latest developments in BIM, IoT, and AI, the study mentioned in [25] suggested an open-source approach for the anticipatory oversight of CO₂ levels in structures, with the goal of enabling digital twin creators to build similar systems. An investigation was conducted regarding the academic structure of a university campus in Norway to confirm the developed semantically improved digital twin model, particularly emphasising the role of HVAC predictive upkeep [26]. Exploring the business uses of digital twin technology for structures, while heavily promoted as digital twin solutions by companies like Bosch and Siemens, the current research on the extent of AI implementation by these organisations is still unclear.

The concept of a digital replica is linked to the well-being and comfort of occupants, as outlined in Section 4.2.2. Currently, most studies utilising digital twin technologies for thermal comfort and energy efficiency in structures mainly focus on temperature and humidity information, overlooking other factors like air speed and radiant temperature [86]. However, virtual replicas could provide diverse information to improve individualised temperature comfort frameworks. The layout and closeness of individuals to different elements of the structure, like openings, air circulators, air distribution devices, or warmth sources, can affect how comfortable individuals feel and ought to be included in comfort frameworks [87]. Illustrating the complex and interconnected relationships between construction systems and the well-being and satisfaction of occupants is challenging [88]. Graph-based neural architectures, a distinct type of neural networks utilising graph structures, could offer a means for examining intricate datasets [87].

Unlike production industries where virtual replicas have been successfully utilised, the use of virtual replicas in the constructed surroundings encounters various challenges. They relate to specific building characteristics, encompassing measurements, along with differences in style, on-location assembly, and the integration of various systems. Structures possess an extended functional duration, requiring attention to deteriorating edifices and related frameworks. Consequently, challenges in adopting digital twin technology within structures could relate to the upkeep costs linked to Information and Communication Technology (ICT) or IoT systems, encompassing software and hardware enhancements essential for the continuous functioning of AI/ML algorithms. There are no uniform approaches available for creating and executing digital replicas in structures. The complexity of information exchange in structures poses a major obstacle to the broad adoption of virtual replicas, as various participants and data owners are involved during the management of a building's life cycle, encompassing construction, property owners, equipment vendors, and service providers.

When put into action, a virtual replica of a structure provides an extensive and evolving information system that artificial intelligence applications can leverage, first during the construction stage prior to the establishment of a tangible edifice, and later during the functioning phase, to enhance building management and resident satisfaction instantaneously. In the construction stage, a digital model of the structure, including information on operational activities, resident actions, and environmental factors, can be employed by management algorithms, like reinforcement learning, to improve their decision-making frameworks and discover the best policies and management approaches. Throughout the active period, an extensive virtual replica of the structure and its components (including climate control and illumination) employing near-instantaneous information and cohesive analysis would enable flexibly enhanced management techniques. These

approaches would adapt to varying construction circumstances, usage trends, and energy requirements, with intelligent agents consistently observing the digital replica's condition and predictions. Additionally, in complex construction settings defined by multiple interconnected systems and data-holding participants, creating a digital replica can provide a unified platform (throughout both the design and operational stages) to enhance interactions among different AI entities, promoting their capacity to work together, negotiate, and achieve common objectives.

4.4 Ethical issues in implementing AI-based control and optimization solutions into real buildings

This study finds that ethical issues are often overlooked when suggesting AI applications for real buildings. Ethics in this context focus on ensuring AI is human-centered, prioritising safety, reliability, security, sustainability, and societal impact. At the same time, concerns about privacy and data protection are central to how data is gathered and managed. Ethical practice requires transparency, supported by clear governance structures and accountability mechanisms. AI systems should remain impartial and actively work to reduce bias and discrimination by recognising human diversity.

This chapter sets out the ethical aspects that should guide the design and implementation of AI in buildings. It introduces the principles of ethical AI and proposes a framework for assessing practical applications. The EU has stressed that digitalisation and growing data use bring major challenges, particularly regarding cybersecurity, privacy, and ethics. Yussuf and Asfour [13] noted that data validation and protection are critical in applying AI in buildings, as faulty or missing inputs can generate illogical controller behaviour [41]. Incorporating ethics systematically into AI design remains essential [89], yet applying high-level principles during software development is still a major challenge despite existing standards [90].

Kop et al. [91] proposed a framework for responsible technology adoption under three themes: Protection, Engagement, and Progress. Protection focuses on safeguarding information and reducing risks. Engagement calls for openness, inclusion, and diversity. Progress highlights sustainability, accountability, and stakeholder involvement. The EU's AI Act, finalised in February 2024, builds on similar ideas by categorising AI systems according to risk [29]. The Act prohibits high-risk practices such as social scoring, manipulative behavioural techniques, and real-time biometric surveillance. It also demands transparency: systems must identify AI-generated content, prevent illegal outputs, and disclose training data sources. GDPR adds further requirements, including fairness, purpose limitation, data minimisation, and accuracy.

Once applied in buildings, the AI Act will enforce this risk-based structure, ensuring compliance and oversight. The EU's Ethics Guidelines introduced the idea of Trustworthy AI, while UNESCO issued ten global principles for responsible and sustainable AI. Responsible Research and Innovation (RRI) in Europe complements this, stressing foresight, inclusivity, reflexivity, and responsiveness in technology adoption [96].

Building on these standards, this work evaluates three AI-based building control solutions: (1) deep reinforcement learning for energy management [35], (2) multivariate occupant-focused control [50], and (3) real-time temperature regulation with deep learning [71]. Our framework groups ethical concerns into six areas:

1. **Human-Centric AI** – ensuring proportionality, fairness, and non-discrimination while preventing harm. Systems should support comfort, balance efficiency with well-being, and remain responsive to diverse occupants.
2. **Safety, Security, and Reliability** – protecting against cyberattacks and failures through robust testing and validation.
3. **Sustainability and Societal Impact** – designing AI to deliver environmental benefits, long-term social value, and economic stability.
4. **Privacy and Data Protection** – safeguarding personal information with secure storage, ownership clarity, and consent mechanisms.

5. **Transparency and Human Oversight** – enabling explainable decisions, user understanding, and education to build trust in AI systems.
6. **Governance and Accountability** – assigning responsibility, ensuring compliance, and fostering collaboration among governments, businesses, and civil society.

All six areas are essential, with categories 2, 4, and 5 closely tied to technology development, while 1, 3, and 6 relate more directly to people and society. Together, they ensure AI in buildings supports human welfare and living conditions.

Adaptive AI controllers have already been tested in limited building scenarios, often supported by existing BMS data for direct control. However, challenges remain, particularly in handling missing or erroneous data, which requires preprocessing and safety mechanisms [41]. Training is a further barrier: most experiments rely on customised simulations, raising doubts about scalability and transferability.

Field trials also tend to be short-term, leaving questions about long-term effectiveness and adaptability during building upgrades or changes in occupant behaviour. Most focus only on temperature control, ignoring other aspects of facility management. Comfort models often depend on detailed user feedback, which is impractical for large populations. Wearable tracking devices could improve data collection, but they raise privacy concerns. Evidence from the UK shows that people perceive smart home devices as riskier and less beneficial than conventional alternatives [97].

To overcome data scarcity, researchers are applying methods such as transfer learning. Gao et al. [98] used field data from similar climates to predict comfort, while Li et al. [99] applied ASHRAE datasets to workplace comfort prediction. These approaches reduce survey requirements while maintaining privacy. Cohort Comfort Models (CCM), as proposed by Quintana et al. [100], use group-level similarities to estimate comfort with minimal prior data, offering a path to scalable, customised systems.

Recent debates also call for standardising digital twin definitions, IoT protocols, and data-sharing practices to accelerate real-world adoption. Sharma et al. [101] proposed evaluation criteria for digital twins, including technical readiness, data reliability, and AI integration. Yet, the ethical dimension remains underexplored. As Vakkuri et al. [102] emphasise, AI ethics is still developing. Salo-Pontinen and Saariluoma argue that national policies too often prioritise technical goals over user-centered concerns. Scholars such as Ahmad et al. [104] highlight the importance of fairness, privacy, and cultural perspectives in AI for smart cities. Sharpe [107] raises concerns about inhabitants effectively becoming test subjects when their private data is collected for experimentation.

In sum, while reinforcement learning and digital twin technologies show strong potential, the literature demonstrates that ethical concerns—privacy, fairness, transparency, and accountability—must be embedded at every stage. Without this, the adoption of AI in buildings risks eroding trust, undermining safety, and failing to meet broader societal needs.

5. Conclusions

This research set out to build a foundation for applying artificial intelligence-driven control and optimisation systems in real buildings. The evaluation reviewed the present state of such systems from four viewpoints: 1) solutions currently offered in the market, 2) obstacles reported in scientific case studies, 3) the role of digital twins in deployment, and 4) ethical issues linked to implementation. A summary of the key findings is provided in the figure.

With respect to RQ1, which focused on market-ready AI solutions, findings indicate that several products are already available in the Nordic region, most of which target improvements in space heating for residential buildings. However, these systems usually function as “black boxes,” offering little transparency or technical detail about the underlying AI methods employed.

Since descriptions of existing commercial tools often lacked clarity, RQ2 shifted attention to experimental reinforcement learning (RL) applications reported in the literature. Here, the primary barrier to wider adoption was ensuring safe adaptation when algorithms adjust to new conditions and verifying their scalability. Trials documented in published studies were typically short in duration and limited in scope. As a result, more comprehensive investigations are needed to determine how these methods affect whole-building performance over time.

RQ3 dealt with the contribution of digital twins to AI-enabled building control. Digital twins have the potential to provide rich, integrated platforms that improve development, testing, and operation. Nevertheless, little evidence exists that the data collected in practice has been fully exploited to create genuine digital twins as described in theory. Advancing this field requires clearer definitions of maturity levels and the development of validation standards specific to buildings. These should cover IoT integration, data-sharing practices, cybersecurity, and AI performance testing.

Finally, RQ4 addressed ethics. Reports of RL-based building control largely omitted systematic ethical assessments. Given RL's limited transparency in decision-making, ensuring human oversight and explainability is critical. This will be important to gain the trust of building owners and operators and to avoid negative impacts on comfort, safety, or equipment. Based on this, the ethical analysis recommends six guiding principles for AI in buildings: 1) Human-Centricity, 2) Safety, Security, and Reliability, 3) Sustainability and Societal Impact, 4) Privacy and Data Protection, 5) Human Oversight and Transparency, and 6) Governance and Accountability.

References

- [1] European Commission, A Europe fit for the digital age. https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age_en, 2024. (Accessed 16 April 2024).
- [2] European Commission, Excellence and trust in artificial intelligence. https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/excellence-and-trust-artificial-intelligence_en, 2021. (Accessed 16 April 2024).
- [3] OECD.AI, OECD's live repository of AI strategies & policies, (n.d.). <https://oecd.ai/en/dashboards/overview> (accessed April 15, 2024).
- [4] Statista. Market Forecast, Artificial Intelligence - Europe, 2024. <https://www.statista.com/outlook/tmo/artificial-intelligence/europe> (Accessed 16 April 2024).
- [5] European Commission, Digitalisation of the energy systems, (n.d.). https://energy.ec.europa.eu/topics/energy-systems-integration/digitalisation-energy-system_en (accessed February 13, 2024).
- [6] Statista, IoT connected devices worldwide 2019-2030. <https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/>, 2024. (Accessed 7 May 2024).
- [7] S. Paiho, N. Wessberg, M. Dubovik, R. Lavikka, S. Naumer, Twin transition in the built environment – policy mechanisms, technologies and market views from a cold climate perspective, *Sustain. Cities Soc.* 98 (2023) 104870, <https://doi.org/10.1016/j.scs.2023.104870>.
- [8] S.K. Baduge, S. Thilakarathna, J.S. Perera, M. Arashpour, P. Sharafi, B. Teodosio, A. Shringi, P. Mendis, Artificial intelligence and smart vision for building and construction 4.0: machine and deep learning methods and applications, *Autom. ConStruct.* 141 (2022) 104440, <https://doi.org/10.1016/J.AUTCON.2022.104440>.
- [9] S. Tuomela, Smart Home Energy Technologies: Adoption, User Experience and Energy Saving Potential, University of Oulu, 2022. <https://oulurepo.oulu.fi/handle/10024/36951>.

- [10] G. Halhoul Merabet, M. Essaïdi, M. Ben Haddou, B. Qolomany, J. Qadir, M. Anan, A. Al-Fuqaha, M.R. Abid, D. Benhaddou, Intelligent building control systems for thermal comfort and energy-efficiency: a systematic review of artificial intelligence-assisted techniques, *Renew. Sustain. Energy Rev.* 144 (2021), <https://doi.org/10.1016/j.rser.2021.110969>.
- [11] L. Zhang, J. Wen, Y. Li, J. Chen, Y. Ye, Y. Fu, W. Livingood, A review of machine learning in building load prediction, *Appl. Energy* 285 (2021) 116452, <https://doi.org/10.1016/j.apenergy.2021.116452>.
- [12] Z. Wang, J. Liu, Y. Zhang, H. Yuan, R. Zhang, R.S. Srinivasan, Practical issues in implementing machine-learning models for building energy efficiency: moving beyond obstacles, *Renew. Sustain. Energy Rev.* 143 (2021) 110929, <https://doi.org/10.1016/j.rser.2021.110929>.
- [13] R.O. Yussuf, O.S. Asfour, Applications of artificial intelligence for energy efficiency throughout the building lifecycle: an overview, *Energy Build.* 305 (2024) 113903, <https://doi.org/10.1016/J.ENBUILD.2024.113903>.
- [14] C. Lu, S. Li, Z. Lu, Building energy prediction using artificial neural networks: a literature survey, *Energy Build.* 262 (2022) 111718, <https://doi.org/10.1016/j.enbuild.2021.111718>.
- [15] K. Alanne, S. Sierla, An overview of machine learning applications for smart buildings, *Sustain. Cities Soc.* 76 (2022) 103445, <https://doi.org/10.1016/J.SCS.2021.103445>.
- [16] S.H. Khajavi, N.H. Motlagh, A. Jaribion, L.C. Werner, J. Holmstrom, Digital Twin: vision, benefits, boundaries, and creation for buildings, *IEEE Access* 7 (2019) 147406–147419, <https://doi.org/10.1109/ACCESS.2019.2946515>.
- [17] J. Zhang, L. Zhao, G. Ren, H. Li, X. Li, Special issue “digital twin technology in the AEC industry,”, *Adv. Civ. Eng.* 2020 (2020) <https://doi.org/10.1155/2020/8842113>.
- [18] P. de Wilde, Building performance simulation in the brave new world of artificial intelligence and digital twins: a systematic review, *Energy Build.* 292 (2023) 113171, <https://doi.org/10.1016/J.ENBUILD.2023.113171>.
- [19] M. Mazhar Rathore, S. Attique Shah, D. Shukla, E. Bentafat, S. Bakiras, The role of AI, machine learning, and Big data in digital twinning: a systematic literature review, challenges, and Opportunities, *IEEE Access* 9 (2021) 32030–32052, <https://doi.org/10.1109/ACCESS.2021.3060863>.
- [20] J. Zhao, H. Feng, Q. Chen, B. Garcia de Soto, Developing a conceptual framework for the application of digital twin technologies to revamp building operation and maintenance processes, *J. Build. Eng.* 49 (2022) 104028, <https://doi.org/10.1016/J.JOBE.2022.104028>.
- [21] A. Clausen, K. Arendt, A. Johansen, F.C. Sangogboye, M.B. Kjærgaard, C.T. Veje, B.N. Jørgensen, A digital twin framework for improving energy efficiency and occupant comfort in public and commercial buildings, *Energy Inform.* 4 (2021) 1–19, <https://doi.org/10.1186/S42162-021-00153-9/FIGURES/11>.
- [22] C. Boje, A. Guerriero, S. Kubicki, Y. Rezgui, Towards a semantic construction digital twin: directions for future research, *Autom. Construct.* 114 (2020) 103179, <https://doi.org/10.1016/J.AUTCON.2020.103179>.
- [23] C. Blume, S. Blume, S. Thiede, C. Herrmann, Data-driven digital twins for technical building services operation in factories: a cooling tower case study, *J. Manuf. Mater. Process.* 4 (2020) 97, <https://doi.org/10.3390/JMMP4040097>.

- [24] M. Manfren, P.A. James, V. Aragon, L. Tronchin, Lean and interpretable digital twins for building energy monitoring – a case study with smart thermostatic radiator valves and gas absorption heat pumps, *Energy AI* 14 (2023) 100304, <https://doi.org/10.1016/J.EGYAI.2023.100304>.
- [25] A. Arsiwala, F. Elghaish, M. Zoher, Digital twin with Machine learning for predictive monitoring of CO2 equivalent from existing buildings, *Energy Build.* 284 (2023) 112851, <https://doi.org/10.1016/J.ENBUILD.2023.112851>.
- [26] H.H. Hosamo, P.R. Svennevig, K. Svidt, D. Han, H.K. Nielsen, A Digital Twin predictive maintenance framework of air handling units based on automatic fault detection and diagnostics, *Energy Build.* 261 (2022) 111988, <https://doi.org/10.1016/J.ENBUILD.2022.111988>.
- [27] Nordic Innovation, Urgent need for ethical AI-initiatives. https://www.nordicinnovation.org/news/urgent-need-ethical-ai-initiatives?utm_source=Nordic+Innovation+Newsletter&utm_campaign=58dd667850-newsletter-02-2022_COPY_01&utm_medium=email&utm_term=0_1629348415-58dd667850-84183003, 2024. (Accessed 11 April 2024).
- [28] Nordic Ethical AI Expert Group, Policy recommendations for ethical and responsible AI. <https://norden.diva-portal.org/smash/get/diva2:1863330/FULLTEXT01.pdf>, 2024.
- [29] European Parliament, EU AI Act: first regulation on artificial intelligence. <https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>, 2023. (Accessed 5 April 2024).
- [30] UNESCO, Ethical Impact Assessment: a Tool of the Recommendation on the Ethics of Artificial Intelligence, UNESCO, 2023, <https://doi.org/10.54678/YTSA7796>.
- [31] UNESCO, Readiness Assessment Methodology: a Tool of the Recommendation on the Ethics of Artificial Intelligence, UNESCO, 2023, <https://doi.org/10.54678/YHAA4429>.
- [32] E. Prem, From ethical AI frameworks to tools: a review of approaches, *AI Ethics* (2023) 699–716, <https://doi.org/10.1007/S43681-023-00258-9>, 2023 33 3.
- [33] B. Svetozarevic, C. Baumann, S. Muntwiler, L. Di Natale, M.N. Zeilinger, P. Heer, Data-driven control of room temperature and bidirectional EV charging using deep reinforcement learning: simulations and experiments, *Appl. Energy* 307 (2022) 118127, <https://doi.org/10.1016/J.APENERGY.2021.118127>.
- [34] B. Chen, Z. Cai, M. Bergés, Gnu-RL: a practical and scalable reinforcement learning solution for building HVAC control using a differentiable MPC policy, *Front. Built Environ.* 6 (2020) 562239, <https://doi.org/10.3389/FBUIL.2020.562239/BIBTEX>.
- [35] A. Naug, M. Quinones-Grueiro, G. Biswas, Deep reinforcement learning control for non-stationary building energy management, *Energy Build.* 277 (2022), <https://doi.org/10.1016/j.enbuild.2022.112584>.
- [36] C. Blad, S. Bøgh, C. Kallesøe, P. Raftery, A laboratory test of an offline-trained multi-agent reinforcement learning algorithm for heating systems, *Appl. Energy* 337 (2023), <https://doi.org/10.1016/j.apenergy.2023.120807>.
- [37] X. Zhong, Z. Zhang, R. Zhang, C. Zhang, End-to-End deep reinforcement learning control for HVAC systems in office buildings, *Design* 6 (2022) 52, <https://doi.org/10.3390/DESIGNS6030052>.

- [38] P. Lissa, M. Schukat, E. Barrett, Transfer learning applied to reinforcement learning-based HVAC control, *SN Comput. Sci.* 1 (2020) 1–12, <https://doi.org/10.1007/S42979-020-00146-7/TABLES/4>.
- [39] X. Yuan, Y. Pan, J. Yang, W. Wang, Z. Huang, Study on the application of reinforcement learning in the operation optimization of HVAC system, *Build. Simulat.* 14 (2021) 75–87, <https://doi.org/10.1007/S12273-020-0602-9/METRICS>.
- [40] [R.S. Sutton, A.G. Barto, Reinforcement Learning. An Introduction, second ed., The MIT Press, 2018.](#)
- [41] J. Luo, C. Paduraru, O. Voicu, Y. Chervonyi, S. Munns, J. Li, C. Qian, P. Dutta, J.Q. Davis, N. Wu, X. Yang, C.-M. Chang, T. Li, R. Rose, M. Fan, H. Nakhost, T. Liu, B. Kirkman, F. Altamura, L. Cline, P. Tonker, J. Gouker, D. Uden, W.B. Bryan, J. Law, D. Fatiha, N. Satra, J. Rothenberg, M. Waraich, M. Carlin, S. Tallapaka, S. Witherspoon, D. Parish, P. Dolan, C. Zhao, D.J. Mankowitz, Controlling commercial cooling systems using reinforcement learning, *ArXiv* (2022) 27. <https://arxiv.org/abs/2211.07357v2>. (Accessed 11 June 2024).
- [42] A. Heidari, D. Khovalyg, DeepValve: development and experimental testing of a Reinforcement Learning control framework for occupant-centric heating in offices, *Eng. Appl. Artif. Intell.* 123 (2023) 106310, <https://doi.org/10.1016/J.ENGAPPAI.2023.106310>.
- [43] H. Kazmi, F. Mehmood, S. Lodeweyckx, J. Driesen, Gigawatt-hour scale savings on a budget of zero: deep reinforcement learning based optimal control of hot water systems, *Energy* 144 (2018) 159–168, <https://doi.org/10.1016/J.ENERGY.2017.12.019>.
- [44] J. Cai, J. Chen, Y. Hu, S. Li, Q. He, Digital twin for healthy indoor environment: a vision for the post-pandemic era, *Front. Eng. Manag.* 10 (2023) 300–318, <https://doi.org/10.1007/S42524-022-0244-Y>.
- [45] S. Altomonte, J. Allen, P.M. Bluyssen, G. Brager, L. Heschong, A. Loder, S. Schiavon, J.A. Veitch, L. Wang, P. Wargocki, Ten questions concerning well-being in the built environment, *Build. Environ.* 180 (2020) 106949, <https://doi.org/10.1016/J.BUILDENV.2020.106949>.
- [46] P. Wolkoff, Indoor air humidity, air quality, and health - an overview, *Int. J. Hyg Environ. Health* 221 (2018) 376–390, <https://doi.org/10.1016/J.IJHEH.2018.01.015>.
- [47] [ASHRAE, ANSI/ASHRAE Standard 55-2023, Thermal Environmental Conditions for Human Occupancy, 2023.](#)
- [48] Z. Liu, X. Zhang, Y. Sun, Y. Zhou, Advanced controls on energy reliability, flexibility and occupant-centric control for smart and energy-efficient buildings, *Energy Build.* 297 (2023) 113436, <https://doi.org/10.1016/J.ENBUILD.2023.113436>.
- [49] J. Ige, P. Pilkington, J. Orme, B. Williams, E. Prestwood, D. Black, L. Carmichael, G. Scally, The relationship between buildings and health: a systematic review, *J. Public Health* 41 (2019) e121–e132, <https://doi.org/10.1093/PUBMED/FDY138>.
- [50] Y. Lei, S. Zhan, E. Ono, Y. Peng, Z. Zhang, T. Hasama, A. Chong, A practical deep reinforcement learning framework for multivariate occupant-centric control in buildings, *Appl. Energy* 324 (2022), <https://doi.org/10.1016/j.apenergy.2022.119742>.
- [51] J. Kim, Y. Zhou, S. Schiavon, P. Raftery, G. Brager, Personal comfort models: predicting individuals' thermal preference using occupant heating and cooling behavior and machine learning, *Build. Environ.* 129 (2018) 96–106, <https://doi.org/10.1016/J.BUILDENV.2017.12.011>.

- [52] Z. Zhang, A. Chong, Y. Pan, C. Zhang, K.P. Lam, Whole building energy model for HVAC optimal control: a practical framework based on deep reinforcement learning, *Energy Build.* 199 (2019) 472–490, <https://doi.org/10.1016/j.enbuild.2019.07.029>.
- [53] K. Heimar Andersen, S. Pommerencke Melgaard, H. Johra, A. Marszal-Pomianowska, R. Lund Jensen, P. Kvols Heiselberg, Barriers and drivers for implementation of automatic fault detection and diagnosis in buildings and HVAC systems: an outlook from industry experts, *Energy Build.* 303 (2024) 113801, <https://doi.org/10.1016/J.ENBUILD.2023.113801>.
- [54] Capacity.io, Make your heat pump smart. <https://www.capacity.io/>, 2023. (Accessed 1 December 2023).
- [55] Leanheat, Smart building control and maintenance powered by IoT. <https://leanheat.com/>, 2023. (Accessed 11 December 2023).
- [56] Fourdeg, Energy efficient district heating. <https://www.fourdeg.com/>, 2022. (Accessed 11 December 2023).
- [57] MeshWorks, Premier Industrial IoT & Automation Solutions, (n.d.). <https://www.meshworkswireless.com/> (accessed December 12, 2023).
- [58] Themo, Choose Affordable and clean electric heating. <https://www.themo.io/>, 2022. (Accessed 12 December 2023).
- [59] OptiWatti, Aurinkosähkön aliyohjaukseen, (n.d.). <https://www.optiwatti.fi/mika-on-optiwatti/aurinkosahko-alyohjaukseen/> (accessed December 12, 2023). [60] OptiWatti, More comfort and cost savings with electric heating, (n.d.). <https://www.optiwatti.com/> (accessed December 12, 2023).
- [61] Ouman, Tekoaly auttaa taloyhtiön öitön sähkönkäytön optimointia kaukolämpöä ja lämmityksen energiakustannusten vähentämiseksi. <https://ouman.fi/tekoaly-auttaa-taloyhtioita-saastamaan-kaukolampon-iso-buumi-kaynnissa/>, 2023. (Accessed 20 December 2023).
- [62] Riksborgen, Sänk varmakostnaden med smart vvarmeoptimering. https://www.riksbyggen.se/fastighetsforvaltning/fastighetsdrift/varmeoptimering/?utm_source=Di.se&utm_medium=native&utm_content=di-det-har-levererat-exit2&utm_campaign=forvaltning-energiboost-v45-49-2022, 2024. (Accessed 17 January 2024).
- [63] Schneider Electric, Schneider Electricin Tekoalioptimointipalvelu. https://go.schneider-electric.com/FI_202312_BMS_AI_yhteydenotot.html?source=Website&Detail=BMS_AI_Contact_2023_FI, 2024. (Accessed 23 January 2024).
- [64] Schneider Electric Suomi, Optimization AI tuo tekoalyn ja koneoppimisen voiman kiinteistön onhallintaan. <https://www.se.com/fi/fi/about-us/newsroom/news/press-releases/optimization-ai-tuo-tekoalyn-ja-koneoppimisen-voiman-kiinteistön-onhallintaan-6580085a5a2d6984410799f5>, 2023. (Accessed 23 January 2024).
- [65] Nuuka Solutions, Nuuka - Tekoaly ja diagnostiikka ohjelmistoratkaisut tulevaisuuden rakennuksille. Tiedotteen julkaisu, (n.d.). <https://www.nuuka.com/fi/home> (accessed January 23, 2024).
- [66] L&T, Kuopiossa kehitettiin koneoppiva ja ennakoiva kiinteistöautomaatio | Kiinteistöjen ot. <https://lassikko.lt.fi/suomen-kuopiossa-kehitetään-huomisenkestävää-energiatietoa-smartti-automaatiin>, 2021. (Accessed 1 February 2024).

- [67] L&T, Asiakkaamme Technopolis palkittiin Vuoden Energianerokas 2023 -tunnustuksella Smartti Automaatio -energiaohjausjärjestelmän pilotoinnista. " <https://www.lt.fi/fi/media/tiedotteet/asiakkaamme-technopolis-palkittiin-vuoden-energianerokas-2023-tunnustuksella>, 2023. (Accessed 1 February 2024).
- [68] L&T, Smartti Automaatio - L&T. <https://www.lt.fi/fi/smartti>, 2024. (Accessed 1 February 2024).
- [69] Y. An, C. Chen, Energy-efficient control of indoor PM2.5 and thermal comfort in a real room using deep reinforcement learning, *Energy Build.* 295 (2023), <https://doi.org/10.1016/j.enbuild.2023.113340>.
- [70] H. Qin, Z. Yu, T. Li, X. Liu, L. Li, Energy-efficient heating control for nearly zero energy residential buildings with deep reinforcement learning, *Energy* 264 (2023), <https://doi.org/10.1016/j.energy.2022.126209>.
- [71] S. Jung, J. Jeoung, T. Hong, Occupant-centered real-time control of indoor temperature using deep learning algorithms, *Build. Environ.* 208 (2022) 108633, <https://doi.org/10.1016/J.BUILDENV.2021.108633>.
- [72] P. Wei, S. Xia, R. Chen, J. Qian, C. Li, X. Jiang, A deep-reinforcement-learning-based recommender system for occupant-driven energy optimization in commercial buildings, *IEEE Internet Things J.* 7 (2020) 6402–6413, <https://doi.org/10.1109/JIOT.2020.2974848>.
- [73] C. Zhang, A Bio-Sensing and Reinforcement Learning Control System for Personalized Thermal Comfort and Energy Efficiency, Carnegie Mellon University, 2019. https://kithub.cmu.edu/articles/thesis/A_Bio-sensing_and_Reinforcement_Learning_Control_System_for_Personalized_Thermal_Comfort_and_Energy_Efficiency/9934130.
- [74] Z. Zhang, K. Poh Lam, Practical implementation and evaluation of deep reinforcement learning control for a radiant heating system, in: 5th ACM Int. Conf. Syst. Energy-Efficient Built Environ. (BuildSys 2018), 2018, pp. 148–157, <https://doi.org/10.1145/3276774.3276775>. Shenzhen, China, ACM.
- [75] G.T. Costanzo, S. Iacovella, F. Ruelens, T. Leurs, B.J. Claessens, Experimental analysis of data-driven control for a building heating system, *Sustain. Energy, Grids Networks* 6 (2016) 81–90, <https://doi.org/10.1016/j.segan.2016.02.002>.
- [76] S. Sierla, H. Ihasalo, V. Vyatkin, A review of reinforcement learning applications to control of heating, ventilation and air conditioning systems, *Energies* 15 (2022), <https://doi.org/10.3390/en15103526>.
- [77] S. Moon, S.S. Kim, B. Choi, Comparative feasibility study of physiological signals from wristband-type wearable sensors to assess occupants' thermal comfort, *Energy Build.* 308 (2024) 114032, <https://doi.org/10.1016/J.ENBUILD.2024.114032>.
- [78] Q. Lu, A.K. Parlikad, P. Woodall, G.D. Ranasinghe, X. Xie, Z. Liang, E. Konstantinou, J. Heaton, J. Schooling, Developing a digital twin at building and city levels: case study of west Cambridge campus, *J. Manag. Eng.* 36 (2020), [https://doi.org/10.1061/\(ASCE\)ME.1943-5479.0000763](https://doi.org/10.1061/(ASCE)ME.1943-5479.0000763).
- [79] P. Marchione, F. Ruperto, Prototyping a digital twin – a case study of a ‘U-shaped’ military building, *Int. J. Energy Prod. Manag.* 7 (2022) 83–94, <https://doi.org/10.2495/EQ-V7-N1-83-94>.
- [80] Y. Peng, M. Zhang, F. Yu, J. Xu, S. Gao, Digital twin hospital buildings: an Exemplary case study through continuous lifecycle integration, *Adv. Civ. Eng.* 2020 (2020), <https://doi.org/10.1155/2020/8846667>.
- [81] L.C. Tagliabue, F.R. Cecconi, S. Maltese, S. Rinaldi, A.L.C. Ciribini, A. Flammini, Leveraging digital twin for sustainability assessment of an educational building, *Sustainability* 13 (2021), <https://doi.org/10.3390/SU13020480>.

- [82] A. Zaballos, A. Briones, A. Massa, P. Centelles, V. Caballero, A smart campus' digital twin for sustainable comfort monitoring, *Sustainability* 12 (2020) 9196, <https://doi.org/10.3390/SU12219196>.
- [83] P. Spudys, N. Afxentiou, P.Z. Georgali, E. Klumbyte, A. Jurelionis, P. Fokaides, Classifying the operational energy performance of buildings with the use of digital twins, *Energy Build.* 290 (2023) 113106, <https://doi.org/10.1016/J.ENBUILD.2023.113106>.
- [84] Bosch, Digital Building Twins | Bosch Energy and Building Solutions Global, (n.d.). <https://www.boschbuildingsolutions.com/xc/en/news-and-stories/digital-building-twins/> (accessed April 15, 2024).
- [85] Siemens, Building Twin - Siemens Global, (n.d.). <https://www.siemens.com/global/en/products/buildings/digital-building-lifecycle/building-twin.html> (accessed April 15, 2024).
- [86] V.A. Arowoiya, R.C. Moehler, Y. Fang, Digital twin technology for thermal comfort and energy efficiency in buildings: a state-of-the-art and future directions, *Energy Built Environ* 5 (2024) 641–656, <https://doi.org/10.1016/J.ENBENV.2023.05.004>.
- [87] M.M. Abdelrahman, A. Chong, C. Miller, Personal thermal comfort models using digital twins: preference prediction with BIM-extracted spatial–temporal proximity data from Build2Vec, *Build. Environ.* 207 (2022) 108532, <https://doi.org/10.1016/J.BUILDENV.2021.108532>.
- [88] V.M. Gnecco, F. Vittori, A.L. Pisello, Digital twins for decoding human-building interaction in multi-domain test-rooms for environmental comfort and energy saving via graph representation, *Energy Build.* 279 (2023) 112652, <https://doi.org/10.1016/J.ENBUILD.2022.112652>.
- [89] P. Brey, Brandt Dainow, Ethics by Design for Artificial Intelligence, *AI Ethics*, vol. 1, 2023, pp. 1–13, <https://doi.org/10.1007/S43681-023-00330-4>.
- [90] V. Vakkuri, Implementing AI Ethics in Software Development, University of Jyväskylä, 2022. <https://jyx.jyu.fi/handle/123456789/81133>.
- [91] M. Kop, M. Aboy, E. De Jong, U. Gasser, T. Minssen, I.G. Cohen, M. Brongersma, T. Quintel, L. Floridi, R. Laflamme, Ten principles for responsible quantum innovation, *Quantum Sci. Technol.* 9 (2024) 035013, <https://doi.org/10.1088/2058-9565/AD3776>.
- [92] Regulation (EU), 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), *Off. J. Eur. Union* (2016) 1–88. <https://eur-lex.europa.eu/eli/reg/2016/679/>. (Accessed 5 April 2024).
- [93] Panel for the Future of Science and Technology (STOA), The impact of the general data protection regulation (GDPR) on artificial intelligence. <https://doi.org/10.2861/293>, 2020.
- [94] The European Commission's High-Level Expert Group on Artificial Intelligence, The Assessment list for Trustworthy artificial intelligence (ALTAI) for self assessment. <https://ec.europa.eu/digital-single-market/en/news/assessment-list-trustworthy-artificial-intelligence-altai-self-assessment>, 2020.
- [95] UNESCO, Recommendation on the Ethics of Artificial Intelligence, UNESCO, 2022. <https://unesdoc.unesco.org/ark:/48223/pf0000381137>. (Accessed 8 May 2024).

- [96] R. Owen, R. von Schomberg, P. Macnaghten, An unfinished journey? Reflections on a decade of responsible research and innovation, *J. Responsible Innov.* 8 (2021) 217–233, <https://doi.org/10.1080/23299460.2021.1948789>.
- [97] J.L. Hunte, M. Neil, N.E. Fenton, M. Osman, C. Bechlivanidis, The effect of risk communication on consumers' risk perception, risk tolerance and utility of smart and non-smart home appliances, *Saf. Sci.* 174 (2024) 106464, <https://doi.org/10.1016/J.SSCI.2024.106464>.
- [98] N. Gao, W. Shao, M.S. Rahaman, J. Zhai, K. David, F.D. Salim, Transfer learning for thermal comfort prediction in multiple cities, *Build. Environ.* 195 (2021) 107725, <https://doi.org/10.1016/J.BUILDENV.2021.107725>.
- [99] K. Li, Y. Liu, L. Chen, W. Xue, Data efficient indoor thermal comfort prediction using instance based transfer learning method, *Energy Build.* 306 (2024) 113920, <https://doi.org/10.1016/J.ENBUILD.2024.113920>.
- [100] M. Quintana, S. Schiavon, F. Tartarini, J. Kim, C. Miller, Cohort comfort models — using occupant's similarity to predict personal thermal preference with less data, *Build. Environ.* 227 (2023) 109685, <https://doi.org/10.1016/J.BUILDENV.2022.109685>.
- [101] A. Sharma, E. Kosasih, J. Zhang, A. Brintrup, A. Calinescu, Digital Twins: state of the art theory and practice, challenges, and open research questions, *J. Ind. Inf. Integr.* 30 (2022) 100383, <https://doi.org/10.1016/J.JII.2022.100383>.
- [102] V. Vakkuri, K.K. Kemell, J. Kultanen, P. Abrahamsson, The current state of Industrial practice in artificial intelligence ethics, *IEEE Softw* 37 (2020) 50–57, <https://doi.org/10.1109/MS.2020.2985621>.
- [103] H. Salo-Pontinen, P. Saariluoma, Reflections on the human role in AI policy formulations: how do national AI strategies view people? *Discov. Artif. Intell.* 2 (2022) 1–24, <https://doi.org/10.1007/S44163-022-00019-3/TABLES/1>.
- [104] K. Ahmad, M. Maabreh, M. Ghaly, K. Khan, J. Qadir, A. Al-Fuqaha, Developing future human-centered smart cities: critical analysis of smart city security, Data management, and Ethical challenges, *Comput. Sci. Rev.* 43 (2022) 100452, <https://doi.org/10.1016/J.COSREV.2021.100452>.
- [105] C. Debrah, A.P.C. Chan, A. Darko, Artificial intelligence in green building, *Autom. Construct.* 137 (2022) 104192, <https://doi.org/10.1016/J.AUTCON.2022.104192>.
- [106] Y. Li, T. Huwan, Applications of uncertainty models as support in smart buildings and ethical computing in Edge computing of smart cities, *Complexity* 2022 (2022) 2961999, <https://doi.org/10.1155/2022/2961999>.
- [107] T. Sharpe, Ethical issues in domestic building performance evaluation studies, *Build. Res. Inf.* 47 (2019) 318–329, <https://doi.org/10.1080/09613218.2018.1471868>.
- [108] Y. Gao, S. Shi, S. Miyata, Y. Akashi, Successful application of predictive information in deep reinforcement learning control: a case study based on an office building HVAC system, *Energy* 291 (2024) 130344, <https://doi.org/10.1016/J.ENERGY.2024.130344>.
- [109] Z. Shi, R. Zheng, J. Zhao, R. Shen, L. Gu, Y. Liu, J. Wu, G. Wang, Towards various occupants with different thermal comfort requirements: a deep reinforcement learning approach combined with a dynamic PMV model for HVAC control in buildings, *Energy Convers. Manag.* 320 (2024) 118995, <https://doi.org/10.1016/J.ENCONMAN.2024.118995>.

- [110] W. Xue, N. Jia, M. Zhao, Multi-agent deep reinforcement learning based HVAC control for multi-zone buildings considering zone-energy-allocation optimization, *Energy Build.* 329 (2025), <https://doi.org/10.1016/j.enbuild.2024.115241>.
- [111] F. Hou, J.C.P. Cheng, H.H.L. Kwok, J. Ma, Multi-source transfer learning method for enhancing the deployment of deep reinforcement learning in multi-zone building HVAC control, *Energy Build.* 322 (2024) 114696, <https://doi.org/10.1016/j.enbuild.2024.114696>.
- [112] A. Silvestri, D. Coraci, S. Brandi, A. Capozzoli, E. Borkowski, J. Kohler, D. Wu, M.N. Zeilinger, A. Schlueter, Real building implementation of a deep reinforcement learning controller to enhance energy efficiency and indoor temperature control, *Appl. Energy* 368 (2024) 123447, <https://doi.org/10.1016/J.APENERGY.2024.123447>.
- [113] D. Coraci, A. Silvestri, G. Razzano, D. Fop, S. Brandi, E. Borkowski, T. Hong, A. Schlueter, A. Capozzoli, A scalable approach for real-world implementation of deep reinforcement learning controllers in buildings based on online transfer learning: the HiLo case study, *Energy Build.* 329 (2025) 115254, <https://doi.org/10.1016/j.enbuild.2024.115254>.