

SURVEY PAPER

# Machine Learning for Smart and Energy-Efficient Buildings

Hari Prasanna Das<sup>1</sup> , Yu-Wen Lin<sup>1</sup>, Utkarsha Agwan<sup>1</sup>, Lucas Spangher<sup>1</sup>, Alex Devonport<sup>1</sup>, Yu Yang<sup>2</sup>, Ján Drgoňa<sup>3</sup>, Adrian Chong<sup>4</sup>, Stefano Schiavon<sup>5</sup> and Costas J. Spanos<sup>1</sup>

<sup>1</sup>Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, USA

<sup>2</sup>School of Automation Science and Engineering, Xi'an Jiaotong University, Xi'an, China

<sup>3</sup>Pacific Northwest National Laboratory, Richland, WA, USA

<sup>4</sup>Department of Building, National University of Singapore, Singapore, Singapore

<sup>5</sup>Center for Built Environment, University of California, Berkeley, CA, USA

**Corresponding author:** Hari Prasanna Das; Email: [hpdas@berkeley.edu](mailto:hpdas@berkeley.edu)

Authors Y.L., U.A., L.S., A.D., and Y.Y. contributed equally to this work.

**Received:** 01 March 2023; **Revised:** 11 October 2023; **Accepted:** 15 November 2023

**Keywords:** energy efficiency; machine learning; smart buildings

## Abstract

Energy consumption in buildings, both residential and commercial, accounts for approximately 40% of all energy usage in the United States, and similar numbers are being reported from countries around the world. This significant amount of energy is used to maintain a comfortable, secure, and productive environment for the occupants. So, it is crucial that energy consumption in buildings must be optimized, all the while maintaining satisfactory levels of occupant comfort, health, and safety. Machine learning (ML) has been proven to be an invaluable tool in deriving important insights from data and optimizing various systems. In this work, we review some of the most promising ways in which ML has been leveraged to make buildings smart and energy-efficient. For the convenience of readers, we provide a brief introduction to the relevant ML paradigms and the components and functioning of each smart building system we cover. Finally, we discuss the challenges faced while implementing machine learning algorithms in smart buildings and provide future avenues for research in this field.

## Impact Statement

In this work, we review the ways in which Machine Learning (ML) has been leveraged to make buildings smart and energy-efficient. In doing so, we touch upon traditional methods already in place and how ML-based solutions often match or outperform them. Our work aims to benefit ML researchers interested in how their expertise can be applied to the smart building sector, and smart building specialists—researchers whose primary interests are in energy efficiency, the quality of the indoor environment, the related occupant experience, advanced control methods, etc.—who wish to solve domain-specific problems using ML. Our work will equip researchers with a bird's-eye view of the ML applications in Smart Buildings, enabling them to visualize the existing interdependencies among different subsystems, and design ML-based implementations for optimizing energy usage and occupant comfort.

## 1. Introduction

We spend 90% of our time every day in indoor environments, so buildings substantially influence our health, well-being, safety, and work and study performance. Also, the energy used in buildings to ensure

individual safety and comfort is a leading contributor to climate change. It is estimated that in the United States, buildings (both residential and commercial) are responsible for approximately 40% of primary energy consumption, 73% of electrical use, and 40% of greenhouse gas emissions (U.S. Department of Energy—Energy Information Administration, Annual Energy Outlook, 2020). It is of utmost importance to improve building energy systems to optimize energy usage and thus limit the greenhouse gas emissions contributed by them, while, at the same time, ensuring an occupant-friendly environment to improve well-being and productivity. Energy use reductions while maintaining occupant comfort and productivity in buildings can be an environmentally sustainable, equitable, cost-effective, and scalable approach to reducing greenhouse gas emissions.

These requirements necessitate bringing to bear some of the latest sensing, computing, and control technologies into the built environment. It is expected that these technologies, appropriately integrated, will make the built environment responsive to occupant needs and environmental factors, delivering both superior occupant experience and unprecedented energy efficiency. In other words, the buildings will be “Smart”, which in this context refers to the union of “having buildings equipped with Internet-of-Things (IoT) devices and sensors to collect real-time data such as temperature, humidity, lighting levels, and air quality” as well as “having the infrastructure to harness the power of tools such as AI and machine learning (ML) to utilize data from buildings for energy efficiency and occupant comfort”.

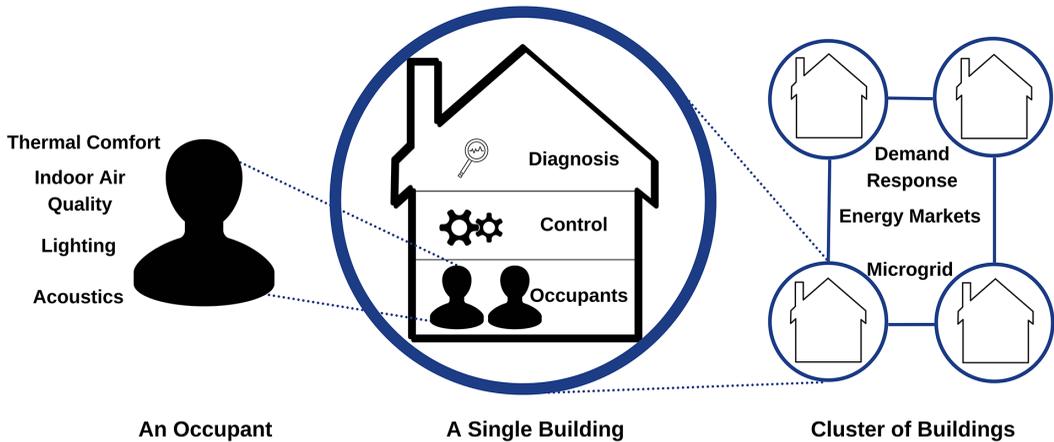
Additionally, as computing has evolved, ML techniques are being shown to be complementary and, in many cases, superior to more classical approaches. The latter require accurate models describing the various aspects of the building and occupant behaviors, and such models are often inaccurate and very hard to build and maintain. As we will show in this paper, ML techniques often fill this gap.

The smart building ecosystem is illustrated in [Figure 1](#). Occupants constitute the basic building block of the ecosystem. Being the consumers of the facilities that a building provides, occupants necessitate the regulation of the building systems to achieve the desired environment. The building comprises structures, devices, and systems in place to control and maintain the desired environment for the occupants, along with diagnostic systems to ensure robust operation. The building operation requires energy, which primarily comes in the form of electricity. The electrical energy is supplied to buildings via a power distribution system, where, with the advent of smart grids, buildings interact and exchange surplus energy and other ancillary services with the energy provider and with each other.

In order to improve energy efficiency and occupant comfort, researchers and industry leaders have attempted to implement intelligent sensing, control, and automation approaches alongside techniques like incentive design and price adjustment to more effectively regulate energy usage. With the growth of IoT devices, and the great variety of user-to-device and device-to-device interactions, there is a need for integration and coordination of the related objectives and actions. Further, in order to derive insights from the vast amount of data in certain scenarios, and from the limited amount of data in others, ML applications are proliferating in smart buildings. These ML-driven insights can be used for downstream tasks such as forecasting, prediction, and control. In this work, we explore the growing application of ML in smart building systems, and examine how ML-based solutions often match or outperform more traditional methods.

### ***1.1. Potential audience for this paper***

Prior studies related to this paper (Rupp et al., 2015; Wang and Hong, 2020; Zhang et al., 2021; Zhao et al., 2021) mainly focus on review studies in a subdomain. Another line of relevant prior studies include study of ML applications in smart buildings (Hong et al., 2020), from the viewpoint of an expert/researcher well aware of the dynamics of smart buildings. In this paper, we summarize and synthesize the research presented in these publications and provide a breadth-first view in the context of smart and energy-efficient buildings. Those who will benefit most from this summary are ML researchers interested in how their expertise can be applied to the smart building sector, and smart building specialists—researchers whose primary interests are in energy efficiency, the quality of the indoor environment and the related occupant experience, advanced control methods, occupancy modeling and monitoring, modeling and simulation, privacy, security, etc.—who wish to solve domain-specific problems using ML.



**Figure 1.** A taxonomy of the smart buildings illustrated at three levels: cluster of buildings, single building, and occupant.

For literature search, we used the Google Scholar platform with keywords specific to various sections presented in this work, as it efficiently indexes and provides research studies at the intersection of multiple domains, which in our case is ML and smart buildings. Unlike many other search engines, we found the above platform be efficient in scoping literature from journals and conference proceedings, since the destination for many smart buildings research studies is journals, while in the case of ML, it is academic conferences.

## 2. Machine Learning

In this section, we briefly cover the fundamentals of some ML algorithms and methods that are commonly in practice in smart buildings. These ML algorithms can be commonly grouped under *supervised*, *unsupervised*, and *reinforcement learning* depending on the learning paradigm, i.e., the way a model gets trained on the available data. We briefly cover the functionalities of the above learning paradigms and various learning methods, while placing an emphasis on methods that we have found particularly useful in the context of optimizing the performance of the built environment, as detailed in the subsequent section.

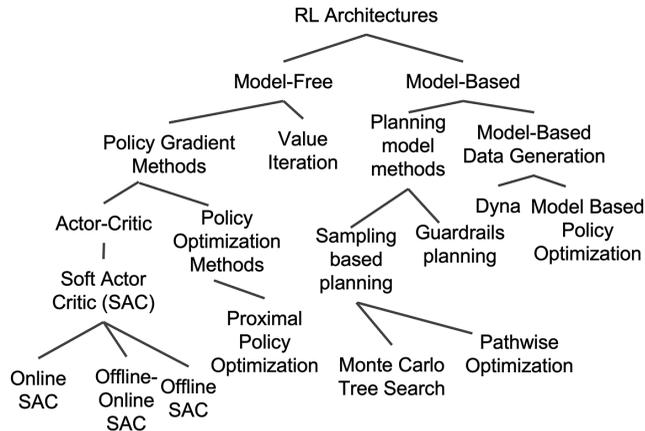
### 2.1. Learning paradigms

#### 2.1.1. Supervised ML

Supervised learning refers to a class of algorithms that learn using data points with known outcomes/labels. The model is trained using an appropriate learning algorithm (such as linear regression, random forests, or neural networks) that typically works through some optimization routine to minimize a loss or error function. The model is trained by feeding it input data as well as correct annotations, whenever available. If the output of the model is continuous, this process is called *regression* and if the output is discrete with finite classes, the process is called *classification*. Supervised learning has been implemented for substantial applications in real life, such as image classification/segmentation, natural language processing, time-series analysis. Since supervised learning requires annotated data, it faces a bottleneck when labeling is expensive or laborious.

#### 2.1.2. Unsupervised ML

In contrast to supervised learning, *unsupervised* learning involves models/algorithms that look for patterns in an unlabeled dataset with minimal human supervision. Some unsupervised learning paradigms include *clustering*, which aims to group data into similar categories (clusters); *dimensionality reduction*, which aims to find a low-dimensional subspace that captures most of the variation in the data; and *generative modeling*, which aims to learn the data distribution in order to generate synthetic data.



**Figure 2.** A taxonomy of reinforcement learning architectures.

### 2.1.3. Reinforcement learning (RL)

RL is a type of agent-based ML where a complex system is controlled through actions that optimize the system in some manner (Sutton and Barto, 2018). The actions ( $a_t$ ) taken based on a probability distribution  $p_\pi$  at a state ( $s_t$ ) and time  $t$  seek to optimize the expected sum of rewards ( $r$ ) based on a policy ( $J$ ) parameterized by  $\theta$ ; i.e.,  $J(\theta) = \sum_{s_t, a_t \sim p_\pi} [r(s_t, a_t)]$ .

RL is based on the Markov property of systems, where the evolution of system states is determined by the current state and the subsequent control sequence, and independent of the past states and past control inputs. Indeed, RL agents might be thought of as Markov decision processes (MDP) with neural nets serving as function approximators helping to decide which state transitions to attempt. RL is useful in when actions and environments are simple or data are plentiful. Early examples include backgammon (Tesauro, 1994), the cart-pole problem, and Atari (Mnih et al., 2013). Recently, much work has been done to extend RL to broader uses.

RL architectures may be grouped in families with different learning methods; a simple one is *policy gradient*. Here, the controller is made up of a single function approximator that maps actions to states. The model is trained along the gradient of expected reward. One subfamily of policy gradient is the *Actor-Critic* that employs two function approximators; the actor estimates actions as in policy gradient, while the critic estimates the long-term value of the actions; these estimates contribute to the training of the actor network. An overview of the RL architecture taxonomy is provided in Figure 2.

We now discuss specific classes of algorithms that are commonly used in the smart building domain.

## 2.2. ML learning methods

### 2.2.1. Kernel-based methods

Kernel algorithms are designed for analyzing patterns in data from explicit feature vector representations provided to them. They work by projecting the onto hyperplanes that make the inference easier. Kernel methods can be used for supervised and unsupervised problems. A commonly used kernel-based method for classification is the *support vector machine* (SVM). Another common method in unsupervised settings is kernel *principal component analysis* (PCA), used for dimensionality reduction in order to facilitate clustering.

### 2.2.2. Ensemble methods

*Ensemble* methods combine multiple ML models to produce an optimal model. Common examples include tree-based algorithms such as decision tree, *bagging* (bootstrap aggregating), *random forest*, and *gradient boosting*.

### 2.2.3. Deep learning

Deep learning-based methods are usually deployed on extremely large datasets, and they use very complex artificial neural networks inspired by information processing in biological neurons for applications in supervised, semisupervised, or unsupervised settings. The neural network architecture depends on the nature of the available data. For example, *convolutional* neural networks typically analyze image data, while *recurrent* networks are usually applied to textual, natural language or time-series data.

### 2.2.4. Time-series forecasting methods

Time-series data comprise a sequence of data points indexed by time, such as stock prices, outdoor temperature, or electricity prices. Time-series analysis is used to extract statistical information about the data and potentially make predictions about future data points. Wide variety of methods can be used for time-series analysis and forecasting. Assuming future data have some relation with historically observed data, linear models such as *moving average* (MA), *auto-regression* (AR), and combinations of these can be used to forecast future data points. Some use cases are in stock price forecasting (Pai and Lin, 2005) and in supply chain management (Aviv, 2003). *Recurrent* neural networks, and *long short-term memory* (LSTM) networks can also be used for time-series with more complex behaviors.

### 2.2.5. Physics-based ML

ML models have shown promising results in learning scientific problems that are yet to be well understood or where it is computationally infeasible with physics-based models. However, pure data-driven models often require very large datasets, offer limited extrapolation abilities, and are unable to provide physically consistent results or to offer intuition about the process they are modeling. On the other hand, pure physics-based models are often complex, include many uncertain parameters, and may be either too simplistic or too computationally intensive for the given application. This has led to the emergence of hybrid models that combine formal domain knowledge with ML. These hybrid models may compensate the weaknesses of pure physics-based models or pure data-driven models, and may be better able to both extrapolate and provide intuition about the process in question. Coupling physics knowledge with ML has been used in several areas, including medicine (Costabal et al., 2020; Kissas et al., 2020), fluid mechanics (Cai et al., 2022), power systems (Misyris et al., 2020), manufacturing (Qi et al., 2019) etc. It has shown potential in better prediction accuracy using a smaller set of sample data and generalizability for out-of-sample scenarios.

### 2.2.6. Learning, dynamics, and control

Traditional control approaches such as *system identification* (Ljung, 1998) and *adaptive control* (Ioannou and Fidan, 2006) have similarities to supervised learning (Ljung et al., 2020) and RL (Matni et al., 2019), respectively. While some system identification is implicit in all control approaches, as a part of how the model is formed, adaptive control never rose to the prominence of RL.

Several approaches to model-based RL use control-theoretic methods. Essentially, the problem is split into two parts: a dynamic model is learned from data, and the model is used to design a controller, with some variations, with *model predictive control* being a popular choice. Similarly, in traditional control applications, learned models are incorporated into modern control synthesis techniques, and standard methods of proving safety are adjusted to incorporate learned models. A learning-based approach that has proved to mesh particularly well with this approach uses *Gaussian process* models (Akametalu et al., 2014; Umlauf et al., 2017; Wang et al., 2018; Devonport et al., 2020, 2021b). Another point of contact is the use of statistical guarantees for controllers based on learned models. In RL, statistical guarantees of correctness such as *regret bounds* and *probably approximately correct* (PAC) bounds are used to ensure that optimal behavior is attained with high probability. By combining these statistical guarantees with control-theoretic techniques, they can also be used to ensure safety with high probability (Devonport and Arcak, 2020; Devonport et al., 2021a, 2021b). Since many learned models are statistical in nature, probabilistic bounds consistent with model accuracy have also been used extensively in recent years (Matni et al., 2019).

	Inter-Building 	Building 	Occupant 
<b>Control strategies</b> • Reinforcement learning • Hard control and model-predictive control	• Power flow coordination • Smart energy storage and distribution	• HVAC, lighting control • Space scheduling • Demand response	• Occupant-centric environmental control
<b>Learning Models</b> • Supervised learning • Physics-based learning • Generative modeling	• Demand response prediction • Building load prediction	• Building control • Fault diagnosis • Occupancy detection	• Thermal comfort • Activity detection • Utility Function Learning
<b>Markets</b> • Automated pricing • Time-series forecasting • Gamification	• Aggregation & mechanism design • Energy allocation	• Demand response cost optimization • Smart load scheduling	• Mechanism design based on occupant utility

**Figure 3.** Illustration of various applications where machine learning methods can be deployed in smart buildings, grouped at the cluster of buildings level, the building level, and the occupant level.

### 3. Applications of ML in Smart Buildings

We organize the ML methods we have used in smart building applications by their targeted level of abstraction, namely at the level of *occupants*, *single building*, and a *cluster of buildings* as shown in Figure 3. Above applications directly or indirectly lead to enhancements in energy efficiency in buildings, while ensuring occupant comfort and productivity.

#### 3.1. Thermal comfort

The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) defines thermal comfort as “the condition of the mind that expresses satisfaction with the thermal environment and is assessed by subjective evaluation” (Standard 55, 2020). We spend more than 90% of time in a day being within a built environment, where our health, well-being, performance, and energy consumption are linked to thermal comfort. However, studies show that only 40% of commercial building occupants are satisfied with their thermal environment (Graham et al., 2021). Ensuring occupants’ thermal comfort involves understanding of the parameters that affect it, developing and using models to predict thermal comfort, and controlling heating, ventilation, and air-conditioning (HVAC) systems to achieve occupant satisfaction. ML has been applied to all of these sectors. In this section, we mainly cover thermal comfort understanding and prediction. In Section 3.4, we cover the control of the thermal environment in buildings. Thermal comfort of building occupants has been studied extensively. These models include traditional ones like the *Predicted Mean Vote* (PMV) and the *adaptive thermal comfort* models, as well as the more recent *personal comfort* models.

The most widely used thermal comfort model is the PMV, developed by Fanger (1970) based on a set of experiments in controlled climate chambers. Fanger established a mathematical formula with six input parameters: air temperature, mean radiant temperature, relative humidity, air velocity, clothing insulation, and metabolic rate. The PMV model adopts a 7-point scale for its thermal sensation output, with values ranging from -3 to +3, and indicating cold, cool, slightly cool, neutral, slightly warm, warm, and hot, respectively. To gauge the level of dissatisfaction of individuals in a space where PMV has been computed, Fanger proposed a *Predicted Percentage of Dissatisfied* (PPD) model (Fanger, 1970), that establishes a quantitative prediction of the percentage of thermally dissatisfied occupants. To comply with ASHRAE 55 standards, the recommended range for PMV on the 7-point scale is between -0.5 and 0.5. The PMV model is designed with a uniform and steady-state conditioned environment and does not

explicitly consider local discomfort, the nonuniformity of the space and dynamic thermal conditions (Zhao et al., 2021). It has also been developed with data from controlled climate chamber experiments, which are not necessarily transferable to the varying thermal environments that exist in buildings around the world (Cheung et al., 2019). Experiments in real buildings showed that the PMV has a low prediction accuracy (around 33%), whereas the PPD is unreliable as a metric (Cheung et al., 2019).

With the booming development of ML techniques and their versatile applications, researchers have attempted to develop data-driven approaches using ML to predict thermal comfort. Often, the Fanger's features, alone or with additional relevant real and synthetic features, are fed into a data-driven model to learn the connections between the features and the thermal preference labels (Kim et al., 2018; Liu et al., 2019). The model is later leveraged to predict the same given raw features. Because the models are trained to learn directly from the data, and not from a rule that was established using prior experiments, they perform better as compared to PMV.

Kernel-based approaches have been widely used for thermal sensation/preference prediction. The list of kernel-based methods popular for thermal comfort prediction includes *SVM* (Chaudhuri et al., 2018; Alsaleem et al., 2020; Chai et al., 2020; Liu et al., 2020; Zhou et al., 2020), *K-nearest neighbors* (KNN) (Liu et al., 2019; Lu et al., 2019; Pigliatile et al., 2020; Lee and Ham, 2021; Cheung et al., 2022), and *ensemble learning* algorithms, such as *random forest* (RF) (Kim et al., 2018; Liu et al., 2019) and *AdaBoost* (Ab). Recently, *feed-forward neural networks* (Zhai and Soh, 2017; Lu et al., 2019; Das et al., 2021), and *time-series based networks* (Chennapragada et al., 2022) have surpassed state-of-the-art kernel-based models in thermal comfort prediction.

Traditional thermal comfort models, as described above, are developed based on aggregated data from a large population. So, rather than predicting the thermal comfort of individuals, they were designed to predict the average thermal comfort of a population, when all its members are exposed to the same environment. This naturally misses the inevitable and sometimes significant differences in how different individuals respond to the same thermal environment. A new approach to thermal comfort modeling that uses personal comfort models instead of the average response of a large population, can be applied to any building thermal control system (Kim et al., 2018). Personal thermal comfort can adapt to the available input variables, such as environmental variables (Cheung et al., 2017), occupant behaviors (Kim et al., 2018), and physiological signals (Liu et al., 2019). Several experiments have been conducted over the years for population groups varying in terms of environmental and personal conditions. ML algorithms ranging from kernel-based methods to neural network-based methods have been proposed as the predictor. For instance, Liu et al. (2019) conducted an experiment to collect physiological signals (skin temperature, heart rate) of 14 subjects (6 female and 8 male adults) and environmental parameters (air temperature, relative humidity) for 2–4 weeks (at least 20 h per day). They developed ML-based predictors of thermal preference for each individual, using synthetic features derived from the above raw features (e.g., mean heart rate over last 5-min from the survey response time) as the input with RF being the best predictor. A data-driven method with indoor environment was applied to classify occupant's personal thermal preference with temperature and humidity sensors (Laftchiev and Nikovski, 2016). Authors in Sim et al. (2018) developed personal thermal sensation models based on watch-type sweat rate sensors. Among all the ML methods, kernel-based algorithms are still the most commonly used ones. However, other ML approaches are also becoming popular in modeling the complex interactions that exist between the features without much feature engineering, e.g., time-series prediction (Somu et al., 2021), artificial neural networks (Das et al., 2021), etc. Better approaches for modeling tabular data in smart buildings, with a focus on thermal comfort datasets are provided in Das and Spanos (2022a).

Although several studies show high prediction accuracy of thermal comfort with various ML methods, there are still several challenges ahead. A common challenge in designing ML-based thermal comfort predictors is the issue of class imbalance in data. Almost all of the thermal comfort datasets (Ličina et al., 2018; Liu et al., 2019) are inherently class-imbalanced, i.e., they have more data belonging to “Prefer No-Change” than “Prefer Warmer” and “Prefer Cooler” thermal preference classes. Researchers have tackled this issue by using weighed loss functions for ML models. A related challenge is the overall lack of sufficient amount of data (Liu et al., 2019). Collecting large amounts of data as required by ML models from humans via real-world experiments is expensive and cumbersome. Generating synthetic data to

augment the training dataset is one of the approaches that has been proposed for tackling the above challenges. Synthetic data generation can be done using classical methods such as SMOTE (Chawla et al., 2002; Quintana et al., 2020) or using advanced neural network-based generative models (Quintana et al., 2020; Das et al., 2021; Yoshikawa et al., 2021; Das and Spanos, 2022b).

Another challenge is domain discrepancy. Thermal comfort, as per PMV, is dependent on six major parameters as described in the beginning of the section or, as per the adaptive thermal comfort model, is dependent on the outside temperature. However, it also varies from person to person, across climatic regions and economic conditions. A literature review of personal comfort models concluded that there is a lack of diversity in terms of building types, climates zone, and participants that are considered in existing thermal comfort studies (Martins et al., 2022). Under such domain discrepancy, models developed in one environment, when used in another target environment may lead to low accuracy or misleading predictions. Also, thermal comfort modeling in general depends largely on self-reporting, which is inherently unreliable. To deal with the data/label insufficiency challenge, domain adaptation methods have been proposed to adapt ML models from one domain to another. Das et al. (2021) propose a transfer learning framework, using Adversarial Domain Adaptation (ADA) to develop personal thermal comfort predictors for target occupants in an unsupervised manner. Gao et al. (2021) propose a transfer learning-based *multilayer perceptron* model for domain adaptation across cities.

In summary, ML models have proven to improve thermal comfort prediction accuracy and they allow a wider and more flexible range of input parameters. They can adapt to the available data streams, and they improve their performance over time. While there are still challenges, there are also many future research opportunities, as listed in Section 5.

### 3.2. Occupancy and activity sensing

Occupancy and activity sensing are key aspects for the observability of a human-in-the-loop building control system. Traditionally, building operation methods that include occupancy as one of their parameters, such as starting heating/cooling from early morning until late in the evening during weekdays assuming maximum occupancy during working hours often have static schedules set for the occupancy, which is far from realistic. Also, how much a building will be occupied depends on several other factors, such as weather, building type, and holiday schedule. Such static policies may lead to a significant waste in energy consumption, because the heating/cooling and ventilation levels are set with no regard for the actual occupancy level. Activity sensing also helps to provide personalized, context-aware services in buildings, thus enhancing overall satisfaction in buildings while creating a safety net for adverse events such as falls in elderly homes (Petroşanu et al., 2019).

Occupancy sensing can be performed using both intrusive and nonintrusive methods. Intrusive methods require the occupants to carry an electronic device whose signature is followed by a central server to infer occupancy/positioning (Lee et al., 2013; Filippoupolitis et al., 2016; Zou et al., 2016, 2018a, 2018b). However, requiring occupants to constantly carry a device is not reliable. This problem gets magnified for the case of elderly population. Hence, nonintrusive methods for occupancy sensing are getting popular. In nonintrusive method, occupancy/activity is detected from data collected from various sensing modalities, and some ground truth during controlled experiments. Machine learning-based methods are very effective for nonintrusive occupancy sensing, allowing for the data-driven construction of a model to map sensor data into an estimate of occupancy.

Nonintrusive methods for occupancy sensing can be divided according to the sensing modalities they employ. A common choice is video feeds from cameras installed in rooms. Regular cameras have not found significant usage because of their privacy issue; infrared sensors or thermal cameras have rather been the choice for occupancy detection. For example, Kraft et al. (2021) use a U-Net-like convolutional neural network on thermal images to infer occupancy. Other studies employing similar ML methods on depth cameras are Brackney et al., 2012; Diraco et al., 2015; and Zhao et al., 2019. But, in general, cameras have other issues such as poor illumination conditions and occlusion. A recent body of work focuses on occupancy and activity detection from Wi-Fi signals (Zou et al., 2019), because of its ubiquitous presence, and better privacy guarantees. Zou et al. (2017, 2018c) use Channel State

Information (CSI) data collected from Wi-Fi sensors (a transmitter and a receiver) and measuring the shape similarity between adjacent time-series CSI curves to infer the occupancy. They improve the detection mechanism in Zou et al. (2019) by using convolutional neural networks on the CSI heatmaps to detect human gestures. Another modality that is used to detect occupancy is CO<sub>2</sub> data in a room. A number of studies (Rahman and Han, 2017; Zuraimi et al., 2017; Arendt et al., 2018; Wei et al., 2022) employ ML methods to map the CO<sub>2</sub> concentration and occupancy. Finally, studies such as Zou et al. (2019) propose sensor fusion, where data from multiple sensing modalities, i.e., RGB camera, and Wi-Fi, are used in tandem to come up with a robust activity detection mechanism.

Many of the sensors used for occupancy and activity detection are spatial in nature, i.e., their data are specific to the design of the spatial environment that they are deployed in. Hence, ML algorithms designed for one environment fails to perform optimally for new environments. Hence, transfer learning methods have been proposed in the literature to adapt ML models across domains and sensor choices. Zou et al. (2019) adapt an ML model to detect gestures using Wi-Fi CSI data between two rooms in a building. Such applications involving transfer learning between domains has been explored in the application of occupancy and activity sensing in several other literature studies (e.g., Pan et al., 2008; Khalil et al., 2021; Li et al., 2021; Dridi et al., 2022; Omeragic et al., 2023). Pinto et al. (2022) delve further into a review of transfer learning application in smart buildings.

To summarize, ML methods have been proposed to map the relation between data from sensors, and occupancy/activity. Once trained, these models can then be deployed to predict the occupancy which then gets fed to building control mechanisms (Esrafilian-Najafabadi and Haghghat, 2021).

### 3.3. Building design and modeling

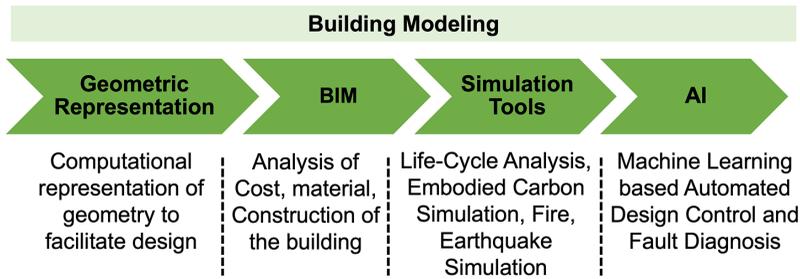
Building information modeling (BIM) was introduced in the 1970s when Eastman (1975) proposed the idea of migrating from hand-drawn building design toward automation and digitization. Early building modeling focused on the computational representation of geometry, which is also known as 3D modeling. Later, a set of new tools focusing on the building performances were developed, including Building Design Advisor (BDA) (Papamichael et al., 1997), Ecotect (Roberts and Marsh, 2001), EnergyPlus (Crawley et al., 2001), ESP-r (Strachan et al., 2008), DEST (Yan et al., 2008), and Modelica (Fritzson and Engelson, 1998). They focused on thermal, energy, lighting, and air quality aspects. They provided feedback and “what-if” analyses. With the emergence of IoT, data have become more easily accessible, which has helped support building simulation tools and allowed development of ML models. The development of building models over the past decades is summarized in Figure 4.

The goal of building modeling is to support the design, construction, and operation of buildings. Early building modeling was mostly used during the design phase, and was not usually maintained or updated after the building was constructed. In recent years, the digital twins (DT) concept (Grieves and Vickers, 2017) has emerged for physical systems where a virtual “twin” is developed that mirror the actual physical system. DT can monitor a building’s condition in real-time and plays an important role in control applications, fault diagnosis, and prognosis. The following sections describe the intersection between ML and building modeling. It is separated into 5 parts: models, design automation, DT, applications, and challenges.

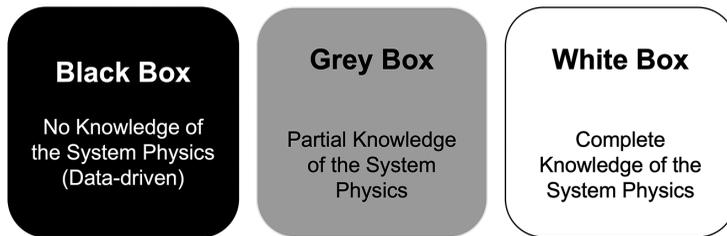
#### 3.3.1 Models

There are three types of building models: white box, black box, and gray box, as illustrated in Figure 5.

**3.3.1.1. White box model.** The “white box” or “first principles” approach models the building system with detailed physics-based equations. Building components and subsystems are modeled in a detailed manner to predict their behaviors. Common Building Performance Simulation (BPS) tools for physics-based modeling are Modelica (Fritzson and Engelson, 1998), EnergyPlus (Crawley et al., 2001), TRNSYS (U. o. W.-M., 1975), and ESP-r (Strachan et al., 2008). Though a white box model can capture the building dynamics well, it requires a lot of information from the physical building and may suffer from high uncertainties.



**Figure 4.** Evolution of building modeling.



**Figure 5.** Three types of building models.

**3.3.1.2. Black box model.** A black box model, also known as a data-driven or ML model, is a modeling approach that constructs a model directly from data without knowledge of the system physics. With the emergence of IoT, sensor measurements are continuously collected at various locations in a building to reflect current operations, which provides an opportunity to build data-driven models. For example, building load prediction using black box models is an important aspect of improving building energy performance. Its applications include control optimization, fault diagnosis, and demand side management (Zhang et al., 2021). The common ML methodologies for building load prediction are ANN, SVM, Gaussian-based regressions, and clustering (Seyedzadeh et al., 2018). The drawbacks of black box models include the need for a large amount of data and the fact that the accuracy of the model is highly dependent on the data quality. Another drawback is their fragility, which occurs when they give results that are not physically possible and may not perform reliably in conditions that have not been seen before. In addition, black box models lack interpretability and may have a high computational cost.

**3.3.1.3. Gray box model.** Gray box modeling is a modeling approach that simplifies the physical equations to simulate the behavior of a building and it combines them with ML techniques. By doing that, it aims to reduce the system complexity while maintaining prediction accuracy. An example of such a model is using thermal networks to model buildings (Bacher and Madsen, 2011), which is a common approach in building modeling. Thermal networks model dynamics of a building through resistor and capacitor (RC) circuits, and the parameters are identified with measurements from the real system. Studies show that the gray box model works well in building control applications (Široký et al., 2011; Oldewurtel et al., 2012a) and grid integration. However, the design of the thermal network is still ambiguous. The best selection of the number of resistors and capacitors for the model is unclear. Bacher and Madsen (2011) analyzed the model selections and found that the 3R3C model is the minimal best fit model for a 120 m<sup>2</sup> building located in Denmark. However, it is inconclusive if all buildings of similar size and use can be modeled with 3R3C with the same accuracy.

Lin et al. (2021) build a building model based on information available to minimize the number of uncertain parameters by combining physics-based equations with a neural network. Other researchers (Daw et al., 2017; Robinson et al., 2022) explored physics-based ML for modeling lake temperature and other dynamical systems. However, there are limited studies in physics-based ML for building modeling.

The definition of gray box models is still ambiguous, where it can refer to a resistor–capacitor network or a hybrid model that utilizes both first principal equations and data-driven methods. While gray box models have shown success in several applications, they are not systematically defined or thoroughly analyzed.

**3.3.2. Design automation.** Buildings are complex systems. Among their objectives is to provide comfortable and safe conditions to people. Building performance is affected by people’s behavior (Hong et al., 2017). Developing a model of a building with these complexities is often time-consuming and costly. Though white box models are widely used in simulation-based studies, gray and black box model are more commonly used in experimental studies of real buildings (Zhan and Chong, 2021). As a result, building design automation is an active area of research to help in developing fast and accurate building models.

Currently, automated design of buildings mostly involves the development of 3D building models. Most of the 3D models are built with Revit/AutoCAD based on site measurements or estimations during the design phase. Sometimes, researchers collect building measurements through LiDAR or airborne stereo imagery to construct 3D building models (Haala and Kada, 2010), which can automatically be imported to BIM or building simulation tools. However, this approach only captures the building at a geometrical level and cannot represent the operational level of buildings. Jia et al. (2018) proposed a platform-based design to reduce redundancy in hardware and software usage in building applications. The design performance is optimized by exploring the design space.

The emergence of building design libraries such as Modelica Building Library (Wetter et al., 2014) opens up opportunities for platform-based design automation of buildings. However, automated building modeling that captures building dynamics is underdeveloped, and most modeling methods still require human intervention.

**3.3.3. Digital twins.** The AIAA and AIA (American Institute of Aeronautics and Astronautics, 2022) recently published a position paper that defines a digital twin as “*A set of virtual information constructs that mimics the structure, context and behavior of an individual/unique physical asset, or a group of physical assets, is dynamically updated with data from its physical twin throughout its life cycle and informs decisions that realize value.*” The same definition can be applied to buildings, specifically how a physics-based model evolves over the course of its lifecycle as information or data about the building becomes available, i.e., as the building goes from the design phase (as designed) to the construction and commissioning phase (as built) to the postoccupancy phase (“as operated”). For buildings, forward BPS tools such as EnergyPlus (Crawley et al., 2001), TRNSYS (U. o. W.-M, 1975), ESP-r (Strachan et al., 2008), and Modelica (Fritzson and Engelson, 1998) are usually used to estimate the energy use of the building and its subsystems. Over the past decade, ML has been increasingly applied to BPS, fueled in part by the emergence of the IoT and the need for computationally tractable predictions.

ML has been used during the design stage to augment generative design and parametric simulations. Deep generative algorithms such as Generative Adversarial Networks (GANs) (Huang and Zheng, 2018; Nauata et al., 2020) have been proposed for generating diverse but realistic architectural floorplans that are known to be a time-consuming iterative process. The automated generation of architectural floorplans can be coupled with BPS tools to systematically explore architectural layouts that optimize building energy efficiency (Gan et al., 2019).

Metamodeling, defined as the practice of using a model to describe another model as an instance (Allemang et al., 2011), is another aspect where ML has been extensively applied to BPS throughout the building lifecycle. Given the complex interaction between different building systems and subsystems, design optimization during early design typically requires exploring high-dimensional design space. Consequently, ML has been used to create metamodels that can be used for optimization and uncertainty analysis (Eisenhower et al., 2012; Bre et al., 2020). Although most studies predict energy consumption and thermal comfort, metamodels have also been proposed for the emulation outputs such as natural ventilation and daylighting to support the modeling of passive design strategies (Chen et al., 2017). A comparison of different metamodeling techniques can be found in Østergård et al. (2018).

Moving from design to postoccupancy, model calibration is often undertaken to improve the model's credibility by reducing the discrepancies between simulation predictions and actual observations. Metamodeling using ML is often used to emulate simulation predictions to alleviate the high computation costs (Coakley et al., 2014). Particularly, metamodeling is almost always applied in Bayesian calibration that is computationally intractable because of the need to perform many model evaluations (Chong et al., 2021). More recently, Bayesian optimization and meta-learning through deep probabilistic neural networks have been proposed to reduce the number of simulations required for model calibration significantly, avoiding the need for metamodels altogether (Chakrabarty et al., 2021; Zhan et al., 2022).

BPS has also been applied to enhance the use of deep RL (DRL) based building controls. For instance, BPS can be used to develop building test cases for the comparison and benchmarking of DRL algorithms (Blum et al., 2021). The deployment of DRL in real-world building controls can be challenging because of the need for large amounts of training data before achieving acceptable performance (Botvinick et al., 2019). Consequently, the exploration during online training can lead to undesirable building operations when the algorithm performs suboptimally. To enhance the practical applications of DRL, BPS has been used as an emulator to pretrain a DRL agent offline. For instance, Zhang et al. (2019) used an automatically calibrated EnergyPlus model for the offline training of an A3C (Asynchronous Advantage Actor-Critic) algorithm that is subsequently deployed in a radiant heating system. It was also shown that a DRL trained offline could achieve similar near-optimal results as model predictive control (Brandi et al., 2022). However, pretraining a DRL agent using a model contradicts the advantages of DRL being "model-free" because BPS models can be time-consuming to develop. Therefore, further research could usefully explore the use of simpler and a lower level of detail models for the offline training of DRL agents for building controls.

**3.3.4. Applications: Design, control, and fault diagnosis.** The main applications for building modeling are facilitating design choice, testing control algorithms, and detecting faults in a system. BPS is used in the design phase of buildings to provide feedback in energy consumption, and greenhouse gas emissions and to guide design choices (Clarke, 2007; Hensen and Lamberts, 2012). Additionally, the development of a robust control system often requires a detailed understanding of system dynamics. Building modeling can provide dynamic feedback to test various control algorithms, which will be discussed in more detail in the next section. Fault detection also plays a crucial part in reducing maintenance costs and increasing the energy efficiency of building operations (Dong et al., 2014). However, faults in an actual building do not occur frequently, and it is hard to collect fault data for analysis. Data sets for fault diagnosis are usually created through testbed experiments or simulations. BPS can provide "what-if" analysis to assist in fault modeling and creating meaningful datasets.

**3.3.5. Challenges.** There are two main challenges in building modeling. First, there is a gap between building models during the design and the operation phases. Moreover, no standard connections are established between BIM and BPS tools, which results in unavoidable human effort to recreate the model for operation after the design phase. Though some established geometrical models can be exported to the BPS software, there are limitations in software compatibility and set up requirements. Second, there is no systematic approach to model selection. Limited studies have explored whether a certain type of model (white, black, or gray box model) is better for a specific application. Developing a level of abstraction of building modeling may benefit the development of a better model selection process for various applications. We anticipate that techniques from machine learning, such as neural architecture search and physics-based ML, can be used here to improve building modeling.

#### **3.4. Building control**

A critical application of a building model is the design of control systems for the building subsystems. These are typically "local models" that capture only aspects of the building that are pertinent to low-level control, such as a physical model heat and ventilation in a collection of rooms near a target temperature. However, a large-scale model is still important for controller design, as it is necessary to understand how

the subsystems of a building interact. For the purposes of this review, the most important subsystem in terms of complexity of control and size of energy demand is arguably the HVAC system, so we will focus our attention there. However, we do want to note that HVAC is by no means the *only* control problem that has received focus; other building control systems enjoy active work and engaged communities, including lighting (Lee and Selkowitz, 1994; Panjaitan and Hartoyo, 2011; ul Haq et al., 2014; Wagiman et al., 2019), window control (Psomas et al., 2017; Yoon et al., 2020), electric vehicle charging (Kontou et al., 2017; Turker and Bacha, 2018; Al-Ogaili et al., 2019), and domestic water heating (Passenberg et al., 2016; Wanjiru et al., 2017; Starke et al., 2020).

Controlling the HVAC system of a building can be, and for many decades has been, accomplished by classical control techniques. Here, the control objective simply being to keep temperatures somewhat close to a setpoint or within a deadband. Two simple controller types have proven to achieve this goal for HVAC systems: (1) threshold-based “bang-bang” controllers (2) and PID controllers. Bang-bang control applies the greatest allowable heating or cooling when the measured temperature leaves a certain band around the target temperature, but otherwise remains inactive. PID controllers use a linear feedback strategy, which decides a control action using a linear function of previous measurements. PID denotes a “proportional-integral-derivative” function formed as the sum of the current measurement (the proportional component) and approximations of the time derivative and integral of the temperature. These simple control policies even achieve optimality in certain cases: for example, the bang-bang strategies are optimal for minimizing the total time that the HVAC system is active, and PID controllers can be tuned to minimize the sensitivity of the controller to exogenous disturbances. However, if we ask for a strategy that is optimal with respect to delivering thermal comfort, or to account for energy load profiles throughout the day, predictively preheat and cool, or try to minimize other metrics such as carbon emissions, these simple control strategies no longer suffice.

The optimal operation of HVAC systems with respect to cost, emissions, comfort, or a combination therein is a multistep optimization problem due to the thermal capacity of buildings and the nonlinear operation of the HVAC system. Specifically, the dynamics of building thermal behavior is a slow-moving process and subject to time delays. This enables the precooling or preheating of buildings (e.g., storing the cooling energy in building mass). Besides, the operation efficiency of HVAC system is nonlinear with respect to supplied air. For example, the amount of energy required to maintain a constant temperature in a zone is a product of supplied zone air flow rates and the difference of the indoor temperature and the supplied temperature. The fan power of air handling unit (AHU) is cubic with respect to the total supplied air. In addition, the operation of HVAC system suffers from time-varying disturbances, such as the outdoor weather variations (i.e., radiation, temperature, humidity, etc.) and the occupancy changes.

HVAC system operation has raised extensive concerns in recent decades due to the continued increase of energy consumption of buildings and the high priority to save building energy consumption. Fortunately, the related work is plentiful and rapid progress has been achieved in this area. Existing studies can be broadly categorized by model architecture into: (1) *Model Predictive Control* (MPC) and (2) *RL*. On top of that, the combination of ML with MPC is developed with the objective to explore the benefits of the classic MPC controllers and the powerful approximation and characterization capabilities of ML tools.

MPC can be used to turn a standard ML algorithm for supervised learning into a model-based RL strategy for optimal control. The key idea is that MPC can be decomposed into two functional components: a state transition model and an optimizer. The ML half of the reaction is responsible for learning a state-space dynamical model— the goal is essentially to learn an application-specific digital twin. The predictive model may be application-agnostic (e.g., a general neural net) or leverage some degree of model knowledge as in the physics-informed model used in Nghiem et al. (2023).

In the following sections, we review the existing studies in ML and building control under three categories.

### 3.4.1. Model predictive control

MPC takes a large step toward improving these aims by incorporating optimization directly into the control policy. We wish to frame briefly our section on MPC: although it would not be considered a classic

case of ML, we find it helpful to thoroughly embed the motivations behind its development and active use into our understanding of the learning-based methods to follow. Thus, we will spend significant time describing it.

The key idea of MPC is to predict the value of the objective function in the future using an auxiliary dynamical model of the building's behavior and to select a control action that optimizes the predicted reward. Generally, if the dynamics model is  $f$ , the action  $u$ , the state  $x$ , the objective  $y$ , and the current timestep  $t$ , then MPC solves:

$$\begin{aligned} \underset{u}{\operatorname{argmin}} \quad & \sum_{\tau=t}^{t+T} -y(x, u) \\ \text{subject to: } & x(\tau+1) = f(x(\tau), u(\tau)); \tau = t, \dots, T-1 \\ & x(t+T) = 0; \tau = t, \dots, T-1 \\ & x(\tau) \in \mathcal{X}; \tau = t, \dots, T-1 \\ & u(\tau) \in \mathcal{U}; \tau = t, \dots, T-1 \end{aligned} \quad (3.1)$$

The condition  $x(\tau) \in \mathcal{X}$  is a state-space constraint, which can be used to ensure safety by attempting to keep the state vector in a region of the state space that is known *a priori* not to be dangerous.

For a given state  $x$ , (3.1) computes a sequence of  $T$  control actions in order to minimize the predicted reward. While the building control policy *could* use all  $T$  actions, so that (3.1) need only be solved every  $T$  time steps, to do so would force the controller to ignore what is happening in the outside world for those time steps. Since the control model used in (3.1) can never be fully accurate, and since dynamical prediction errors compound very quickly, what happens during those time steps is likely very different from what the controller predicts in (3.1). Moreover, the controller would be effectively blind to unanticipated dangers while executing the sequence: by the time  $T$  time steps are up, it may be too late to bring the system back to safety.

Thus, the second key idea of MPC is to resist the temptation to use the full input sequence: an MPC controller solves (3.1) at each time step, *and applies only the first computed input*. Using only the first control action is how MPC mitigates potential dangers while retaining optimal behavior.

MPC has been widely used to achieve anticipatory control of HVAC system for energy saving while ensuring human comfort (Drgoňa et al., 2020). MPC-based control methods require the solving of a multistep optimization problem so as to achieve the control sequence. The control output at the current stage is exerted and this process is repeated with the evolving of time. MPC-based control methods are online control methods that rely on the dynamic predictions of system disturbance. For HVAC control, several critical issues are related to the development of MPC-based controller: modeling complexity, modeling accuracy, and prediction accuracy. MPC-based controllers have been popular with HVAC control (see (Oldewurtel et al., 2012b; Afram and Janabi-Sharifi, 2014; Drgoňa et al., 2020; Kathirgamanathan et al., 2021) for a comprehensive review). The MPC controllers generally use either physics-based models (also known as “analytical first principle” or “forward models”) or data-driven models (also known as “black box” or “inverse models”) to predict system output.

MPC-based controllers have many advantages over rule-based or PID controllers. Since the energy saving and human comfort targets may be directly expressed in the MPC's objective and constraints and the time-varying disturbances are directly considered, MPC shows superior performance in energy efficiency and in ensuring human comfort. In addition, MPC-based controllers are robust to the time-varying disturbance, quick transient response to the environment changes, and consistent performance under a wide range of varying operating conditions.

However, MPC's advantages are achieved at the cost of frequently solving comprehensive optimization problems online. The intrinsic problem related to the HVAC control is nonlinear and nonconvex. This leads to the high online computation cost. Many efforts have been made to address the computational challenges. One solution is to restrict attention to *linear models* of the form  $x(t+1) = A(t)x + B(t)u$ , where

$A(t)$  and  $B(t)$  are matrix functions whose values may be time-dependent but not state-dependent. This model simplification, in conjunction with a quadratic objective, reduces (3.1) to a *quadratic program*, a convex problem that even embedded devices can solve quickly. However, a simpler model will necessarily lead to predictions of lower accuracy, and thereby worse performance. If we do not wish to simplify the model, an alternative is to develop explicit MPC controllers (Drgoña et al., 2013; Klaučo and Kvasnica, 2014; Parisio et al., 2014). The key idea is to drive some close-form solution for the MPC problem and then identify the parameterized control outputs online. To achieve the objective, simplified or approximated linear models for HVAC control are generally used. The benefit of explicit MPC is that it can be implemented on simple hardware and yields low online computation cost. Another line of work to address the computation burden is to develop distributed solvers. When MPC-based controllers are applied to large-scale commercial buildings, we require to solve large-scale optimization problem that couples the control of the HVAC system supplying multiple zones or rooms. Centralized methods tend to suffer from the high computation cost and are not scalable. Many studies have made effort to develop distributed MPC controllers (Radhakrishnan et al., 2016, 2017; Zhang et al., 2017; Yang et al., 2020, 2021). The major challenge to develop distributed MPC controllers for HVAC system is to handling the couplings across the multiple zones, including the thermal couplings (i.e., heat transfer), the composite objective (i.e., the fan power is quadratic with respect to the total zone mass flow rates), and operating constraints). The optimization problem related to multizone HVAC system is a nonlinear and nonconvex optimization problem subject to nonlinear couplings (nonlinear and nonconvex). The existing distributed solution methods cannot be directly applied to solve such a class of problems. To overcome the challenges and enable the decomposition of the problem, (Zhang et al., 2017; Yang et al., 2020, 2021) has explored the convex relaxation technique to solve such problem. Particularly, Yang et al. (2020, 2021) have applied the well-known alternating direction method of multipliers (ADMM) to solve the resulting relaxed problems, which have nonconvex objectives but subject to linear couplings. The superior performance both in energy saving and computation efficiency have been demonstrated via simulations. Slightly different, Radhakrishnan et al. (2016, 2017) have explored the decomposition of the problem with respect to the zone controller via hierarchical optimization and relaxation. Particularly, these works have relied on predictions to predict the thermal couplings across the zones. Besides, the objective components are optimized in two steps. These studies have comprehensively discussed the challenges to develop distributed control methods for large-scale commercial buildings due to the problem complexity.

#### 3.4.2. Reinforcement learning

Time-varying and uncertain parameters such as the weather variations and occupancy represent one major challenge for HVAC control. MPC controllers have relied on short-term dynamic predictions to address such problems but at the cost of high online computation burden and an accurate model to predict the system state evolution. These are the two main obstacles for the wide deployment of MPC controllers in practice. As one main branch of ML, RL explores to optimize the operation of system through the interaction with environment as long as the performance can be observed and quantified.

RL being one of the mainstream tools for multistage decision making under uncertainties (Sutton and Barto, 2018) makes it well suited for HVAC control. Indeed, RL has been widely used for advanced control of HVAC systems considering the uncertainties and multistage problem features (Sun et al., 2012, 2015; Yang et al., 2021). RL can enable efficient online computation by learning the control policies offline. Specifically, by learning the control policies offline, the online implementation of RL only requires to search the table to identify the mapping from the state space to the action space.

However, the obstacle to effective implementation of RL is that the learning may be computationally intensive or intractable when state and action space are large, either in terms of dimensionality (number of observations) or enumeration of action space (continuous vs. categorical.) To overcome the challenges, DRL has been popular in this domain. DRL proposes to combine the approximation and characterization capabilities of deep neural networks with RL techniques. Specially, deep neural networks are used to characterize the state or state action value functions. The benefits of DRL include reducing the quantity of

data required for training and reducing the memory required to store the policy for online implementation. The proliferation of smart meters has boosted the development of DRL-based HVAC controllers for smart buildings. We direct the readers to (Han et al., 2019; Mason and Grijalva, 2019; Wang and Hong, 2020; Yang et al., 2020; Yu et al., 2021). As an earlier work, Wei et al. (2017) has explored deep Q-network-based RL for single zone HVAC control. A Q-network is trained to approximate the state-value functions to overcome the computation burden related to the large state space. Yu et al. (2019) has explored the application of Deep Deterministic Policy Gradients (DDPG) (Lillicrap et al., 2015) for the management of building energy system including an HVAC system. The benefit of DDPG is that it can enable the decision making on the continuous state and action space so as to improve the control accuracy. While most of the existing works have focus on the control of HVAC system for single zone or single room, several recent works have explored the development of distributed RL for multizone commercial buildings (Yu et al., 2020; Hanumaiah and Genc, 2021). For example, Yu et al. (2020) studied the application of multiagent actor-attention-critic RL (Iqbal and Sha, 2019) for multizone HVAC control considering both thermal comfort and indoor air quality.

One major advantage of RL-based HVAC controllers over MPC is that it does not depend on explicit and accurate models, as in classical RL methods, a model of the environment may only ever be implicitly encoded in the inner layers of the policy or value prediction. In this way, RL methods alleviate the burden on the engineer for articulating an accurate model of the environment. It also allows for environment adaptation, i.e., flexibility to an environment that changes over time. The RL agent would just shift the value prediction it makes from actions over time.

RL methods are generally offline tools and require a rich set of data to compute the optimal control policies for a wide range of possible operation conditions. Moreover, the existing RL models including DRLs are still suffering from the computation intensity and scalability issues. Classical RL works best in cases where failure, at least at first, is relatively cheap. However, numerous studies elucidate ways that RL can be modified to reduce variability and increase safety. Arnold et al use Surprise Minimizing RL, a type of RL that rewards the agent for achieving states closer to those seen before, in optimizing building demand response (Arnold et al., 2021). Augmenting RL with a system simulation is important, as almost always a large part of the training can take place in a simulation environment which can use a rules-based heuristic as a starting point, and use exogenous parameters to create a distribution of unique systems to train on (i.e., domain randomization). Finn firmly established the field of meta-learning in RL, arguing that a technique like model agnostic meta-learning can train in the different simulations to approximate a starting distribution for policy network weight initializations (Finn et al., 2017). Jang et al. studies offline learning through offline-online RL (Jang et al., 2021) and Model Agnostic Meta Learning (MAML) (Jang et al., 2021) to get RL agents “hit the ground ready” using offline datasets and simulations.

### 3.4.3. *Learning-based MPC*

From the above literature, we note that MPC and RL tools both have pros and cons for HVAC control. MPC-based controllers can yield superior performance both in energy saving and ensuring human comfort due to the capability to incorporate the dynamic predictions regarding the disturbances. However, this is achieved at the cost of high online computation burden. Moreover, MPC-based controllers generally suffer from the modeling complexity. For example, the more elaborate thermal comfort model, such as the PMV is hard to be incorporated in the MPC-based control due the challenge to solve the optimization problem (Xu et al., 2017). In contrast, RL-based models can accommodate the modeling complexities but require a huge amount of data to ensure the consistence of performance under uncertainties. The current research trends have been the combination of MPC and RL so as to enjoy both of their pros. The researches are mainly in the following two aspects. One is to use ML to mimic the control policies output by MPC controllers. Specifically, the control of MPC controllers can be regarded as the parameterized policy with respect to the disturbance. To reduce the online computation burden and enable the deployment of MPC controller on simple hardware, the characterization capacity of ML, such as deep neural networks can be leverage to train parameterized control policies. This can yield control

policies of high energy-saving performance and human comfort with MPC at lower online implementation cost. For example, Drgoña et al. (2018) proposed an approximate MPC where deep time delay neural networks (TDNN) and regression tree-based regression models are used to mimic the control policy of MPC controller. Simulation results show that the approximate MPC only incurs a slight performance loss (3%) but reduces the online computation burden by a factor of 7. Later, a real-time implementation of the approximate MPC is developed and demonstrated in Drgoña et al. (2019) in an office building located in Hasselt, Belgium. Similarly, Karg and Lucia (2018) proposed to use deep learning networks to learn the control policies of MPC controller to enable low-cost implementation. Using ML tools to learn the control policies of MPC controllers can be regarded as the extension of the explicit MPC discussed above. When the problem is simple and the closed-form solution is available, the implementation of explicit MPC only requires to compute the parameters of the parameterized policy and therefore the online implementation only requires low online computation cost. By leveraging ML tools, we are able to achieve the same objective by training the parameterized policies using deep neural networks when the problem is complex and the explicit closed-form solution is not accessible. Another line of work to combine ML with MPC is to use ML to establish the models for some complex components with building or HVAC system. We direct the readers to Afram et al. (2017) for a comprehensive review. For example, neural networks can be trained to efficiently predict PMV (Ferreira et al., 2012; Ku et al., 2014) for a MPC controller. Since neural networks are generally characterized by nonlinear functions, the nonlinear solvers, such as genetic algorithms and particle swarm algorithms are often used to solve the resulting optimization problem. Compared with the original PMV models, neural network approximation can enable more efficient evaluation of the PMV index without requiring any iterations.

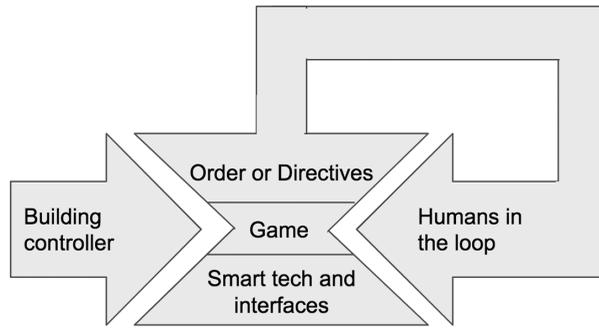
#### 3.4.4. Summary

Though MPC and RL are two main tools for HVAC control, there has not been significant implementation of them in buildings, where the primary controllers have been PID and rule-based ones. There are still many practical issues required to be addressed to enable the deployment of MPC and RL-based methods. Challenges related to classic MPC controllers include the development of models that can properly trade off the modeling complexity and accuracy. In addition, the performance of MPC-based controllers is closely related to the prediction accuracy of disturbance and the prediction horizon. Last but not the least, MPC-based controllers suffer from the high online computation cost and the implementation requires advanced micro-controllers. In contrast, RL-based controllers are free from the establishment of models and can enable efficient and low-cost implementation but requires enormous data to train the agent. Moreover, RL-based controllers cannot incorporate the possibly available dynamic predictions to enhance the performance. That is why most of the existing RL-based controllers are not comparable with MPC-based controllers in terms of energy-saving performance. Another critical issue related to MPC controller is that the safety cannot be ensured. Since the agent is trained on limited data, RL-based controller may draw the system to any possible state which may be prohibitive in practice. Combining ML with MPC-based controllers seems a promising solution as it can enjoy both the pros of ML and MPC. However, how well the control policies can be learned and how the network should be designed remain to be investigated. The operation of HVAC system is expected to provide a comfortable indoor environment for the occupants.

#### 3.4.5. Mechanism design

Compliance is the ability for people, in aggregate, to follow guidelines or recommendations. In any system that seeks to influence human behavior, lack of compliance is generally a large problem. Energy systems are no different: an empirical study of demand response found that noncompliance weakened the overall effect of the programs (Cappers et al., 2010).

Gamification of systems, i.e., the adding of an additional layer of intrinsic motivation to a core program, may be an important tool to boost compliance. The additional layer of motivation may come in the form of engagement-derived pleasure (i.e., interesting and exciting gameplay), competition



**Figure 6.** An example of a system in which gamification can be part of a control strategy.

(i.e., equally matched competitors), or social status (i.e., badges, ranks, and advancement). Hamari et al., 2014 rightly note the explosion in interest in gamification, from gamification around code learning to gamification around energy demand reduction and perform an extensive review of its efficacy, finding that generally gamification provides positive results. Figure 6 provides a visualization of the gamification pipeline and the purpose it serves.

There are some examples of gamifying office-related energy systems in the literature. Although not extensively, ML algorithms have been utilized in the gamification process to improve some individual components of the game. In one example, a game called “NTU Social Game” has been formulated around reductions in energy in a residential dormitory. Here, players compete against each other to have the lowest energy, and accrue points depending on how much less energy they consume. Prizes are not simply doled out according to whoever has the highest energy reductions; instead, players are entered into a prize pool if they are among the top third performing participants. A small number of prizes are allocated randomly among the winning pool as per the Vickrey Clarke Groves (VCG) auction mechanism. A design like this is shown to boost user engagement and directive compliance (Konstantakopoulos et al., 2019a, 2019b). The NTU Social Game has been used for energy reduction and to test different visualizations on player outcome (Spangher et al., 2019). It has also been used to conduct segmentation analysis of players into low, medium, and high energy-efficient groups for intelligent incentive planning (Das et al., 2019, 2020). The authors use a hybrid approach of supervised classification and unsupervised clustering, to create the characteristic clusters that satisfy the energy efficiency level. They also have incorporated explainability in the model by combining both the ML approaches.

An extension of this work, “OfficeLearn”, considers prices that are distributed across a day. Here, a building controller sets energy prices throughout the day, and players compete with each other on having the lowest cost of energy throughout the day. As before, the top performing individuals are entered into a prize pool and selected based on a VCG auction mechanism (Spangher et al., 2020). This paper presents an OpenAI gym environment for testing demand response with occupant level building dynamics, useful for standardizing RL algorithms implemented for the same purpose.

Other examples of energy-based games exist. In one game called “energy chickens”, competitors save energy from smart appliances for the health of virtual chickens; depending on how healthy the chickens are, they will generate probabilistically more or less eggs (Orland et al., 2014). In another game, “Energy Battle”, players in simulations of identical residential homes strive to reduce energy and players compete one to one without a VCG mechanism (Geelen et al., 2012). In a third game, “EnerGAware”, players also compete in a building simulation in which players, who are virtual cats, decide how to operate a house’s appliances with greater energy efficiency while not reducing any overall functions (Casals et al., 2020). The use of ML in energy social games, although not prominent before, is showing signs of increase, and is adequately justified since ML can have a profound use in coordinating multiagent systems such as a game.

### 3.5. Integration of buildings with power grids, microgrids, and energy markets

So far, we have discussed how modeling and control methods can be used to improve occupant comfort and increase the energy efficiency of buildings. However, buildings do not function in isolation and their operation can actually have significant impacts on other energy systems. The energy utilized in buildings come from power grids, and recently, from renewable sources installed at the local building/community level as well. In order to mitigate climate change, buildings must maximize their energy consumption from clean sources.

One way that buildings can achieve the above is by providing load flexibility to facilitate clean energy sources such as wind and solar that have variable generation and cannot be controlled at will to meet electricity demand, a mechanism commonly called demand response (DR). By making their electricity consumption flexible, buildings can adapt to variable generation curves and enable greater amounts of clean generation on the power grid. Buildings can capitalize on their load flexibility using two methods: by participating directly in DR programs run by utilities or by forming aggregations and participating in broader energy markets.

Another way that smart buildings can mitigate carbon emissions is by investing in local clean generation resources (like solar). Smart buildings are often equipped with distributed energy resources (DERs) such as rooftop solar, backup batteries, and electric vehicles (and chargers) to reduce demand charges or optimize grid consumption. If these resources are associated with a single building, they may not be fully utilized. Therefore, buildings should function as *prosumers*, i.e., use their electricity production capability to provide services to other buildings or to the power grid. Doing so can increase the utilization of clean energy sources, and incentivize additional investment in DERs thus reducing the reliance on grid-wide CO<sub>2</sub> emission. In the following parts, we will discuss DR in further details, and how ML plays a role in achieving the same. We will also cover on the prosumer nature of buildings.

#### 3.5.1. Demand response

We start off by discussing how individual buildings can use their load flexibility as DR resources. DR is a broad term for a variety of methods used to modulate electricity demand. These can include load response to time-varying electricity rates, financial incentives, or direct control of loads (like appliances and air-conditioning) by utilities.

**3.5.1.1. DR sources and timescales.** There are two primary sources of flexible loads in smart buildings: lighting and the HVAC system (Watson et al., 2006). Modifying energy consumption for either of these has a direct impact on the building occupants, as covered in the thermal comfort section (Section 3.1). Lighting control is a great DR resource since it does not have a *rebound* effect, i.e., curtailing lighting load during a DR event does not cause a spike after the DR period ends.

Flexible load can be modulated over timescales ranging from a few seconds to a few hours (Alstone et al., 2017). For example, smart buildings can *shift* load by preheating or precooling their indoor spaces, and can *shed* load altogether by reducing their lighting levels. They can also modulate their instantaneous load by adjusting fan speeds, and respond to high-frequency regulation signals in electricity markets.

**3.5.1.2. DR program mechanisms.** The two most common types of DR schemes are time-of-use (ToU) rates and curtailment contracts. Time-based rates aim to *shape* the energy consumption curve by driving long-term behavioral change and energy efficiency investments (Alstone et al., 2017). Utilities also run DR programs to procure short-term load curtailment, e.g., in situations when the electricity price spikes and it is too expensive to procure additional generation. These programs are typically structured as contracts where the buildings promise to curtail their energy consumption by a certain amount when called upon to do so, and the utility pays an incentive upfront based on the curtailment size. If the building is unable to curtail its consumption by that amount, it has to pay a penalty to the utility proportional to the shortfall. Readers are encouraged to refer to Albadi and El-Saadany (2007) for an overview of common DR mechanisms.

*3.5.1.3. ML for forecasting loads and baselines.* Buildings procure energy from utilities, who buy it from generators through long-term contracts or in wholesale energy markets. Energy consumption is variable, and utilities employ many ML methods to forecast electricity consumption accurately. Authors in Federal Energy Regulatory Commission (2020) provide an overview of the functioning of electricity markets. The electricity consumption forecasts are usually made for the aggregate load, and individual building loads are not as important for utilities. However, there are a couple of use cases for load forecasting at the building level.

First, in order to implement DR programs and incentivize loads to modulate their energy consumption, it is important to estimate the counterfactual energy consumption: what would the load have consumed without an incentive/intervention? This is considered as the baseline, and needs to be developed at the load level, i.e., building level. There are a number of ML methods used for developing baselines. Deb et al. (2017) conduct a review of time-series models for forecasting building energy consumption. They found that a variety of ML models have been used for this purpose: classical time-series methods such as ARIMA, SVM regression for short-term and long-term forecasts, as well as neural network-based methods. Jazaeri et al. (2016) compare different methods for computing baselines, such as regression, neural networks, polynomial interpolation, and classical methods (X of last Y days) used by utilities. They find that the ML methods emerge as the state-of-the-art for most of the cases.

A second use case of forecasting load at the building level is peak shaving, where buildings aim to forecast and then reduce their peak energy consumption in order to reduce the demand charges that they have to pay as a part of their electricity bill. Xu et al. (2019) develop a probabilistic forecasting methodology to predict the time of occurrence of magnitude of the peak load, where they quantify the probabilistic occurrence and magnitude of the peak load. Kim et al. (2019) use a variety of time-series methods to forecast the peak load demand of institutional buildings, with a dataset of buildings in South Korea.

*3.5.1.4. ML for estimating DR capability.* Once the baseline energy consumption is estimated, the utility/system operator can compare it with the actual consumption to determine the impact of the DR program, and in turn determine payments and penalties. However, most energy consumption by buildings is variable, and depends on occupancy and weather conditions. Additionally, the impact of load reduction measures such as reducing lighting and changing temperature setpoints are difficult to estimate ahead of time since the amount of energy curtailed will depend on the building state before the intervention. Some work has been done to quantify the effect of DR actions, e.g., Nghiem and Jones (2017) develop a Gaussian process model of how buildings respond to DR signals. In DR contracts, there might be a penalty associated with curtailment shortfall, and accurately estimating curtailment capability becomes an important task. Mathieu et al. (2011) use linear regression to develop models for estimating DR capability using time-of-week and outdoor air temperature buckets. It is important to have such DR capability estimates over longer time horizons, and Jung et al. (2014) develop a look-up table approach to predict future DR capability. They relate power consumption data to parameters such as time-of-day, day-of-week, occupancy, and outdoor temperatures and determine the load flexibility at a point when the forecast value of these parameters is available.

### *3.5.2. DR aggregations*

Load flexibility can be procured through utility programs as discussed previously, and also through direct market participation by loads. There has been a move in recent years to allow loads to submit bids for load reduction into wholesale electricity markets, and treat those bids similar to generation bids. Federal regulations, specifically Federal Energy Regulatory Commission (FERC) Order 2222 in the United States have mandated that distributed resources be allowed to participate in wholesale markets. The state of California has instituted a market mechanism called DR auction mechanism to allow loads to submit bids as ‘negative generation’. However, the magnitude of load flexibility provided by individual buildings is often too small for them to participate directly in electricity

markets. This is where aggregations can play a role: aggregators can bundle the load modulation capability of multiple buildings and use the aggregate capacity to submit bids to electricity markets. DR can also be coupled with other resources like generation and storage, and form a part of a *virtual power plant*.

**3.5.2.1. Value proposition of aggregations.** Aggregations can add value beyond simply providing a means of market participation, by spreading the risk of shortfall of curtailment capability. Burger et al. (2016) discuss the value of aggregators in electricity systems, and identify that apart from opportunistic and transitory value, aggregations also provide fundamental value, e.g., through economies of scale, that is independent of market or regulatory context. Since the load curtailment capability of buildings is variable, and each building faces a risk of having to pay penalties if it curtails less power than it promised to. However, different buildings will have different variability in their load curtailment capability, and aggregations can capitalize on this uncorrelatedness to minimize the penalty risk for all the participants (Agwan et al., 2021). Optimizing the complementary nature of buildings can significantly increase the marginal value of forming an aggregation.

**3.5.2.2. Optimizing market bids.** Once an aggregation is formed, it can optimize its bid in electricity markets using price forecasts and predictions of the available curtailment capacity. A variety of time-series methods are used by commercial entities that participate in these markets, and improving on these methods is an active area of research in both academia and industry as it can directly lead to higher revenues. Shinde and Amelin (2019) conduct a review of electricity price forecasting methods. Brusafferri et al. (2019) develop probabilistic forecasting methods for electricity prices using Bayesian deep learning methods, as opposed to commonly used point forecasts.

### 3.5.3. Prosumers and local energy markets

So far, we covered how buildings can use their load flexibility to provide services to the power grid, both individually and through aggregations. Additionally, buildings can directly consume energy from onsite clean resources. Smart buildings are often equipped with a variety of generation and storage resources such as rooftop solar, battery backups, and electric vehicle chargers, which can be used to reduce demand charges or optimize grid consumption. They can profit from these resources by supplying energy back to the grid, i.e., functioning as *prosumers*. A prosumer is an entity that is capable of energy production as well as consumption through the presence of local generation and energy storage devices.

**3.5.3.1. Value proposition of prosumers.** Investing in local energy resources is not always a profitable proposition: buildings have to invest in over-capacity generation to meet peak loads, and storage might only be used for peak shaving on a few days in a month or a few hours in a day. If these resources can be shared across multiple users, their utilization levels can increase which can further incentivize investment in clean energy sources. One way to facilitate such sharing is through local energy markets, where prosumers can trade energy with each other before procuring the aggregate balance from an outside source, such as the utility (Agwan, 2020).

**3.5.3.2. Organization of local energy markets.** There are two main challenges that come up while organizing local energy markets: selecting prosumers to form a market, and then managing the trades within the market. ML methods are used in tackling both of these.

First, we consider the problem of setting up a local energy market, and the methods used to address this challenge. Saad et al. (2011) use coalitional game theory to optimally aggregate prosumers in a cooperative microgrid. Wu et al. (2020) uses a mixed-integer linear programming to discuss the optimal sizing of energy resources in a microgrid. In the study by Quashie et al. (2018), the microgrid resource planning problem is cast as the upper level of a bilevel program which aims to optimize investment costs

along with system reliability. Research in this area mainly focuses on optimizing resource investments, and the problem of evaluating a new prosumer using numerical metrics is not as widely studied. While forming a local energy market, it is valuable to select participants that complement each other's net load curves. For example, a building that has a net load at a certain time would benefit from being in a market with a net generator at that time. The complementarity of market participants is essentially a measure of how well the electricity generation and demand curves match up. There have been game theoretic approaches to evaluating complementary participants, and the authors in Agwan et al. (2021) have developed a complementarity metric that uses net generation and load curves to evaluate the marginal value of adding new participants to the local energy markets.

Once a market has been set up, we need to facilitate energy trades between participants. This can be achieved by soliciting demand and supply bids from participants, and then settling the market. However, it may be an onerous task for buildings to generate such bids for a market. An alternative to this approach is to set a price for energy and then settle all trades at those prices. A number of previous studies develop price-based controls: Kim and Giannakis (2016) develop a price to minimize load variation, and Liu et al. (2017) develop a price based on supply-demand ratio to incentivize energy sharing.

While this price can be set in an iterative manner using a variety of distributed optimization techniques, it might be preferable to set it in a one-shot manner to avoid the time lag of back-and-forth communications. This method has been investigated in the literature, e.g., Jia et al. (2013) develop an online learning method to dynamically price electricity for retail consumers. They try to learn the response function of a number of price-sensitive consumers who modify their HVAC load to optimize their consumption. Agwan et al. (2021) approach this problem using RL, and develop a controller that can set day ahead electricity prices for a local market with a variety of prosumers.

In summary, ML can aid in predicting the energy load of buildings, enable DR schemes via load flexibility, and can accelerate the energy market set up with buildings as prosumers.

#### **4. Limitations of ML in the Context of Smart Buildings**

Although ML techniques have proliferated successfully across many fields, the lack of success of ML in some applications bears note. There are some notable applications where ML does not work well such as safety critical applications, data-poor applications; we describe the fundamental barriers to success in these domains. Arguably, most of these barriers also exist when ML is employed in smart buildings.

##### **4.1. 51/49 vs. 99/01 problems**

Does the area of interest require 51% and 49% failure of the control algorithm in its trials, or does it require 99% success and 1% failure? Examples of the former may include stock investing, product recommendation, efficient allocation of computational tasks, and image classification. Examples of the latter may include medical decisions, decisions in autonomous driving vehicles etc.

In smart buildings, the applications fall into both the above categories, and it often becomes harder for ML algorithms to achieve 99/01 performance without jeopardizing robustness and safety. Advancements in ML that may help to address 99/01 problems in buildings are safety-guarantees in ML, robust worst case minimizing RL, proper explainability, etc.

##### **4.2. Data availability**

Does the domain provide rich, reliable, and accurate data? The shining examples of ML success often occur in domains where this is the case: RL techniques started to shine when Atari game simulators were developed and could be run millions of times, image ML began to flourish when large image datasets were provided and standardized, and natural language processing began to bear fruit when datasets like the entirety of Wikipedia could be processed and trained on. Most recently, large language models (LLMs) have captured public interest with their performance, largely based on their training on an incredibly large

corpus of written materials. In contrast, building energy datasets are not nearly as large, rich, or diverse, and so many of the techniques that work well in different domains may not translate well into the built environment (Ličina et al., 2018; Fierro et al., 2019; Miller et al., 2020; Dong et al., 2022). A fruitful area of work may include preparing and standardizing large datasets for ease of energy research. Also, given that each building is unique, techniques such as domain transfer learning may be leveraged to enhance the utility of building-specific datasets. Also, advancements in RL that can work well in data-poor regimes may include generative models, regularized and sparse models, etc.

### 4.3. Pathway to physical implementation

Does the domain have a physical interface that can use the output of ML models? Oftentimes, we may put a lot of work into making a prediction model accurate, but a building from the 1970s might not have the physical infrastructure to house the model, respond to its predictions, or reliably provide state information. A smart appliance may be limited in its scope and effect. To this extent, building operating systems (OS) are newer developments that seek to unify buildings into one software control scheme. Another important function of building OSes should be to provide to the outside world an API that does not need to be aware of inner complexities of the building but can provide certain data about the building (aggregate energy usage) and receive certain grid information like price and state. A third important function of the building OS is to provide unified controls. To this extent, ML research in controls is inherently closer to many of the applications that a building may need.

### 4.4. Computational cost

Researchers need to be aware of the tradeoff between energy saved via efficient building operations, and the carbon emissions from the compute used to achieve the above. Machine learning is notoriously compute-heavy, and some ML models can require large amounts of energy for training and deployment. The information technology sector is responsible for 2–3% of global greenhouse gas emissions, and the benefits of applying ML in buildings to reduce emissions should be compared to the emissions caused by these models themselves, whether in training or in deployment. Further analysis of these emissions can be found in Kaack et al. (2022).

## 5. Conclusion and Future Work

During the life cycle of a building including the design, construction, operation and maintenance, retrofitting, and disposal, there are many factors that contribute toward a building's energy consumption, efficiency, and embodied carbon. In this paper, we primarily evaluated improving energy efficiency in buildings via ML during design, operation, and maintenance stage. From the well-being of an individual, to building's occupants, to the building operations and maintenance, and finally to the electric grid impact and energy market economic response, we discussed all the different facets of the smart building ecosystem.

While ML shows promising results in facilitating the energy improvement of building operation and maintenance, there are still several challenges and works to alleviate them ahead. The foremost challenge of all ML approaches is the time and effort to collect reliable data. ML models require a great amount of data, and the quality and the selection of the data greatly affects the quality of the output model. In thermal comfort modeling, data are even harder to come by as it uses intrusive methods that require occupants' input. Additionally, thermal comfort data are inherently class-imbalanced, which may result in ill-conditioned models. In building modeling, realistic data such as occupancy, temperature settings, and schedules are crucial, and without them, the outcomes of models may be inaccurate or have a high level of uncertainty. Furthermore, the realistic data required may be completely dependent on the application in which the data are needed. Although many researchers have been collecting this realistic data, there is no consistency between companies or research entities. As such, utilizing the collected data across many

applications is time-consuming and ultimately infeasible. A potential solution for the above data issue is to develop a standard for collecting and storing building data, potentially as a public data repository from which crowd-sourced data can be posted and requested. An example of a data standard that could suffice is the BRICK schema (Balaji et al., 2018), which is a metadata schema that represents buildings' sensor information and the relationships between them. Gray box models developed on the data along with building physics can be the ultimate goal. Synthetic data-generation methods are also being proposed as a solution to the data insufficiency problem (Das and Spanos, 2022b).

Generalization is another open challenge for building applications. Firstly, thermal physiology and expectations vary significantly between human occupants. If all occupants' thermal comfort models need to be individually trained, it will be computationally expensive. As a result, the need for more generalized models that can represent all occupants arises. Secondly, as all buildings are inherently different in design and operation, models, and control strategies from one building are difficult to apply to other buildings. Hence, transfer learning and domain adaptation are some open areas of research that help create more generalized building models.

Finally, thermal comfort models, building models, control algorithms, and grid applications exist on different platforms and do not have a standard interface through which they can all be used. Developing a standardized platform for each application and a standardize interface protocol for the above information can benefit the research communities, as the application of the above techniques are correlated.

As data become more widely available through IoT and ML becomes more computationally feasible, more and more success can be achieved in enhancing occupant comfort and building energy performance. In the future, ML could provide resiliency in the built environment and assist in optimizing occupants' satisfaction, energy use and cost.

**Author contribution.** All the authors conceptualized the project. H.P.D. coordinated the ideation, analysis, and writing effort, and other project administration efforts. H.P.D., Y-W.L., U.A., L.S., A.D., Y.Y., J.D., and A.C. were responsible for all formal analysis, investigation, and writing of the original draft. S.S. and C.J.S. supervised the project, and were responsible for reviewing and editing of the paper.

**Competing interest.** The authors declare none.

**Data availability statement.** Review papers used in this study are available publicly.

**Ethics statement.** The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

**Funding statement.** This research is funded by the Republic of Singapore's National Research Foundation through a grant to the Berkeley Education Alliance for Research in Singapore (BEARS) for the Singapore-Berkeley Building Efficiency and Sustainability in the Tropics (SinBerBEST) Program. BEARS has been established by the University of California, Berkeley as a center for intellectual excellence in research and education in Singapore.

## References

- American Institute of Aeronautics and Astronautics** (2022) Digital twin: Definition & value - An AIAA and AIA position paper. Available at <https://www.aiaa-aerospace.org/report/digital-twin-paper/>.
- Afram A and Janabi-Sharifi F** (2014) Theory and applications of HVAC control systems—A review of model predictive control (MPC). *Building and Environment* 72, 343–355.
- Afram A, Janabi-Sharifi F, Fung AS and Raahemifar K** (2017) Artificial neural network (ANN) based model predictive control (MPC) and optimization of HVAC systems: A state of the art review and case study of a residential HVAC system. *Energy and Buildings* 141, 96–113.
- Agwan U** (2020) Optimal prosumer aggregations: Design and modeling.
- Agwan U, Poolla K and Spanos CJ** (2021) Optimal composition of prosumer aggregations. In *2021 IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe)*. IEEE, pp. 1–5.
- Agwan U, Spangher L, Arnold W, Srivastava T, Poolla K and Spanos CJ** (2021) Pricing in prosumer aggregations using reinforcement learning. In *Proceedings of the Twelfth ACM International Conference on Future Energy Systems*, pp. 220–224.
- Agwan U, Spanos CJ and Poolla K** (2021) Asset participation and aggregation in incentive-based demand response programs. In *IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*. IEEE, pp. 89–94.

- Akametalu AK, Fisac JF, Gillula JH, Kaynama S, Zeilinger MN and Tomlin CJ** (2014) Reachability-based safe learning with Gaussian processes. In *53rd IEEE Conference on Decision and Control*. IEEE, pp. 1424–1431.
- Albadi MH and El-Saadany EF** (2007) Demand response in electricity markets: An overview. In *2007 IEEE Power Engineering Society General Meeting*. IEEE, pp. 1–5.
- Allemang D, Hendlar J and Gandon F** (2011) Expert modeling in owl.
- Al-Ogaili AS, Hashim TJT, Rahmat NA, Ramasamy AK, Marsadek MB, Faisal M and Hannan MA** (2019) Review on scheduling, clustering, and forecasting strategies for controlling electric vehicle charging: Challenges and recommendations. *IEEE Access* 7, 128353–128371.
- Alsalem F, Tesfay MK, Rafaie M, Sinkar K, Besarla D and Arunasalam P** (2020) An IoT framework for modeling and controlling thermal comfort in buildings. *Frontiers in Built Environment* 6, 87.
- Alstone P, Potter J, Piette MA, Schwartz P, Berger MA, Dunn LN, Smith SJ, Sohn MD, Aghajanzadeh A, Stensson S, et al.** (2017) 2025 California demand response potential study-charting california's demand response future. Final report on phase 2 results, Tech. Report, Lawrence Berkeley National Lab (LBNL), Berkeley, CA, USA.
- Arendt K, Johansen A, Jørgensen BN, Kjærgaard MB, Mattera CG, Sangogboye FC, Schwegel JH and Veje CT** (2018) Room-level occupant counts, airflow and CO<sub>2</sub> data from an office building. In *Proceedings of the First Workshop on Data Acquisition to Analysis*, pp. 13–14.
- Arnold W, Srivastava T, Spangher L, Agwan U and Spanos C** (2021) Adapting surprise minimizing reinforcement learning techniques for transactive control. In *Proceedings of the Twelfth ACM International Conference on Future Energy Systems, e-Energy '21*. New York, NY: Association for Computing Machinery, pp. 488–492. <https://doi.org/10.1145/3447555.3466590>.
- Aviv Y** (2003) A time-series framework for supply-chain inventory management. *Operations Research* 51(2), 210–227.
- Bacher P and Madsen H** (2011) Identifying suitable models for the heat dynamics of buildings. *Energy and Buildings* 43(7), 1511–1522.
- Balaji B, Bhattacharya A, Fierro G, Gao J, Gluck J, Hong D, Johansen A, Koh J, Ploennigs J, Agarwal Y, et al.** (2018) Brick: Metadata schema for portable smart building applications. *Applied Energy* 226, 1273–1292.
- Blum D, Arroyo J, Huang S, Drgoña J, Jorissen F, Walnum HT, Chen Y, Benne K, Vrabie D, Wetter M, et al.** (2021) Building optimization testing framework (BOPTTEST) for simulation-based benchmarking of control strategies in buildings. *Journal of Building Performance Simulation* 14(5), 586–610.
- Botvinick M, Ritter S, Wang JX, Kurth-Nelson Z, Blundell C and Hassabis D** (2019) Reinforcement learning, fast and slow. *Trends in Cognitive Sciences* 23(5), 408–422.
- Brackney LJ, Florita AR, Swindler AC, Polese LG and Brunemann GA** (2012) Design and performance of an image processing occupancy sensor. In *Proceedings of the Second International Conference on Building Energy and Environment 2012/987 Topic 10. Intelligent Buildings and Advanced Control Techniques*.
- Brandi S, Fiorentini M and Capozzoli A** (2022) Comparison of online and offline deep reinforcement learning with model predictive control for thermal energy management. *Automation in Construction* 135, 104128.
- Bre F, Roman N and Fachinotti VD** (2020) An efficient metamodel-based method to carry out multi-objective building performance optimizations. *Energy and Buildings* 206, 109576.
- Brusaferrri A, Matteucci M, Portolani P and Vitali A** (2019) Bayesian deep learning based method for probabilistic forecast of day-ahead electricity prices. *Applied Energy* 250, 1158–1175.
- Burger S, Chaves-Avila JP, Battle C, Perez-Arriaga I** (2016) The value of aggregators in electricity systems.
- Cai S, Mao Z, Wang Z, Yin M and Karniadakis GE** (2022) Physics-informed neural networks (pinns) for fluid mechanics: A review. *Acta Mechanica Sinica* 37, 1727–1738.
- Cappers P, Goldman C and Kathan D** (2010) Demand response in U.S. electricity markets: Empirical evidence. *Energy* 35(4), 1526–1535. <https://doi.org/10.1016/j.energy.2009.06.029>; Available at <https://www.sciencedirect.com/science/article/pii/S0360544209002461>.
- Casals M, Gangoletts M, Macarulla M, Forcada N, Fuertes A and Jones RV** (2020) Assessing the effectiveness of gamification in reducing domestic energy consumption: Lessons learned from the EnerGAware project. *Energy and Buildings* 210, 109753.
- Chai Q, Wang H, Zhai Y and Yang L** (2020) Using machine learning algorithms to predict occupants' thermal comfort in naturally ventilated residential buildings. *Energy and Buildings* 217, 109937.
- Chakrabarty A, Maddalena E, Qiao H and Laughman C** (2021) Scalable Bayesian optimization for model calibration: Case study on coupled building and HVAC dynamics. *Energy and Buildings* 253, 111460.
- Chaudhuri T, Zhai D, Soh YC, Li H and Xie L** (2018) Thermal comfort prediction using normalized skin temperature in a uniform built environment. *Energy and Buildings* 159, 426–440.
- Chawla NV, Bowyer KW, Hall LO and Kegelmeyer WP** (2002) Smote: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research* 16, 321–357.
- Chen X, Yang H and Sun K** (2017) Developing a meta-model for sensitivity analyses and prediction of building performance for passively designed high-rise residential buildings. *Applied Energy* 194, 422–439.
- Chennapragada A, Periyakoil D, Das HP and Spanos CJ** (2022) Time series-based deep learning model for personal thermal comfort prediction. In *Proceedings of the Thirteenth ACM International Conference on Future Energy Systems*, pp. 552–555.
- Cheung T, Graham LT and Schiavon S** (2022) Impacts of life satisfaction, job satisfaction and the big five personality traits on satisfaction with the indoor environment. *Building and Environment* 212, 108783.

- Cheung TC, Schiavon S, Gall ET, Jin M and Nazaroff WW** (2017) Longitudinal assessment of thermal and perceived air quality acceptability in relation to temperature, humidity, and CO<sub>2</sub> exposure in Singapore. *Building and Environment* 115, 80–90.
- Cheung T, Schiavon S, Parkinson T, Li P and Brager G** (2019) Analysis of the accuracy on PMV–PPD model using the ASHRAE global thermal comfort database II. *Building and Environment* 153, 205–217.
- Chong A, Gu Y and Jia H** (2021) Calibrating building energy simulation models: A review of the basics to guide future work. *Energy and Buildings* 253, 111533.
- Clarke JA** (2007) *Energy Simulation in Building Design*. London: Routledge.
- Coakley D, Raftery P and Keane M** (2014) A review of methods to match building energy simulation models to measured data. *Renewable and Sustainable Energy Reviews* 37, 123–141.
- Costabal FS, Yang Y, Perdikaris P, Hurtado DE and Kuhl E** (2020) Physics-informed neural networks for cardiac activation mapping. *Frontiers in Physics* 8. <https://doi.org/10.3389/fphy.2020.00042>. Available at <https://www.frontiersin.org/article/10.3389/fphy.2020.00042>.
- Crawley DB, Lawrie LK, Winkelmann FC, Buhl WF, Huang YJ, Pedersen CO, Strand RK, Liesen RJ, Fisher DE, Witte MJ, et al.** (2001) Energyplus: Creating a new-generation building energy simulation program. *Energy and Buildings* 33(4), 319–331.
- Das HP, Konstantakopoulos IC, Manasawala AB, Veeravalli T, Liu H and Spanos CJ** (2019) A novel graphical lasso based approach towards segmentation analysis in energy game-theoretic frameworks. In *18th IEEE International Conference On Machine Learning And Applications (ICMLA)*. IEEE, pp. 1702–1709.
- Das HP, Konstantakopoulos I, Manasawala AB, Veeravalli T, Liu H and Spanos CJ** (2020) Do occupants in a building exhibit patterns in energy consumption? Analyzing clusters in energy social games. In *Workshop on Tackling Climate Change with Machine Learning, Conference on Neural Information Processing Systems (NeurIPS)*.
- Das HP, Schiavon S and Spanos CJ** (2021) Unsupervised personal thermal comfort prediction via adversarial domain adaptation. In *Proceedings of the 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pp. 230–231.
- Das HP and Spanos C** (2022a) Improved dequantization and normalization methods for tabular data pre-processing in smart buildings. In *ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys)*.
- Das HP and Spanos CJ** (2022b) Conditional synthetic data generation for personal thermal comfort models. Preprint, [arXiv:2203.05242](https://arxiv.org/abs/2203.05242).
- Das HP, Tran R, Singh J, Yue X, Tison G, Sangiovanni-Vincentelli A and Spanos CJ** (2021) Conditional synthetic data generation for robust machine learning applications with limited pandemic data. Preprint, [arXiv:2109.06486](https://arxiv.org/abs/2109.06486).
- Daw A, Karpatne A, Watkins W, Read J and Kumar V** (2017) Physics-guided neural networks (PGNN): An application in lake temperature modeling. Preprint, [arXiv:1710.11431](https://arxiv.org/abs/1710.11431).
- Deb C, Zhang F, Yang J, Lee SE and Shah KW** (2017) A review on time series forecasting techniques for building energy consumption. *Renewable and Sustainable Energy Reviews* 74, 902–924.
- Devonport A and Arcak M** (2020) Estimating reachable sets with scenario optimization. In *Learning for Dynamics and Control*. PMLR, pp. 75–84.
- Devonport A, Yang F, El Ghaoui L and Arcak M** (2021a) Data-driven reachability analysis with christoffel functions. In *60th IEEE Conference on Decision and Control (CDC), 2021*, pp. 5067–5072. <https://doi.org/10.1109/CDC45484.2021.9682860>
- Devonport A, Yang F, Ghaoui LE and Arcak M** (2021b) Data-driven reachability analysis and support set estimation with christoffel functions. Preprint, [arXiv:2112.09995](https://arxiv.org/abs/2112.09995).
- Devonport A, Yin H and Arcak M** (2020) Bayesian safe learning and control with sum-of-squares analysis and polynomial kernels. In *2020 59th IEEE Conference on Decision and Control (CDC)*, IEEE, pp. 3159–3165.
- Diraco G, Leone A and Siciliano P** (2015) People occupancy detection and profiling with 3D depth sensors for building energy management. *Energy and Buildings* 92, 246–266.
- Dong B, Liu Y, Mu W, Jiang Z, Pandey P, Hong T, Olesen B, Lawrence T, O’Neil Z, Andrews C, et al.** (2022) A global building occupant behavior database. *Scientific Data* 9(1), 1–15.
- Dong B, O’Neill Z and Li Z** (2014) A BIM-enabled information infrastructure for building energy fault detection and diagnostics. *Automation in Construction* 44, 197–211.
- Drgoňa J, Arroyo J, Figueroa IC, Blum D, Arendt K, Kim D, Ollé EP, Oravec J, Wetter M, Vrabie DL and Helsen L** (2020) All you need to know about model predictive control for buildings, *Annual Reviews in Control* 50, 190–232. <https://doi.org/10.1016/j.arcontrol.2020.09.001>. Available at <https://www.sciencedirect.com/science/article/pii/S1367578820300584>.
- Drгона J, Helsen L and Vrabie D** (2019) Stripping off the implementation complexity of physics-based model predictive control for buildings via deep learning. *Advances in Neural Information Processing Systems*.
- Drgoňa J, Kvasnica M, Klaučo M and Fikar M** (2013) Explicit stochastic MPC approach to building temperature control. In *52nd IEEE Conference on Decision and Control*, pp. 6440–6445. <https://doi.org/10.1109/CDC.2013.6760908>.
- Drgoňa J, Picard D, Kvasnica M and Helsen L** (2018) Approximate model predictive building control via machine learning. *Applied Energy* 218, 199–216.
- Dridi J, Amayri M and Bouguila N** (2022) Transfer learning for estimating occupancy and recognizing activities in smart buildings. *Building and Environment* 217, 109057.
- Eastman C** (1975) The use of computers instead of drawings in building design. *AIA Journal* 63(3), 46–50.
- Eisenhower B, O’Neill Z, Narayanan S, Fonoberov VA and Mezić I** (2012) A methodology for meta-model based optimization in building energy models. *Energy and Buildings* 47, 292–301.

- Esrafilian-Najafabadi M and Haghghat F** (2021) Occupancy-based HVAC control systems in buildings: A state-of-the-art review. *Building and Environment* 197, 107810.
- Fanger PO** (1970) *Thermal Comfort: Analysis and Applications in Environmental Engineering*. New York: McGraw-Hill.
- Federal Energy Regulatory Commission** (2020) *Energy Primer: A Handbook for Energy Market Basics*. Available at [https://www.ferc.gov/sites/default/files/2020-06/energy-primer-2020\\_0.pdf](https://www.ferc.gov/sites/default/files/2020-06/energy-primer-2020_0.pdf).
- Ferreira P, Ruano A, Silva S and Conceicao E** (2012) Neural networks based predictive control for thermal comfort and energy savings in public buildings. *Energy and Buildings* 55, 238–251.
- Fierro G, Pritoni M, AbdelBaky M, Lengyel D, Leyden J, Prakash A, Gupta P, Raftery P, Peffer T, Thomson G, et al.** (2019) Mortar: An open testbed for portable building analytics. *ACM Transactions on Sensor Networks (TOSN)* 16(1), 1–31.
- Filippopolitis A, Oliff W and Loukas G** (2016) Bluetooth low energy based occupancy detection for emergency management. In *2016 15th International Conference on Ubiquitous Computing and Communications and 2016 International Symposium on Cyberspace and Security (IUC-CSS)*. IEEE, pp. 31–38.
- Finn C, Abbeel P and Levine S** (2017) Model-agnostic meta-learning for fast adaptation of deep networks. CoRR abs/1703.03400. Preprint, [arXiv:1703.03400](https://arxiv.org/abs/1703.03400). Available at <http://arxiv.org/abs/1703.03400>.
- Fritzson P and Engelson V** (1998) Modelica—A unified object-oriented language for system modeling and simulation. In *European Conference on Object-Oriented Programming*. Springer, pp. 67–90.
- Gan VJ, Wong H, Tse KT, Cheng JC, Lo IM and Chan CM** (2019) Simulation-based evolutionary optimization for energy-efficient layout plan design of high-rise residential buildings. *Journal of Cleaner Production* 231, 1375–1388.
- Gao N, Shao W, Rahaman MS, Zhai J, David K and Salim FD** (2021) Transfer learning for thermal comfort prediction in multiple cities. *Building and Environment* 195, 107725.
- Geelen D, Keyson D, Boess S and Brezet H** (2012) Exploring the use of a game to stimulate energy saving in households, *Journal of Design Research* 10 (1-2), 102–120.
- Graham LT, Parkinson T and Schiavon S** (2021) Lessons learned from 20 years of CBE’s occupant surveys. *Buildings and Cities* 2(1), 166–184.
- Grieves M and Vickers J** (2017) *Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems*. Cham: Springer International Publishing, pp. 85–113. [https://doi.org/10.1007/978-3-319-38756-7\\_4](https://doi.org/10.1007/978-3-319-38756-7_4).
- Haala N and Kada M** (2010) An update on automatic 3D building reconstruction. *ISPRS Journal of Photogrammetry and Remote Sensing* 65(6), 570–580.
- Hamari J, Koivisto J and Sarsa H** (2014) Does gamification work? – A literature review of empirical studies on gamification. In *47th Hawaii International Conference on System Sciences, 2014*, pp. 3025–3034. <https://doi.org/10.1109/HICSS.2014.377>.
- Han M, May R, Zhang X, Wang X, Pan S, Yan D, Jin Y and Xu L** (2019) A review of reinforcement learning methodologies for controlling occupant comfort in buildings. *Sustainable Cities and Society* 51, 101748.
- Hanumaiah V and Genc S** (2021) Distributed multi-agent deep reinforcement learning framework for whole-building HVAC control. Preprint, [arXiv:2110.13450](https://arxiv.org/abs/2110.13450).
- Hensen JL and Lamberts R** (2012) *Building Performance Simulation for Design and Operation*. London: Routledge.
- Hong T, Wang Z, Luo X and Zhang W** (2020) State-of-the-art on research and applications of machine learning in the building life cycle. *Energy and Buildings* 212, 109831.
- Hong T, Yan D, D’Oca S and Chen Cf** (2017) Ten questions concerning occupant behavior in buildings: The big picture. *Building and Environment* 114, 518–530. <https://doi.org/10.1016/j.buildenv.2016.12.006>. Available at <https://www.sciencedirect.com/science/article/pii/S0360132316304851>.
- Huang W and Zheng H** (2018) *Architectural Drawings Recognition and Generation through Machine Learning*. ACADIA.
- Ioannou P and Fidan B** (2006) *Adaptive Control Tutorial*. SIAM.
- Iqbal S and Sha F** (2019) Actor-attention-critic for multi-agent reinforcement learning. In *International Conference on Machine Learning*. PMLR, pp. 2961–2970.
- Jang D, Spangher L, Khattar M, Agwan U and Spanos C** (2021) Using meta reinforcement learning to bridge the gap between simulation and experiment in energy demand response. In *Proceedings of the Twelfth ACM International Conference on Future Energy Systems, e-Energy ’21*. New York, NY: Association for Computing Machinery, pp. 483–487. <https://doi.org/10.1145/3447555.3466589>.
- Jang D, Spangher L, Srivastava T, Khattar M, Agwan U, Nadarajah S and Spanos C** (2021) Offline-online reinforcement learning for energy pricing in office demand response: Lowering energy and data costs. In *Proceedings of the 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, BuildSys ’21*. New York, NY: Association for Computing Machinery, pp. 131–139. <https://doi.org/10.1145/3486611.3486668>.
- Jazaeri J, Alpcan T, Gordon R, Brandao M, Hoban T and Seeling C** (2016) Baseline methodologies for small scale residential demand response. In *2016 IEEE Innovative Smart Grid Technologies-Asia (ISGT-Asia)*. IEEE, pp. 747–752.
- Jia R, Jin B, Jin M, Zhou Y, Konstantakopoulos IC, Zou H, Kim J, Li D, Gu W, Arghandeh R, et al.** (2018) Design automation for smart building systems. *Proceedings of the IEEE* 106(9), 1680–1699.
- Jia L, Zhao Q and Tong L** (2013) Retail pricing for stochastic demand with unknown parameters: An online machine learning approach. In *51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, pp. 1353–1358.
- Jung D, Krishna VB, Temple WG and Yau DK** (2014) Data-driven evaluation of building demand response capacity. In *IEEE International Conference on Smart Grid Communications (SmartGridComm)*. IEEE, pp. 541–547.

- Kaack LH, Donti PL, Strubell E, Kamiya G, Creutzig F and Rolnick D** (2022) Aligning artificial intelligence with climate change mitigation. *Nature Climate Change* 12(6), 518–527.
- Karg B and Lucia S** (2018) Deep learning-based embedded mixed-integer model predictive control. In *European Control Conference (ECC)*. IEEE, pp. 2075–2080.
- Kathirgamanathan A, De Rosa M, Mangina E and Finn DP** (2021) Data-driven predictive control for unlocking building energy flexibility: A review. *Renewable and Sustainable Energy Reviews* 135, 110120.
- Khalil M, McGough S, Pourmirza Z, Pazhoohesh M and Walker S** (2021) Transfer learning approach for occupancy prediction in smart buildings In *12th International Renewable Engineering Conference (IREC)*. IEEE, pp. 1–6.
- Kim S-J and Giannakis GB** (2016) An online convex optimization approach to real-time energy pricing for demand response. *IEEE Transactions on Smart Grid* 8(6), 2784–2793.
- Kim J, Schiavon S and Brager G** (2018) Personal comfort models—a new paradigm in thermal comfort for occupant-centric environmental control. *Building and Environment* 132, 114–124.
- Kim Y, Son H-g and Kim S** (2019) Short term electricity load forecasting for institutional buildings. *Energy Reports* 5, 1270–1280.
- Kim J, Zhou Y, Schiavon S, Raftery P and Brager G** (2018) Personal comfort models: Predicting individuals' thermal preference using occupant heating and cooling behavior and machine learning. *Building and Environment* 129, 96–106.
- Kissas G, Yang Y, Hwuang E, Witschey WR, Detre JA and Perdikaris P** (2020) Machine learning in cardiovascular flows modeling: Predicting arterial blood pressure from non-invasive 4D flow MRI data using physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering* 358, 112623. <https://doi.org/10.1016/j.cma.2019.112623>. Available at <https://www.sciencedirect.com/science/article/pii/S0045782519305055>.
- Klaučo M and Kvasnica M** (2014) Explicit MPC approach to PMV-based thermal comfort control In *53rd IEEE Conference on Decision and Control*. IEEE, pp. 4856–4861.
- Konstantakopoulos IC, Barkan AR, He S, Veeravalli T, Liu H and Spanos C** (2019a) A deep learning and gamification approach to improving human-building interaction and energy efficiency in smart infrastructure. *Applied Energy* 237, 810–821.
- Konstantakopoulos IC, Das HP, Barkan AR, He S, Veeravalli T, Liu H, Manasawala AB, Lin Y-W and Spanos CJ** (2019b) Design, benchmarking and explainability analysis of a game-theoretic framework towards energy efficiency in smart infrastructure. Preprint, [arXiv:1910.07899](https://arxiv.org/abs/1910.07899).
- Kontou E, Yin Y and Ge Y-E** (2017) Cost-effective and ecofriendly plug-in hybrid electric vehicle charging management. *Transportation Research Record* 2628(1), 87–98.
- Kraft M, Aszkowski P, Pieczyński D and Fularz M** (2021) Low-cost thermal camera-based counting occupancy meter facilitating energy saving in smart buildings. *Energies* 14(15), 4542.
- Ku K, Liaw J, Tsai M and Liu T** (2014) Automatic control system for thermal comfort based on predicted mean vote and energy saving. *IEEE Transactions on Automation Science and Engineering* 12(1), 378–383.
- Lafthiev E and Nikovski D** (2016) An IoT system to estimate personal thermal comfort. In *2016 IEEE 3rd World Forum on Internet of Things (WF-IoT)*. IEEE, pp. 672–677.
- Lee S, Chon Y, Kim Y, Ha R and Cha H** (2013) Occupancy prediction algorithms for thermostat control systems using mobile devices. *IEEE Transactions on Smart Grid* 4(3), 1332–1340.
- Lee J and Ham Y** (2021) Physiological sensing-driven personal thermal comfort modelling in consideration of human activity variations. *Building Research & Information* 49(5), 512–524.
- Lee ES and Selkowitz SE** (1994) The design and evaluation of integrated envelope and lighting control strategies for commercial buildings.
- Li P, Cui H, Khan A, Raza U, Piechocki R, Doufexi A and Farnham T** (2021) Deep transfer learning for WiFi localization. In *IEEE Radar Conference (RadarConf21)*. IEEE, pp. 1–5.
- Ličina VF, Cheung T, Zhang H, De Dear R, Parkinson T, Arens E, Chun C, Schiavon S, Luo M, Brager G, et al.** (2018) Development of the ASHRAE global thermal comfort database ii. *Building and Environment* 142, 502–512.
- Lillicrap TP, Hunt JJ, Pritzel A, Heess N, Erez T, Tassa Y, Silver D and Wierstra D** (2015) Continuous control with deep reinforcement learning. Preprint, [arXiv:1509.02971](https://arxiv.org/abs/1509.02971).
- Lin Y-W, Tang TLE and Spanos CJ** (2021) Hybrid approach for digital twins in the built environment. In *Proceedings of the Twelfth ACM International Conference on Future Energy Systems, e-Energy '21*. New York, NY: Association for Computing Machinery, pp. 450–457. <https://doi.org/10.1145/3447555.3466585>
- Liu K, Nie T, Liu W, Liu Y and Lai D** (2020) A machine learning approach to predict outdoor thermal comfort using local skin temperatures. *Sustainable Cities and Society* 59, 102216.
- Liu S, Schiavon S, Das HP, Jin M and Spanos CJ** (2019) Personal thermal comfort models with wearable sensors. *Building and Environment* 162, 106281.
- Liu N, Yu X, Wang C, Li C, Ma L and Lei J** (2017) Energy-sharing model with price-based demand response for microgrids of peer-to-peer prosumers. *IEEE Transactions on Power Systems* 32(5), 3569–3583.
- Ljung L** (1998) System identification. In *Signal Analysis and Prediction*. Springer, pp. 163–173.
- Ljung L, Andersson C, Tiels K and Schön TB** (2020) Deep learning and system identification. *IFAC-PapersOnLine* 53(2), 1175–1181.
- Liu S, Wang W, Lin C and Hameen EC** (2019) Data-driven simulation of a thermal comfort-based temperature set-point control with ASHRAE RP884. *Building and Environment* 156, 137–146.

- Martins LA, Soebarto V and Williamson T** (2022) A systematic review of personal thermal comfort models. *Building and Environment* 207, 108502.
- Mason K and Grijalva S** (2019) A review of reinforcement learning for autonomous building energy management. *Computers & Electrical Engineering* 78, 300–312.
- Mathieu JL, Price PN, Kilicote S and Piette MA** (2011) Quantifying changes in building electricity use, with application to demand response. *IEEE Transactions on Smart Grid* 2(3), 507–518.
- Matni N, Proutiere A, Rantzer A and Tu S** (2019) From self-tuning regulators to reinforcement learning and back again. In *IEEE 58th Conference on Decision and Control (CDC)*. IEEE, pp. 3724–3740.
- Miller C, Kathirgamanathan A, Picchetti B, Arjunan P, Park JY, Nagy Z, Raftery P, Hobson BW, Shi Z and Meggers F** (2020) The building data genome project 2, energy meter data from the ASHRAE great energy predictor iii competition. *Scientific Data* 7(1), 1–13.
- Misyris GS, Venzke A and Chatzivasileiadis S** (2020) Physics-informed neural networks for power systems. In *2020 IEEE Power Energy Society General Meeting (PESGM)*, pp. 1–5. <https://doi.org/10.1109/PESGM41954.2020.9282004>.
- Mnih V, Kavukcuoglu K, Silver D, Graves A, Antonoglou I, Wierstra D and Riedmiller M** (2013) Playing Atari with deep reinforcement learning. Preprint, [arXiv:1312.5602](https://arxiv.org/abs/1312.5602).
- Nauata N, Chang K-H, Cheng C-Y, Mori G and Furukawa Y** (2020) House-gan: Relational generative adversarial networks for graph-constrained house layout generation. In *European Conference on Computer Vision*. Springer, pp. 162–177.
- Nghiem TX, Drgoňa J, Jones C, Nagy Z, Schwan R, Dey B, Chakrabarty A, Cairano SD, Paulson JA, Carron A, Zeilinger MN, Cortez WS and Vrabie DL** (2023) Physics-informed machine learning for modeling and control of dynamical systems. In *American Control Conference (ACC)*, pp. 3735–3750. <https://doi.org/10.23919/ACC55779.2023.10155901>.
- Nghiem TX and Jones CN** (2017) Data-driven demand response modeling and control of buildings with Gaussian processes. In *American Control Conference (ACC)*. IEEE, pp. 2919–2924.
- Oldewurtel F, Parisio A, Jones CN, Gyalistras D, Gwerder M, Stauch V, Lehmann B and Morari M** (2012a) Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings* 45, 15–27.
- Oldewurtel F, Parisio A, Jones CN, Gyalistras D, Gwerder M, Stauch V, Lehmann B and Morari M** (2012b) Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings* 45, 15–27.
- Omeragic E, Orhan O, Uzunovic T and Golubovic E** (2023) Analysing transfer learning efficacy with different feature sets for occupancy detection. In *XXIX International Conference on Information, Communication and Automation Technologies (ICAT)*. IEEE, pp. 1–6.
- Orland B, Ram N, Lang D, Houser K, Kling N and Coccia M** (2014) Saving energy in an office environment: A serious game intervention. *Energy and Buildings* 74, 43–52.
- Østergård T, Jensen RL and Maagaard SE** (2018) A comparison of six metamodeling techniques applied to building performance simulations. *Applied Energy* 211, 89–103.
- Pai P-F and Lin C-S** (2005) A hybrid Arima and support vector machines model in stock price forecasting. *Omega* 33(6), 497–505.
- Pan SJ, Zheng VW, Yang Q and Hu DH** (2008) Transfer learning for wifi-based indoor localization. In *Association for the Advancement of Artificial Intelligence (AAAI) Workshop*, Vol. 6.
- Panjaitan SD and Hartoyo A** (2011) A lighting control system in buildings based on fuzzy logic. *TELKOMNIKA (Telecommunication Computing Electronics and Control)* 9(3), 423–432.
- Papamichael K, LaPorta J and Chauvet H** (1997) Building design advisor: Automated integration of multiple simulation tools. *Automation in Construction* 6(4), 341–352. [https://doi.org/10.1016/S0926-5805\(97\)00043-5](https://doi.org/10.1016/S0926-5805(97)00043-5); Available at <https://www.sciencedirect.com/science/article/pii/S0926580597000435>.
- Parisio A, Fabietti L, Molinari M, Varagnolo D and Johansson KH** (2014) Control of HVAC systems via scenario-based explicit MPC. In *53rd IEEE Conference on Decision and Control*. IEEE, pp. 5201–5207.
- Passenberg C, Meyer D, Feldmaier J and Shen H** (2016) Optimal water heater control in smart home environments. In *IEEE International Energy Conference (ENERGYCON)*. IEEE, pp. 1–6.
- Petroşanu D-M, Caruţaşu G, Caruţaşu NL and Pirjan A** (2019) A review of the recent developments in integrating machine learning models with sensor devices in the smart buildings sector with a view to attaining enhanced sensing, energy efficiency, and optimal building management. *Energies* 12(24), 4745.
- Pigliautile I, Casaccia S, Morresi N, Arnesano M, Pisello AL and Revel GM** (2020) Assessing occupants' personal attributes in relation to human perception of environmental comfort: Measurement procedure and data analysis. *Building and Environment* 177, 106901.
- Pinto G, Wang Z, Roy A, Hong T and Capozzoli A** (2022) Transfer learning for smart buildings: A critical review of algorithms, applications, and future perspectives. *Advances in Applied Energy* 5, 100084.
- Psomas T, Fiorentini M, Kokogiannakis G and Heiselberg P** (2017) Ventilative cooling through automated window opening control systems to address thermal discomfort risk during the summer period: Framework, simulation and parametric analysis. *Energy and Buildings* 153, 18–30.
- Qi X, Chen G, Li Y, Cheng X and Li C** (2019) Applying neural-network-based machine learning to additive manufacturing: Current applications, challenges, and future perspectives. *Engineering* 5(4), 721–729. <https://doi.org/10.1016/j.eng.2019.04.012>; Available at <https://www.sciencedirect.com/science/article/pii/S2095809918307732>.
- Quashie M, Marnay C, Bouffard F and Joós G** (2018) Optimal planning of microgrid power and operating reserve capacity. *Applied Energy* 210, 1229–1236.

- Quintana M, Schiavon S, Tham KW and Miller C** (2020) Balancing thermal comfort datasets: We can, but should we?. CoRR abs/2009.13154. Preprint, arXiv:2009.13154.
- Radhakrishnan N, Srinivasan S, Su R and Poolla K** (2017) Learning-based hierarchical distributed HVAC scheduling with operational constraints. *IEEE Transactions on Control Systems Technology* 26(5), 1892–1900.
- Radhakrishnan N, Su Y, Su R and Poolla K** (2016) Token based scheduling for energy management in building HVAC systems. *Applied Energy* 173, 67–79.
- Rahman H and Han H** (2017) Occupancy estimation based on indoor CO<sub>2</sub> concentration: Comparison of neural network and Bayesian methods. *International Journal of Air-Conditioning and Refrigeration* 25(3), 1750021.
- Roberts A and Marsh A** (2001) Ecotect: Environmental prediction in architectural education.
- Robinson H, Pawar S, Rasheed A and San O** (2022) Physics guided neural networks for modelling of non-linear dynamics. Preprint, arXiv:2205.06858.
- Rupp RF, Vásquez NG and Lamberts R** (2015) A review of human thermal comfort in the built environment. *Energy and Buildings* 105, 178–205.
- Saad W, Han Z and Poor HV** (2011) Coalitional game theory for cooperative micro-grid distribution networks. In *IEEE International Conference on Communications Workshops (ICC)*. IEEE, pp. 1–5.
- Seyedzadeh S, Rahimian FP, Glesk I and Roper M** (2018) Machine learning for estimation of building energy consumption and performance: A review. *Visualization in Engineering* 6(1), 1–20.
- Shinde P and Amelin M** (2019) A literature review of intraday electricity markets and prices. *IEEE Milan PowerTech 2019*, 1–6.
- Sim JK, Yoon S and Cho Y-H** (2018) Wearable sweat rate sensors for human thermal comfort monitoring. *Scientific Reports* 8(1), 1–11.
- Široký J, Oldewurtel F, Cigler J and Prívvara S** (2011) Experimental analysis of model predictive control for an energy efficient building heating system. *Applied Energy* 88(9), 3079–3087.
- Somu N, Sriram A, Kowli A and Ramamritham K** (2021) A hybrid deep transfer learning strategy for thermal comfort prediction in buildings. *Building and Environment* 204, 108133.
- Spangher L, Gokul A, Khattar M, Palakapilly J, Tawade A, Bouyamourn A, Devonport A and Spanos C** (2020) Prospective experiment for reinforcement learning on demand response in a social game framework. In *Proceedings of the Eleventh ACM International Conference on Future Energy Systems*, pp. 438–444.
- Spangher L, Tawade A, Devonport A and Spanos C** (2019) Engineering vs. ambient type visualizations: Quantifying effects of different data visualizations on energy consumption. In *Proceedings of the 1st ACM International Workshop on Urban Building Energy Sensing, Controls, Big Data Analysis, and Visualization, UrbSys '19*. New York, NY: Association for Computing Machinery, pp. 14–22. <https://doi.org/10.1145/3363459.3363527>.
- Standard 55** (2020) *Thermal Environmental Conditions for Human Occupancy*. Atlanta, GA: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.
- Starke M, Munk J, Zandi H, Kuruganti T, Buckberry H, Hall J and Leverette J** (2020) Real-time MPC for residential building water heater systems to support the electric grid. In *IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*. IEEE, pp. 1–5.
- Strachan P, Kokogiannakis G and Macdonald I** (2008) History and development of validation with the ESP-r simulation program. *Building and Environment* 43(4), 601–609.
- Sun B, Luh PB, Jia Q-S, Jiang Z, Wang F and Song C** (2012) Building energy management: Integrated control of active and passive heating, cooling, lighting, shading, and ventilation systems. *IEEE Transactions on Automation Science and Engineering* 10(3), 588–602.
- Sun B, Luh PB, Jia Q-S and Yan B** (2015) Event-based optimization within the Lagrangian relaxation framework for energy savings in HVAC systems. *IEEE Transactions on Automation Science and Engineering* 12(4), 1396–1406.
- Sutton RS and Barto AG** (2018) *Reinforcement Learning: An Introduction*. MIT Press.
- Tesauro G** (1994) Td-Gammon, a self-teaching backgammon program, achieves master-level play. *Neural Computation* 6(2), 215–219.
- Turker H and Bacha S** (2018) Optimal minimization of plug-in electric vehicle charging cost with vehicle-to-home and vehicle-to-grid concepts. *IEEE Transactions on Vehicular Technology* 67(11), 10281–10292.
- U. o. W.-M** (1975) *Solar Energy Laboratory, TRNSYS, A Transient Simulation Program*. Madison, WI: The Laboratory.
- U.S. Department of Energy—Energy Information Administration, Annual Energy Outlook** (2020) Table 18. Energy-Related Carbon Dioxide Emissions by Sector and Source. Available at <https://www.eia.gov/outlooks/aeo>.
- ul Haq MA, Hassan MY, Abdullah H, Rahman HA, Abdullah MP, Hussin F and Said DM** (2014) A review on lighting control technologies in commercial buildings, their performance and affecting factors. *Renewable and Sustainable Energy Reviews* 33, 268–279.
- Umlauf J, Lederer A and Hirche S** (2017) Learning stable Gaussian process state space models. In *American Control Conference (ACC)*. IEEE, pp. 1499–1504.
- Wagiman KR, Abdullah MN, Hassan MY and Radzi NHM** (2019) A review on sensing-based strategies of interior lighting control system and their performance in commercial buildings. *Indonesian Journal of Electrical Engineering and Computer Science* 16(1), 208–215.
- Wang Z and Hong T** (2020) Reinforcement learning for building controls: The opportunities and challenges. *Applied Energy* 269, 115036.

- Wang L, Theodorou EA and Egerstedt M** (2018) Safe learning of quadrotor dynamics using barrier certificates. In *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 2460–2465.
- Wanjiru EM, Sichilalu SM and Xia X** (2017) Optimal control of heat pump water heater-instantaneous shower using integrated renewable-grid energy systems. *Applied Energy* 201, 332–342.
- Watson DS, Kiliccote S, Motegi N and Piette MA** (2006) Strategies for demand response in commercial buildings.
- Wei S, Tien PW, Chow TW, Wu Y and Calautit JK** (2022) Deep learning and computer vision based occupancy CO2 level prediction for demand-controlled ventilation (DCV). *Journal of Building Engineering* 56, 104715.
- Wei T, Wang Y and Zhu Q** (2017) Deep reinforcement learning for building HVAC control. In *Proceedings of the 54th Annual Design Automation Conference*, pp. 1–6.
- Wetter M, Zuo W, Nouidui TS and Pang X** (2014) Modelica buildings library. *Journal of Building Performance Simulation* 7(4), 253–270. <https://doi.org/10.1080/19401493.2013.765506>.
- Wu D, Ma X, Huang S, Fu T and Balducci P** (2020) Stochastic optimal sizing of distributed energy resources for a cost-effective and resilient microgrid. *Energy* 198, 117284.
- Xu Z, Hu G, Spanos CJ and Schiavon S** (2017) PMV-based event-triggered mechanism for building energy management under uncertainties. *Energy and Buildings* 152, 73–85.
- Xu L, Wang S and Tang R** (2019) Probabilistic load forecasting for buildings considering weather forecasting uncertainty and uncertain peak load. *Applied Energy* 237, 180–195.
- Yan D, Xia J, Tang W, Song F, Zhang X and Jiang Y** (2008) Dest—An integrated building simulation toolkit part I: Fundamentals. In *Building Simulation*, Vol. 1. Springer, pp. 95–110.
- Yang Y, Hu G and Spanos CJ** (2020) HVAC energy cost optimization for a multizone building via a decentralized approach. *IEEE Transactions on Automation Science and Engineering* 17(4), 1950–1960.
- Yang Y, Hu G and Spanos CJ** (2021) Stochastic optimal control of HVAC system for energy-efficient buildings. *IEEE Transactions on Control Systems Technology* 30(1), 376–383.
- Yang Y, Srinivasan S, Hu G and Spanos CJ** (2021) Distributed control of multizone HVAC systems considering indoor air quality. *IEEE Transactions on Control Systems Technology* 29(6), 2586–2597.
- Yang T, Zhao L, Li W and Zomaya AY** (2020) Reinforcement learning in sustainable energy and electric systems: A survey. *Annual Reviews in Control* 49, 145–163.
- Yoon N, Piette MA, Han JM, Wu W and Malkawi A** (2020) Optimization of window positions for wind-driven natural ventilation performance. *Energies* 13(10), 2464.
- Yoshikawa H, Uchiyama A and Higashino T** (2021) Data balancing for thermal comfort datasets using conditional Wasserstein gan with a weighted loss function. In *Proceedings of the 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pp. 264–267.
- Yu L, Qin S, Zhang M, Shen C, Jiang T and Guan X** (2021) A review of deep reinforcement learning for smart building energy management. *IEEE Internet of Things Journal* 12046–12063.
- Yu L, Sun Y, Xu Z, Shen C, Yue D, Jiang T and Guan X** (2020) Multi-agent deep reinforcement learning for HVAC control in commercial buildings. *IEEE Transactions on Smart Grid* 12(1), 407–419.
- Yu L, Xie W, Xie D, Zou Y, Zhang D, Sun Z, Zhang L, Zhang Y and Jiang T** (2019) Deep reinforcement learning for smart home energy management. *IEEE Internet of Things Journal* 7(4), 2751–2762.
- Zhai D and Soh YC** (2017) Balancing indoor thermal comfort and energy consumption of ACMV systems via sparse swarm algorithms in optimizations. *Energy and Buildings* 149, 1–15.
- Zhan S and Chong A** (2021) Data requirements and performance evaluation of model predictive control in buildings: A modeling perspective. *Renewable and Sustainable Energy Reviews* 142, 110835.
- Zhan S, Wichern G, Laughman C, Chong A and Chakrabarty A** (2022) Calibrating building simulation models using multi-source datasets and meta-learned Bayesian optimization. *Energy and Buildings* 270, 112278.
- Zhang Z, Chong A, Pan Y, Zhang C and Lam KP** (2019) Whole building energy model for HVAC optimal control: A practical framework based on deep reinforcement learning. *Energy and Buildings* 199, 472–490.
- Zhang X, Shi W, Yan B, Malkawi A and Li N** (2017) Decentralized and distributed temperature control via HVAC systems in energy efficient buildings. Preprint, [arXiv:1702.03308](https://arxiv.org/abs/1702.03308).
- Zhang L, Wen J, Li Y, Chen J, Ye Y, Fu Y and Livingood W** (2021) A review of machine learning in building load prediction. *Applied Energy* 285, 116452.
- Zhao Q, Lian Z and Lai D** (2021) Thermal comfort models and their developments: A review. *Energy and Built Environment* 2(1), 21–33.
- Zhao Y, Tu P and Chang M-C** (2019) Occupancy sensing and activity recognition with cameras and wireless sensors. In *Proceedings of the 2nd Workshop on Data Acquisition to Analysis*, pp. 1–6.
- Zhou X, Xu L, Zhang J, Niu B, Luo M, Zhou G and Zhang X** (2020) Data-driven thermal comfort model via support vector machine algorithms: Insights from ASHRAE RP-884 database. *Energy and Buildings* 211, 109795.
- Zou H, Das HP, Yang J, Zhou Y and Spanos C** (2019) Machine learning empowered occupancy sensing for smart buildings.
- Zou H, Jiang H, Luo Y, Zhu J, Lu X and Xie L** (2016) Bluedetect: An iBeacon-enabled scheme for accurate and energy-efficient indoor-outdoor detection and seamless location-based service. *Sensors* 16(2), 268.

- Zou H, Yang J, Das HP, Liu H, Zhou Y and Spanos CJ** (2019) WiFi and vision multimodal learning for accurate and robust device-free human activity recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*.
- Zou H, Zhou Y, Jiang H, Chien S-C, Xie L and Spanos CJ** (2018a) Winlight: A wifi-based occupancy-driven lighting control system for smart building. *Energy and Buildings* 158, 924–938.
- Zou H, Zhou Y, Yang J, Gu W, Xie L and Spanos C** (2017) Freedetector: Device-free occupancy detection with commodity WiFi. In *2017 IEEE International Conference on Sensing, Communication and Networking (SECON Workshops)*. IEEE, pp. 1–5.
- Zou H, Zhou Y, Yang J, Liu H, Das HP and Spanos CJ** (2019) Consensus adversarial domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33, pp. 5997–6004.
- Zou H, Zhou Y, Yang J and Spanos CJ** (2018b) Unsupervised WiFi-enabled IoT device-user association for personalized location-based service. *IEEE Internet of Things Journal* 6(1), 1238–1245.
- Zou H, Zhou Y, Yang J and Spanos CJ** (2018c) Device-free occupancy detection and crowd counting in smart buildings with WiFi-enabled IoT. *Energy and Buildings* 174, 309–322.
- Zuraimi M, Pantazaras A, Chaturvedi K, Yang J, Tham K and Lee S** (2017) Predicting occupancy counts using physical and statistical CO<sub>2</sub>-based modeling methodologies. *Building and Environment* 123, 517–528.