

Machine Learning for Smart and Energy-Efficient Buildings

Hari Prasanna Das[†], Yu-Wen Lin^{*†}, Utkarsha Agwan^{*†}, Lucas Spangher^{*†}, Alex Devonport^{*†}, Yu Yang^{*‡},
Ján Dragoňa[§], Adrian Chong^ξ, Stefano Schiavon[⌘], Costas J. Spanos[†]

[†]Department of Electrical Engineering and Computer Sciences, University of California, Berkeley

[‡]School of Automation Science and Engineering, Xi'an Jiaotong University

[§]Pacific Northwest National Laboratory

^ξDepartment of Building, National University of Singapore

[⌘]Center for Built Environment, University of California, Berkeley

^{*}Authors Contributed Equally

Corresponding Author: Hari Prasanna Das (hpdas@berkeley.edu)

Abstract

Energy consumption in buildings, both residential and commercial, accounts for approximately 40% of all energy usage in the U.S., 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 the energy consumption in buildings must be optimized, all the while maintaining satisfactory levels of occupant comfort, health, and safety. Recently, Machine Learning has been proven to be an invaluable tool in deriving important insights from data and optimizing various systems. In this work, we review the ways in which machine learning has been leveraged to make buildings smart and energy-efficient. For the convenience of readers, we provide a brief introduction of several machine learning paradigms and the components and functioning of each smart building system we cover. Finally, we discuss challenges faced while implementing machine learning algorithms in smart buildings and provide future avenues for research at the intersection of smart buildings and machine learning.

Keywords: Smart Buildings, Machine Learning, Energy Efficiency

1. Introduction

Buildings and their indoor environment substantially influence our health, well-being, safety, and work and study performance. We spend 90% of our time every day in indoor environments. 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 [1]. 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 being able to maintain occupant comfort and

productivity in buildings can be an environmentally sustainable, equitable, cost-effective, and scalable approach to reducing greenhouse gas emissions.

The smart building ecosystem is illustrated in Fig 1. Occupants constitute the basic building block of the ecosystem. Being the consumer of the facilities that a building provides, occupants necessitate regulation of the building systems to achieve the desired environment. The building comprises of devices and system in place to control and maintain the desired environment for the occupants, along with diagnostic systems to ensure a 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 with the energy provider and with each other.

In order to improve energy efficiency and occu-

November 29, 2022

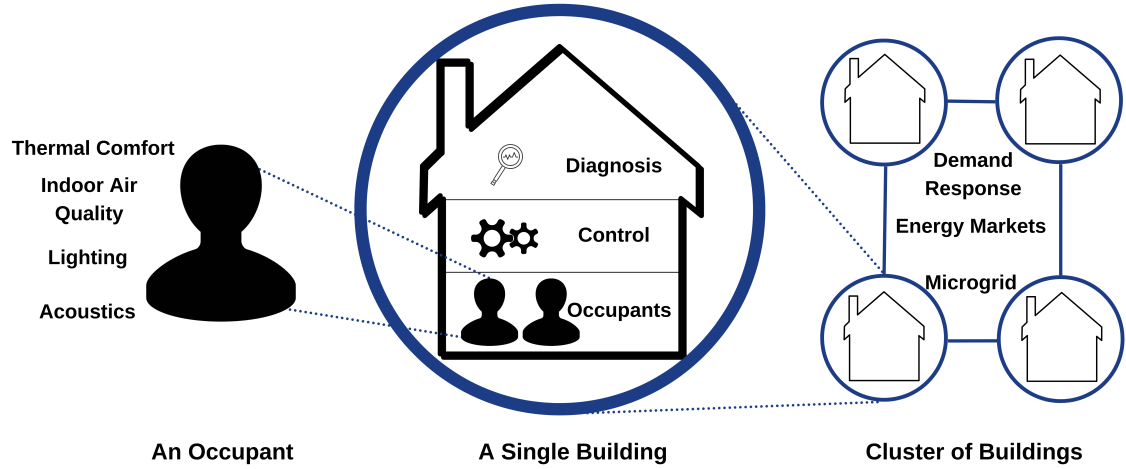


Figure 1: A taxonomy of Smart Buildings ecosystem with building operations illustrated at three levels, at a cluster of buildings level, at a single building level, and at a single occupant level.

partment 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 internet-of-things (IoT) devices in buildings, the heterogeneity of user-to-device and device-to-device interactions in buildings has become considerable in variety and thus necessitates a system that can adequately account for it. Deriving insights from the vast amount of data in certain scenarios, and from the limited amount of data in other scenarios have also been crucial. Fueled by the above challenges, Machine Learning (ML) as a tool has been implemented in smart buildings, with increasingly important implications. ML constitutes algorithms to process and extract meaningful insights from data and to use it for downstream tasks such as forecasting, prediction, and control. In this work, we explore the status quo of ML applications in smart building systems. In doing so, we touch upon traditional methods already in place, and on how ML-based solutions often match or outperform them.

2. Machine Learning

In this section, we briefly cover the fundamentals of various machine learning algorithms and methods that are commonly in practice in smart buildings. Machine Learning algorithms can be commonly grouped under supervised, unsupervised, and reinforcement learning depending on the way a model gets trained on the data-available. We

briefly cover the functionalities of the above learning paradigms and various learning methods.

2.1. Learning Paradigms

2.1.1. Supervised Machine Learning

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 work through some optimization routine to minimize a loss or error function. The model is trained to function 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 etc. Since supervised learning required annotated labels as an input, it faces a bottleneck when labels are expensive or require significant labor to obtain.

2.1.2. Unsupervised Machine Learning

Unsupervised learning involves models/algorithms that look for patterns in a dataset with no labels and with minimal human supervision. This is in contrast to supervised learning techniques, such as classification or regression, where a model is given a training set of inputs and a set of observations/labels, and must learn a mapping from

the inputs to the observations. Some unsupervised learning paradigms include clustering techniques, that aim 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 and generate synthetic data, etc.

2.1.3. Reinforcement Learning

RL is a type of agent-based machine learning where a complex system is controlled through actions that optimize the system in some manner [2]. The actions (a_t) taken based on a probability distribution p_π at states (s_t) at time t seek to optimize the expected sum of rewards (r) based on a policy (J) parameterized by θ ; 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, i.e., the evolution of system states is entirely 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 a Markov Decision Processes (MDP) with neural nets serving as function approximators helping to decide which state transitions to attempt. RL is useful in contexts where actions and environments are simple or data is plentiful, with early use cases being optimizing the control of backgammon [3], the cart-pole problem, and Atari [4]. Recently, much work has been done to extend RL to much broader use cases.

RL architectures may be grouped in families between which learning differs significantly; 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 sub-family of policy gradient is the family of Actor-Critic architectures. Here the RL controller is made up of two function approximators; one estimates actions (i.e. the actor), as in policy gradient, but another estimates the long term value of the actions (i.e. the critic); these estimates contribute to the training of the actor network. An overview of the reinforcement learning architecture taxonomy is provided in Fig 2.

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

2.2. Learning Methods

2.2.1. Kernel-based Methods

Kernel algorithms are a class of ML algorithms that are designed for analyzing patterns in data from

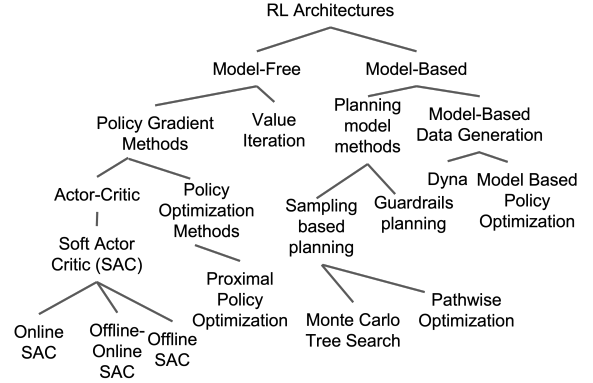


Figure 2: A taxonomy of Reinforcement Learning architectures.

explicit feature vector representations provided to them. They map the input data into a different space to make the inference of the given data easier by finding hyperplanes in that space that divide the data. Kernel methods can be used for supervised and unsupervised problems. A commonly used kernel-based method for classification is Support Vector Machines (SVM). Similarly, a common method in unsupervised settings is kernel Principal Component Analysis (PCA), that is used for spectral clustering.

2.2.2. Ensemble Methods

Ensemble methods combine multiple machine learning 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 use artificial neural networks inspired by information processing in biological neurons for applications in supervised, semi-supervised or unsupervised settings. The neural networks are typically designed based on the data they are employed to process, such as convolutional neural networks to analyse image data, recurrent networks to analyse textual, natural language or time-series data etc.

2.2.4. Time Series Forecasting Methods

Time series data comprises of a series of data points indexed by time, such as the 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, forecasting, and predictions. Assuming future data has some relation with historically observed data, models such as Moving Average (MA), Auto-Regression (AR), and combinations of these can be used to forecast future data points. Some common use cases are stock price forecasting [5] and supply chain [6]. Recurrent neural networks, and Long Short Term Memory (LSTM) networks can be used for time-series predictions.

2.2.5. *Physics-based Machine Learning*

Machine Learning (ML) models have shown promising results in learning scientific problems that are yet to be well understood or where it's computationally infeasible with physics based models. However, pure data-driven models often require large datasets, offer limited extrapolation abilities, and are unable to provide physically consistent results. On the other hand, pure physics-based models are often complex, include large amount of uncertain parameters, and may suffer from burdens of computational cost. As a result, research communities have begun to explore the combination of physics domain knowledge and ML models. The hybrid models may compensate the weaknesses of pure physics-based models or pure data-driven models, and provide potential for achieving generalizability and explainability our of the models. Coupling physics knowledge with ML has been used in several areas, including medicine [7, 8], fluid mechanics [9], power systems [10], manufacturing [11] etc. It has shown potentials 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*

Control theory has at times anticipated the need to learn about the system from data: this is collected in the disciplines of system identification [12] and adaptive control [13], which have similarities to supervised learning [14] and reinforcement learning [15] 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 prominence as reinforcement learning has.

From the reinforcement learning side, several approaches to model-based reinforcement learning use control-theoretic methods. Essentially, the problem is split into two parts: a dynamical model is learned

from data, and the model is used to design a controller, with variations on model-predictive control being a popular choice. From the control theory side, learned models are incorporated into modern control synthesis techniques, and the standard methods of proving safety are adjusted to incorporate learned models. A learning-based model that has proved to mesh particularly well with this approach are Gaussian process models, which have experienced a renaissance among learning-based controls researchers [16, 17, 18, 19, 20]. Another point of contact is the use of statistical guarantees for controllers based on learned models. In reinforcement learning, 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 provide guarantees of safety with high probability [20, 21, 22]. Since many learned models are statistical in nature, such a probabilistic bound is often the best possible without overestimating the accuracy of the model, which has also become popular in recent years [15].

3. Applications of Machine Learning in Smart Buildings

We organize the ML applications in smart buildings into different building sub-systems and other desiderata. A summary of ML applications at the level of occupants, single building and a cluster of buildings is provided in Fig. 3.

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 express satisfaction with the thermal environment and is assessed by subjective evaluation” [23]. Humans spend more than 90% of their days within a built environment, where their health, well-being, performance and energy consumption are linked to thermal comfort. But, studies show that only 40% of commercial building occupants are satisfied with their thermal environment [24]. Ensuring occupants’ thermal comfort involves understanding of the parameters that affect it, developing and using predictive models to predict thermal comfort, and controlling HVAC systems to achieve occupant satisfaction. Machine learning has been applied to all

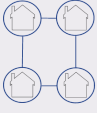


	Inter-Building 	Building 	Occupant 
Control strategies • Reinforcement learning • Hard control and model-predictive control	• Power flow coordination • Smart energy storage and distribution	• HVAC, lighting control • Space scheduling • Demand response	• Occupant-centric environmental control
Learning Models • 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
Markets • 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.

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 understanding and prediction for occupants in a building have been studied extensively in past. The prediction models can be broadly divided based on the features used to develop the model, as traditional models like Predicted Mean Vote (PMV) and the adaptive thermal comfort models and the more innovative personal comfort models.

The most widely used thermal comfort model is the PMV, developed by Fanger in 1970 [25] 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, 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 [25], that establishes a quantitative prediction of the percentage of thermally dissatisfied occupants. To com-

ply 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 non-uniformity of the space and dynamic thermal conditions [26]. 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 [27]. Experiments in real buildings showed that the PMV has a low prediction accuracy (around 33%), whereas the PPD is unreliable as a metric, hence it is not advisable to be used [27].

With the booming development of Machine Learning (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 [28, 29]. The model is later leveraged to predict the same given raw features. Because the models are trained to learn from the data itself, and not from a rule that was established using prior experiments, they perform better as compared to PMV.

Kernel-based approaches have been the widely-used methods for thermal sensation/preference prediction. The list of kernel-based methods popular for thermal comfort prediction are Support Vector Machine (SVM) [30, 31, 32, 33, 34], K-Nearest Neighbors (KNN) [35, 36, 37, 38, 28], and Ensemble Learning algorithms, such as Random Forest (RF) [28, 29] and AdaBoost (Ab). Recently, neural network such as feed-forward networks [39, 37, 40], and time-series based networks [41] have proven to be capable of surpassing the 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. They were designed to predict the average thermal comfort of the entire population rather than an individual. Personal thermal comfort models, which is a new approach to thermal comfort modeling that predicts individuals' thermal comfort responses instead of the average response of a large population, can be applied to any building thermal control system [29]. Personal thermal comfort can adapt to the available input variables, such as environmental variables [42], occupant behaviors [43] and physiological signals [28]. 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, [28] 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 Random Forest 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 [44]. Authors in [45] 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 [46], artificial neural networks [39], etc. Better approaches for

modeling tabular data in smart buildings, with a focus on thermal comfort datasets are provided in [47].

Although several studies show high prediction accuracy in predicting 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 [48, 28] 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 weighed loss function for ML models. A related challenge is the overall lack of sufficient amount of data [28]. Collecting large amounts of data as required by ML models from humans via real-world experiments is cumbersome and resource-intensive. Generating synthetic data to augment the training dataset is one of the approaches that has been proposed for tackling with the above challenges. The synthetic data generation can be done using classical methods such as SMOTE [49, 50] or using advanced neural network-based generative models [50, 51, 52, 53].

Another challenge is domain discrepancy. Thermal comfort, as per PMV, is dependent upon 6 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 [54]. Under such domain discrepancy, using models developed in one environment in another target environment may lead to low accuracy or misleading predictions. Thermal comfort modeling in the target domain independently requires a large amount of self-reported thermal comfort data. To deal with the data/label insufficiency challenge, domain adaptation methods have been proposed to adapt ML models from one domain to another. [39] propose a transfer learning framework, using Adversarial Domain Adaptation (ADA) to develop personal thermal comfort predictors for target occupants in an unsupervised manner. [55] propose a transfer learning-based multilayer perceptron model for domain adaptation across cities.

Summarizing, machine learning models have proven to increase thermal comfort predictions accu-

racy 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. There are several future research in this direction that we enlist in Section 5.

3.2. Occupancy and Activity Sensing

Occupancy and activity sensing are key aspects for 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 till 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 [56].

Occupancy sensing can be performed using both intrusive and non-intrusive methods. Intrusive methods require the occupants to carry an electronic device whose signature is followed by a central server to infer occupancy/positioning [57, 58, 59, 60, 61]. However, requiring occupants to constantly carry a device is not reliable. This problem gets magnified for the case of elderly population. Hence, non-intrusive methods for occupancy sensing are getting popular. In non-intrusive 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 non-intrusive occupancy sensing, allowing for the data-driven construction of a model to map sensor data into an estimate of occupancy.

Non-intrusive 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, rather infrared sensors or thermal cameras have been the choice for occupancy detection. For example, authors in [62] use an U-Net like convolutional neural network on thermal images to infer occupancy.




Figure 4: Evolution of building modeling.

Other works employing similar machine learning methods on depth cameras are [63, 64, 65]. 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 WiFi signals [66], because of its ubiquitous presence, and better privacy guarantees. Authors in [67, 68] use Channel State Information (CSI) data collected from WiFi 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 [69] by using convolutional neural networks on the CSI heatmaps to detect human gestures. Another modality that is used to detect occupancy is CO₂ data in a room. A number of works [70, 71, 72, 73] employ machine learning methods to map the CO₂ concentration and occupancy. Finally, works such as [74] propose sensor fusion, where data from multiple sensing modalities, i.e. RGB camera, and WiFi are used in tandem to come up with a robust activity detection mechanism.

To summarize, machine learning 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 [75].

3.3. Building Design and Modeling

Building Information Modeling (BIM) was introduced in the 1970s when Eastman [76] proposed the idea of migrating from hand-drawn building design towards 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, among them Building Design Advisor (BDA) [77], Ecotect [78], EnergyPlus [79], ESP-r [80], DEST [81], and Modelica [82]. They focused on thermal, energy,




Figure 5: Three types of building models

lighting and air quality aspects. They provided feedback and “what-if” analyses. With the emergence of IoT, data has become more easily accessible, which has helped support building simulation tools and allowed development of machine learning 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 [83] 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 machine learning and building modeling. It is separated into five parts: models, design automation, digital twins, applications, and challenges.

3.3.1. Models

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

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 [82], EnergyPlus [79], TRNSYS [84], and ESP-r [80]. 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.

Black box model. A black box model, also known as a data-driven or machine learning model, is a modeling approach that constructs a model directly from

data without knowledge of the system physics. With the emergence of Internet of Things (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 [85]. The common machine learning methodologies for building load prediction are ANN, SVM, Gaussian-based regressions, and clustering [86]. 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.

Grey box model. Grey box modeling is a modeling approach that simplifies the physical equations to simulate the behavior of a building and it combines them with machine learning 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 [87], 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 grey box model works well in building control applications [88, 89] 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 et al. [87] analyzed the model selections and found that the 3R3C model is the minimal best fit model for a $120m^2$ 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. [90] 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 [91, 92] explored physics-based machine learning for modeling lake temperature and other dynamical systems. However, there are limited studies in physics-based ML for building modeling.

The definition of grey box models is still ambiguous, where it can refer to a Resistor-Capacitor Network or a hybrid model that utilizes both first principle equations and data-driven methods. While grey 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 [93]. 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, grey and black box model are more commonly used in experimental studies of real buildings [94]. 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 [95], 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. [96] 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 [97] 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) [98] 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 post-occupancy phase (“as operated”). For buildings, forward building performance simulation (BPS) tools such as EnergyPlus [79], TRNSYS [84], ESP-r [80], and Modelica [82] are usually used to estimate the energy use of the building and its subsystems. Over the past decade, machine learning has been increasingly applied to BPS, fueled in part by the emergence of the Internet of Things (IoT) and the need for computationally tractable predictions.

Machine learning has been used during the design stage to augment generative design and parametric simulations. Deep generative algorithms such as Generative Adversarial Networks (GANs) [99, 100] 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 [101].

Metamodeling, defined as the practice of using a model to describe another model as an instance [102], is another aspect where machine learning has been extensively applied to BPS throughout the building lifecycle. Given the complex interaction between different building systems and sub-systems, design optimization during early design typically requires exploring high-dimensional design space. Consequently, machine learning has been used to create metamodels that can be used for optimization and uncertainty analysis [103, 104]. 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 [105]. A comparison of different metamodeling techniques can be found in [106].

Moving from design to post-occupancy, model calibration is often undertaken to improve the model's credibility by reducing the discrepancies between simulation predictions and actual observations. Metamodeling using machine learning is often used to emulate simulation predictions to alleviate the high computation costs [107]. Particularly, metamodeling is almost always applied in Bayesian

calibration that is computationally intractable because of the need to perform many model evaluations [108]. 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 [109, 110].

BPS has also been applied to enhance the use of deep reinforcement learning (DRL) based building controls. For instance, BPS can be used to develop building test cases for the comparison and benchmarking of DRL algorithms [111]. 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 [112]. Consequently, the exploration during online training can lead to undesirable building operations when the algorithm performs sub-optimally. To enhance the practical applications of DRL, BPS has been used as an emulator to pre-train a DRL agent offline. For instance, Zhang et al. [113] 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 [114]. However, pre-training 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 [115, 116]. 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 [117]. However, faults

in an actual building don’t 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 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 grey 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 Heating, Ventilation and Air Conditioning (HVAC) system, so we will focus our attention there. However, we do want to note that HVAC is by no means the *only* controls problems that has received focus; other building control systems enjoy active work and engaged communities, including lighting [118, 119, 120, 121], window control [122, 123], electric vehicle charging [124, 125, 126], and domestic water heating [127, 128, 129].

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. (1) 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. (2) 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 pre-heat 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 multi-step 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 w.r.t. 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 con-

cerns 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 works can be broadly categorized by model architecture into: (1) *Model Predictive Control* (MPC) and (2) *Reinforcement Learning* (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.

In what follows, we review the existing works in the three categories.

3.4.1. Model Predictive Control (MPC)

Model predictive control takes a large step towards 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 classical 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 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:

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

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 , (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 (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 (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 (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 (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.

Model predictive control (MPC) has been widely used to achieve anticipatory control of HVAC system for energy saving while ensuring human comfort [130]. MPC-based control methods require the solving of a multi-step 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 [130, 131, 132, 133] 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 on-line. The intrinsic problem related to the HVAC control is non-linear and non-convex. This leads to the high online computa-

tion 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 $B(t)$ matrix functions which whose values may be time-dependent but not state-dependent. This model simplification, in conjunction with a quadratic objective, reduces (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 [134, 135, 136]. The key idea is to drive some close-form solution for the MPC problem and then identify the parameterized control outputs on-line. 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 on-line 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 works have made effort to develop distributed MPC controllers [137, 138, 139, 140, 141]. 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 w.r.t. the total zone mass flow rates), and operating constraints). The optimization problem related to multi-zone HVAC system is a nonlinear and nonconvex optimization problem subject to nonlinear couplings (nonlinear and non-convex). The existing distributed solution methods can not be directly applied to solve such a class of problems. To overcome the challenges and enable the decomposition of the problem, [139, 137, 138] has explored the convex relaxation technique to solve such problem. Particularly, [137, 138] 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, [140, 141] have explored the decomposition

of the problem w.r.t. 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 a two-step. These works 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 on-line 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, reinforcement learning (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 multi-stage decision making under uncertainties [2] makes it well suited for HVAC control. Indeed, RL has been widely used for advanced control of HVAC systems considering the uncertainties and multi-stage problem features [142, 143, 144]. RL can enable efficient on-line computation by learning the control policies off-line. Specifically, by learning the control policies off-line, the on-line 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 computational 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, deep reinforcement learning (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 on-line implementation. The proliferation of smart meters have boosted the development of DRL-based HVAC

controllers for smart buildings. We refer the readers to [145, 146, 147, 148, 149]. As an earlier work, [150] has explored deep Q-network based RL for single zone HVAC control. An Q-network is trained to approximate the state-value functions to overcome the computation burden related to the large state space. [151] has explored the application of Deep Deterministic Policy Gradients (DDPG) [152] 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 multi-zone commercial buildings [153, 154]. For example, [153] studied the application of multi-agent actor-attention-critic RL [155] for multi-zone 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 off-line 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 work 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 [156]. Augmenting reinforcement learning 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 estab-

lished 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 [157]. Jang et al. studies offline learning through offline-online RL [158] and Model Agnostic Meta Learning (MAML) [159] to get RL agents “hit the ground ready” using offline datasets and simulations.

3.4.3. Learning-based Model Predictive Control

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 on-line computation burden. Moreover, MPC-based controllers generally suffer from the modeling complexity. For example, the more elaborate thermal comfort model, such as the Predicted Mean Vote (PMV) is hard to be incorporated in the MPC-based control due the challenge to solve the optimization problem [160]. 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 has 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 on-line computation burden and enable the deployment of MPC controller on simple hardware, the characterization capacity of ML, such as deep neutral 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 on-line implementation cost. For example, [161] 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 3% of performance loss but reduces the on-line computation burden by a factor of 7. Later, a real-time implementation of the approximate MPC is developed and demonstrated in [162] in an office building located in Hasselt, Belgium.

Similarly, [163] 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 on-line implementation only requires low on-line 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 refer the readers to [164] for a comprehensive review. For example, neural networks can be trained to efficiently predict PMV [165, 166] 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 on-line 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 can not incorporate the possibly available dynamic predictions to enhance the performance. That’s 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 can not 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 non-compliance weakened the overall effect of the programs [167].

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 (i.e. equally matched competitors), or social status (i.e. badges, ranks, and advancement.) Hamari et al. [168] 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. For a visualization of the purpose that gamifying provides, please see Figure 6.

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, an game called “NTU Social Game” has been formulated around reductions in energy in the 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




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

prizes are allocated amongst 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 [169, 170]. The NTU social game has been used for energy reduction and to test different visualizations on player outcome [171]. It has also been used to conduct segmentation analysis of players into low, medium and high energy-efficient groups for intelligent incentive planning [172, 173]. 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 [174]. 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 [175]. 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 [176]. 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 [177]. 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 multi-agent 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 can not 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 demand response 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 CO₂ emission. In the following parts, we will discuss demand response in further details, and how machine learning plays a role in achieving the

same. We will also cover on the prosumer nature of buildings.

3.5.1. Demand Response (DR)

We start off by discussing how individual buildings can use their load flexibility as demand response resources. Demand response 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.

DR sources and timescales. There are two primary sources of flexible loads in smart buildings: lighting and the HVAC system [178]. Modifying energy consumption for either of these has a direct impact on the building occupants, as covered in the thermal comfort section (Sec.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 [179]. For example, smart buildings can *shift* load by pre-heating or pre-cooling 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.

DR program mechanisms. The two most common types of demand response 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 [179]. 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 [180] for an overview of common demand response mechanisms.

ML for forecasting loads and baselines. Buildings procure energy from utilities, who buy it from from generators through long-term contracts or in wholesale energy markets. Energy consumption is variable, and utilities employ many machine learning methods to forecast electricity consumption accurately. Authors in [181] 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. [182] 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, support vector machine (SVM) regression for short-term and long-term forecasts, as well as neural network based methods. [183] 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 machine learning 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. [184] 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. [185] use a variety of time series methods to forecast the peak load demand of institutional buildings, with a dataset of buildings in South Korea.

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 vari-

able, 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. [186] 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. [187] use linear regression to develop models for estimating demand response capability using time-of-week and outdoor air temperature buckets. It is important to have such DR capability estimates over longer time horizons, and [188] 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, 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 Demand Response 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. Demand response can also be coupled with other resources like generation and storage, and form a part of a *virtual power plant*.

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. [189] 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 [190]. Optimizing the complementary nature of buildings can significantly increase the marginal value of forming an aggregation.

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. [191] conduct a review of electricity price forecasting methods. [192] 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.

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 [193].

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. Machine learning 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. The authors in [194] use coalitional game theory to optimally aggregate prosumers in a cooperative microgrid. [195] uses a mixed-integer linear programming to discuss the optimal sizing of energy resources in a microgrid. In [196], 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 [197] have developed a complementarity metric that uses the 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 that prices. A number of previous works develop price based controls: [198] develop a

price to minimize load variation, and [199] 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 literature, e.g. [200] 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. [201] approaches this problem using reinforcement learning, and develop a controller that can set day ahead electricity prices for a local market with a variety of prosumers.

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

4. Limitations of Machine Learning in the context of Smart Buildings

Although machine learning techniques have proliferated across many fields to great success, the lack of success of machine learning in some applications bears note. There are some notable applications where ML doesn't work well: 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 machine learning that may help to address 99/01 problems in buildings are safety-guarantees in machine learning, robust worst case minimizing reinforcement learning, proper explainability etc.

4.2. Data Problems

Does the domain provide rich, reliable, and accurate data? The shining examples of machine learning success often occur in domains where this is the case: reinforcement learning 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. Oftentimes, building energy datasets are not as large, rich, or diverse, and so many of the techniques that work well in different domains may not translate well to energy [48, 202, 203, 204]. A fruitful area of work may include preparing and standardizing large datasets for ease of energy research. 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 machine learning models? Oftentimes we may put a lot of work into making a prediction model accurate, but a building from the 70's 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, machine learning research in controls is inherently closer to many of the applications that a building may need.

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 is even harder to come by as it uses intrusive methods that require occupants’ input. Additionally, thermal comfort data is 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 is 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 [205], which is a metadata schema that represents buildings’ sensor information and the relationships between them. Grey-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 [53].

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 dif-

ferent 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 becomes 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.

6. Acknowledgements

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.

7. Declaration of Competing Interest

The authors declare no sources of competing interest(s).

References

- [1] U.S. Department of Energy—Energy Information Administration, Annual Energy Outlook 2020. Table 18. Energy-Related Carbon Dioxide Emissions by Sector and Source, <https://www.eia.gov/outlooks/aeo> (2020).
- [2] R. S. Sutton, A. G. Barto, Reinforcement learning: An introduction, MIT press, 2018.
- [3] G. Tesauro, Td-gammon, a self-teaching backgammon program, achieves master-level play, *Neural computation* 6 (2) (1994) 215–219.
- [4] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, M. Riedmiller, Playing atari with deep reinforcement learning, *arXiv preprint arXiv:1312.5602* (2013).

- [5] P.-F. Pai, C.-S. Lin, A hybrid arima and support vector machines model in stock price forecasting, *Omega* 33 (6) (2005) 497–505.
- [6] Y. Aviv, A time-series framework for supply-chain inventory management, *Operations Research* 51 (2) (2003) 210–227.
- [7] F. Sahli Costabal, Y. Yang, P. Perdikaris, D. E. Hurtado, E. Kuhl, Physics-informed neural networks for cardiac activation mapping, *Frontiers in Physics* 8 (2020). doi:10.3389/fphy.2020.00042. URL <https://www.frontiersin.org/article/10.3389/fphy.2020.00042>
- [8] G. Kissas, Y. Yang, E. Hwuang, W. R. Witschey, J. A. Detre, P. Perdikaris, 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 (2020) 112623. doi:https://doi.org/10.1016/j.cma.2019.112623. URL <https://www.sciencedirect.com/science/article/pii/S0045782519305055>
- [9] S. Cai, Z. Mao, Z. Wang, M. Yin, G. E. Karniadakis, Physics-informed neural networks (pinns) for fluid mechanics: A review, *Acta Mechanica Sinica* (2022) 1–12.
- [10] G. S. Misyris, A. Venzke, S. Chatzivasilieiadis, Physics-informed neural networks for power systems, in: 2020 IEEE Power Energy Society General Meeting (PESGM), 2020, pp. 1–5. doi:10.1109/PESGM41954.2020.9282004.
- [11] X. Qi, G. Chen, Y. Li, X. Cheng, C. Li, Applying neural-network-based machine learning to additive manufacturing: Current applications, challenges, and future perspectives, *Engineering* 5 (4) (2019) 721–729. doi:https://doi.org/10.1016/j.eng.2019.04.012. URL <https://www.sciencedirect.com/science/article/pii/S2095809918307732>
- [12] L. Ljung, System identification, in: *Signal analysis and prediction*, Springer, 1998, pp. 163–173.
- [13] P. Ioannou, B. Fidan, Adaptive control tutorial, SIAM, 2006.
- [14] L. Ljung, C. Andersson, K. Tiels, T. B. Schön, Deep learning and system identification, *IFAC-PapersOnLine* 53 (2) (2020) 1175–1181.
- [15] N. Matni, A. Proutiere, A. Rantzer, S. Tu, From self-tuning regulators to reinforcement learning and back again, in: 2019 IEEE 58th Conference on Decision and Control (CDC), IEEE, 2019, pp. 3724–3740.
- [16] J. Umlauft, A. Lederer, S. Hirche, Learning stable Gaussian process state space models, in: 2017 American Control Conference (ACC), IEEE, 2017, pp. 1499–1504.
- [17] L. Wang, E. A. Theodorou, M. Egerstedt, Safe learning of quadrotor dynamics using barrier certificates, in: 2018 IEEE International Conference on Robotics and Automation (ICRA), IEEE, 2018, pp. 2460–2465.
- [18] A. K. Akametalu, J. F. Fisac, J. H. Gillula, S. Kaynama, M. N. Zeilinger, C. J. Tomlin, Reachability-based safe learning with Gaussian processes, in: 53rd IEEE Conference on Decision and Control, IEEE, 2014, pp. 1424–1431.
- [19] A. Devonport, H. Yin, M. Arcak, Bayesian safe learning and control with sum-of-squares analysis and polynomial kernels, in: 2020 59th IEEE Conference on Decision and Control (CDC), IEEE, 2020, pp. 3159–3165.
- [20] A. Devonport, F. Yang, L. E. Ghaoui, M. Arcak, Data-driven reachability analysis and support set estimation with christoffel functions, arXiv preprint arXiv:2112.09995 (2021).
- [21] A. Devonport, M. Arcak, Estimating reachable sets with scenario optimization, in: *Learning for dynamics and control*, PMLR, 2020, pp. 75–84.
- [22] A. Devonport, F. Yang, L. El Ghaoui, M. Arcak, Data-driven reachability analysis with christoffel functions, in: 2021 60th IEEE Conference on Decision and Control (CDC), 2021, pp. 5067–5072. doi:10.1109/CDC45484.2021.9682860.
- [23] A. Standard, Standard 55-2020. thermal environmental conditions for human occupancy, American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.: Atlanta, GA, USA (2020).
- [24] L. T. Graham, T. Parkinson, S. Schiavon, Lessons learned from 20 years of CBE’s occupant surveys, *Buildings and Cities* 2 (1) (2021).
- [25] P. O. Fanger, Thermal comfort. analysis and applications in environmental engineering., *Thermal comfort. Analysis and applications in environmental engineering*. (1970).
- [26] Q. Zhao, Z. Lian, D. Lai, Thermal comfort models and their developments: A review, *Energy and Built Environment* 2 (1) (2021) 21–33.
- [27] T. Cheung, S. Schiavon, T. Parkinson, P. Li, G. Brager, Analysis of the accuracy on pmv–ppd model using the ASHRAE global thermal comfort database ii, *Building and Environment* 153 (2019) 205–217.
- [28] S. Liu, S. Schiavon, H. P. Das, M. Jin, C. J. Spanos, Personal thermal comfort models with wearable sensors, *Building and Environment* 162 (2019) 106281.
- [29] J. Kim, S. Schiavon, G. Brager, Personal comfort models—a new paradigm in thermal comfort for occupant-centric environmental control, *Building and Environment* 132 (2018) 114–124.
- [30] X. Zhou, L. Xu, J. Zhang, B. Niu, M. Luo, G. Zhou, X. Zhang, Data-driven thermal comfort model via support vector machine algorithms: Insights from ASHRAE rp-884 database, *Energy and Buildings* 211 (2020) 109795.
- [31] Q. Chai, H. Wang, Y. Zhai, L. Yang, Using machine learning algorithms to predict occupants’ thermal comfort in naturally ventilated residential buildings, *Energy and Buildings* 217 (2020) 109937.
- [32] F. Alsaleem, M. K. Tesfay, M. Rafaie, K. Sinkar, D. Besarla, P. Arunasalam, An iot framework for modeling and controlling thermal comfort in buildings, *Frontiers in Built Environment* 6 (2020) 87.
- [33] T. Chaudhuri, D. Zhai, Y. C. Soh, H. Li, L. Xie, Thermal comfort prediction using normalized skin temperature in a uniform built environment, *Energy and Buildings* 159 (2018) 426–440.
- [34] K. Liu, T. Nie, W. Liu, Y. Liu, D. Lai, A machine learning approach to predict outdoor thermal comfort using local skin temperatures, *Sustainable Cities and Society* 59 (2020) 102216.
- [35] I. Pigliautile, S. Casaccia, N. Morresi, M. Arnesano, A. L. Pisello, G. M. Revel, Assessing occupants’ personal attributes in relation to human perception of environmental comfort: Measurement procedure and data analysis, *Building and Environment* 177 (2020) 106901.
- [36] T. Cheung, L. T. Graham, S. Schiavon, Impacts of life satisfaction, job satisfaction and the big five personality traits on satisfaction with the indoor environment,

- Building and Environment 212 (2022) 108783.
- [37] S. Lu, W. Wang, C. Lin, E. C. Hameen, Data-driven simulation of a thermal comfort-based temperature set-point control with ASHRAE rp884, *Building and Environment* 156 (2019) 137–146.
 - [38] J. Lee, Y. Ham, Physiological sensing-driven personal thermal comfort modelling in consideration of human activity variations, *Building Research & Information* 49 (5) (2021) 512–524.
 - [39] H. P. Das, S. Schiavon, C. J. Spanos, 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*, 2021, pp. 230–231.
 - [40] D. Zhai, Y. C. Soh, Balancing indoor thermal comfort and energy consumption of acmv systems via sparse swarm algorithms in optimizations, *Energy and Buildings* 149 (2017) 1–15.
 - [41] A. Chennapragada, D. Periyakoil, H. P. Das, C. J. Spanos, Time series-based deep learning model for personal thermal comfort prediction, in: *Proceedings of the Thirteenth ACM International Conference on Future Energy Systems*, 2022, pp. 552–555.
 - [42] T. C. Cheung, S. Schiavon, E. T. Gall, M. Jin, W. W. Nazaroff, Longitudinal assessment of thermal and perceived air quality acceptability in relation to temperature, humidity, and co2 exposure in singapore, *Building and Environment* 115 (2017) 80–90.
 - [43] J. Kim, Y. Zhou, S. Schiavon, P. Raftery, G. Brager, Personal comfort models: Predicting individuals’ thermal preference using occupant heating and cooling behavior and machine learning, *Building and Environment* 129 (2018) 96–106.
 - [44] E. Laftchiev, D. Nikovski, An iot system to estimate personal thermal comfort, in: *2016 IEEE 3rd World Forum on Internet of Things (WF-IoT)*, IEEE, 2016, pp. 672–677.
 - [45] J. K. Sim, S. Yoon, Y.-H. Cho, Wearable sweat rate sensors for human thermal comfort monitoring, *Scientific reports* 8 (1) (2018) 1–11.
 - [46] N. Somu, A. Sriram, A. Kowli, K. Ramamritham, A hybrid deep transfer learning strategy for thermal comfort prediction in buildings, *Building and Environment* 204 (2021) 108133.
 - [47] H. P. Das, C. Spanos, 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)*, 2022.
 - [48] V. F. Ličina, T. Cheung, H. Zhang, R. De Dear, T. Parkinson, E. Arens, C. Chun, S. Schiavon, M. Luo, G. Brager, et al., Development of the ASHRAE global thermal comfort database ii, *Building and Environment* 142 (2018) 502–512.
 - [49] N. V. Chawla, K. W. Bowyer, L. O. Hall, W. P. Kegelmeyer, Smote: synthetic minority over-sampling technique, *Journal of artificial intelligence research* 16 (2002) 321–357.
 - [50] M. Quintana, S. Schiavon, K. W. Tham, C. Miller, Balancing thermal comfort datasets: We gan, but should we?, *CoRR abs/2009.13154* (2020). [arXiv:2009.13154](https://arxiv.org/abs/2009.13154). URL <https://arxiv.org/abs/2009.13154>
 - [51] H. Yoshikawa, A. Uchiyama, T. Higashino, 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*, 2021, pp. 264–267.
 - [52] H. P. Das, R. Tran, J. Singh, X. Yue, G. Tison, A. Sangiovanni-Vincentelli, C. J. Spanos, Conditional synthetic data generation for robust machine learning applications with limited pandemic data, *arXiv preprint arXiv:2109.06486* (2021).
 - [53] H. P. Das, C. J. Spanos, Conditional synthetic data generation for personal thermal comfort models, *arXiv preprint arXiv:2203.05242* (2022).
 - [54] L. A. Martins, V. Soebarto, T. Williamson, A systematic review of personal thermal comfort models, *Building and Environment* 207 (2022) 108502.
 - [55] N. Gao, W. Shao, M. S. Rahaman, J. Zhai, K. David, F. D. Salim, Transfer learning for thermal comfort prediction in multiple cities, *Building and Environment* 195 (2021) 107725.
 - [56] D.-M. Petroşanu, G. Caruţaşu, N. L. Caruţaşu, A. Pirjan, 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) (2019) 4745.
 - [57] S. Lee, Y. Chon, Y. Kim, R. Ha, H. Cha, Occupancy prediction algorithms for thermostat control systems using mobile devices, *IEEE Transactions on Smart Grid* 4 (3) (2013) 1332–1340.
 - [58] H. Zou, Y. Zhou, J. Yang, C. J. Spanos, Unsupervised wifi-enabled iot device-user association for personalized location-based service, *IEEE Internet of Things Journal* 6 (1) (2018) 1238–1245.
 - [59] H. Zou, H. Jiang, Y. Luo, J. Zhu, X. Lu, L. Xie, Bluedetect: An ibeacon-enabled scheme for accurate and energy-efficient indoor-outdoor detection and seamless location-based service, *Sensors* 16 (2) (2016) 268.
 - [60] H. Zou, Y. Zhou, H. Jiang, S.-C. Chien, L. Xie, C. J. Spanos, Winlight: A wifi-based occupancy-driven lighting control system for smart building, *Energy and Buildings* 158 (2018) 924–938.
 - [61] A. Filippoupolitis, W. Oliff, G. Loukas, Bluetooth low energy based occupancy detection for emergency management, in: *2016 15th international conference on ubiquitous computing and communications and 2016 International Symposium on Cyberspace and Security (IUCC-CSS)*, IEEE, 2016, pp. 31–38.
 - [62] M. Kraft, P. Aszkowski, D. Pieczyński, M. Fularz, Low-cost thermal camera-based counting occupancy meter facilitating energy saving in smart buildings, *Energies* 14 (15) (2021) 4542.
 - [63] Y. Zhao, P. Tu, M.-C. Chang, Occupancy sensing and activity recognition with cameras and wireless sensors, in: *Proceedings of the 2nd Workshop on Data Acquisition To Analysis*, 2019, pp. 1–6.
 - [64] G. Diraco, A. Leone, P. Siciliano, People occupancy detection and profiling with 3d depth sensors for building energy management, *Energy and Buildings* 92 (2015) 246–266.
 - [65] L. J. Brackney, A. R. Florita, A. C. Swindler, L. G. Polese, G. A. Brunemann, Design and performance of an image processing occupancy sensor, in: *Proceedings: The Second International Conference on Building Energy and Environment 2012987 Topic 10. Intelligent buildings and advanced control techniques*, Citeseer,

- 2012.
- [66] H. Zou, H. P. Das, J. Yang, Y. Zhou, C. Spanos, Machine learning empowered occupancy sensing for smart buildings (2019).
- [67] H. Zou, Y. Zhou, J. Yang, C. J. Spanos, Device-free occupancy detection and crowd counting in smart buildings with wifi-enabled iot, *Energy and Buildings* 174 (2018) 309–322.
- [68] H. Zou, Y. Zhou, J. Yang, W. Gu, L. Xie, C. Spanos, Freedetector: Device-free occupancy detection with commodity wifi, in: 2017 IEEE International Conference on Sensing, Communication and Networking (SECON Workshops), IEEE, 2017, pp. 1–5.
- [69] H. Zou, Y. Zhou, J. Yang, H. Liu, H. P. Das, C. J. Spanos, Consensus adversarial domain adaptation, in: Proceedings of the AAAI conference on artificial intelligence, Vol. 33, 2019, pp. 5997–6004.
- [70] M. Zuraimi, A. Pantazaras, K. Chaturvedi, J. Yang, K. Tham, S. Lee, Predicting occupancy counts using physical and statistical co2-based modeling methodologies, *Building and Environment* 123 (2017) 517–528.
- [71] K. Arendt, A. Johansen, B. N. Jørgensen, M. B. Kjærsgaard, C. G. Mattera, F. C. Sangogboye, J. H. Schwee, C. T. Veje, Room-level occupant counts, airflow and co2 data from an office building, in: Proceedings of the First Workshop on Data Acquisition To Analysis, 2018, pp. 13–14.
- [72] H. Rahman, H. Han, Occupancy estimation based on indoor co2 concentration: comparison of neural network and bayesian methods, *International Journal of Air-Conditioning and Refrigeration* 25 (03) (2017) 1750021.
- [73] S. Wei, P. W. Tien, T. W. Chow, Y. Wu, J. K. Calautit, Deep learning and computer vision based occupancy co2 level prediction for demand-controlled ventilation (dcv), *Journal of Building Engineering* 56 (2022) 104715.
- [74] H. Zou, J. Yang, H. Prasanna Das, H. Liu, Y. Zhou, C. J. Spanos, 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, 2019, pp. 0–0.
- [75] M. Esrafilian-Najafabadi, F. Haghighat, Occupancy-based hvac control systems in buildings: A state-of-the-art review, *Building and Environment* 197 (2021) 107810.
- [76] C. Eastman, The use of computers instead of drawings in building design, *AIA journal* 63 (3) (1975) 46–50.
- [77] K. Papamichael, J. LaPorta, H. Chauvet, Building design advisor: automated integration of multiple simulation tools, *Automation in Construction* 6 (4) (1997) 341–352, design computation: Reasoning and collaboration. doi:[https://doi.org/10.1016/S0926-5805\(97\)00043-5](https://doi.org/10.1016/S0926-5805(97)00043-5). URL <https://www.sciencedirect.com/science/article/pii/S0926580597000435>
- [78] A. Roberts, A. Marsh, Ecotect: environmental prediction in architectural education (2001).
- [79] D. B. Crawley, L. K. Lawrie, F. C. Winkelmann, W. F. Buhl, Y. J. Huang, C. O. Pedersen, R. K. Strand, R. J. Liesen, D. E. Fisher, M. J. Witte, et al., Energyplus: creating a new-generation building energy simulation program, *Energy and buildings* 33 (4) (2001) 319–331.
- [80] P. Strachan, G. Kokogiannakis, I. Macdonald, History and development of validation with the esp-r simulation program, *Building and Environment* 43 (4) (2008) 601–609.
- [81] D. Yan, J. Xia, W. Tang, F. Song, X. Zhang, Y. Jiang, Dest—an integrated building simulation toolkit part i: Fundamentals, in: *Building Simulation*, Vol. 1, Springer, 2008, pp. 95–110.
- [82] P. Fritzson, V. Engelson, Modelica—a unified object-oriented language for system modeling and simulation, in: *European Conference on Object-Oriented Programming*, Springer, 1998, pp. 67–90.
- [83] M. Grieves, J. Vickers, Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems, Springer International Publishing, Cham, 2017, pp. 85–113. doi:10.1007/978-3-319-38756-7_4. URL https://doi.org/10.1007/978-3-319-38756-7_4
- [84] U. o. W.-M. Solar Energy Laboratory, TRNSYS, a transient simulation program, Madison, Wis.: The Laboratory, 1975, 1975.
- [85] L. Zhang, J. Wen, Y. Li, J. Chen, Y. Ye, Y. Fu, W. Livingood, A review of machine learning in building load prediction, *Applied Energy* 285 (2021) 116452.
- [86] S. Seyedzadeh, F. P. Rahimian, I. Glesk, M. Roper, Machine learning for estimation of building energy consumption and performance: a review, *Visualization in Engineering* 6 (1) (2018) 1–20.
- [87] P. Bacher, H. Madsen, Identifying suitable models for the heat dynamics of buildings, *Energy and Buildings* 43 (7) (2011) 1511–1522.
- [88] F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, M. Morari, Use of model predictive control and weather forecasts for energy efficient building climate control, *Energy and Buildings* 45 (2012) 15–27.
- [89] J. Široký, F. Oldewurtel, J. Cigler, S. Prívvara, Experimental analysis of model predictive control for an energy efficient building heating system, *Applied energy* 88 (9) (2011) 3079–3087.
- [90] Y.-W. Lin, T. L. E. Tang, C. J. Spanos, Hybrid approach for digital twins in the built environment, in: Proceedings of the Twelfth ACM International Conference on Future Energy Systems, e-Energy '21, Association for Computing Machinery, New York, NY, USA, 2021, p. 450–457. doi:10.1145/3447555.3466585. URL <https://doi.org/10.1145/3447555.3466585>
- [91] A. Daw, A. Karpatne, W. Watkins, J. Read, V. Kumar, Physics-guided neural networks (pgnn): An application in lake temperature modeling, *arXiv preprint arXiv:1710.11431* (2017).
- [92] H. Robinson, S. Pawar, A. Rasheed, O. San, Physics guided neural networks for modelling of non-linear dynamics, *arXiv preprint arXiv:2205.06858* (2022).
- [93] T. Hong, D. Yan, S. D'Oca, C. fei Chen, Ten questions concerning occupant behavior in buildings: The big picture, *Building and Environment* 114 (2017) 518–530. doi:<https://doi.org/10.1016/j.buildenv.2016.12.006>. URL <https://www.sciencedirect.com/science/article/pii/S0360132316304851>
- [94] S. Zhan, A. Chong, Data requirements and performance evaluation of model predictive control in buildings: A modeling perspective, *Renewable and Sustainable Energy Reviews* 142 (2021) 110835.
- [95] N. Haala, M. Kada, An update on automatic 3d build-

- ing reconstruction, *ISPRS Journal of Photogrammetry and Remote Sensing* 65 (6) (2010) 570–580.
- [96] R. Jia, B. Jin, M. Jin, Y. Zhou, I. C. Konstantakopoulos, H. Zou, J. Kim, D. Li, W. Gu, R. Arghandeh, et al., Design automation for smart building systems, *Proceedings of the IEEE* 106 (9) (2018) 1680–1699.
- [97] M. Wetter, W. Zuo, T. S. Nouidui, X. Pang, Modelica Buildings library, *Journal of Building Performance Simulation* 7 (4) (2014) 253–270. doi:10.1080/19401493.2013.765506. URL <https://doi.org/10.1080/19401493.2013.765506>
- [98] A. AIAA, Digital twin: Definition & value - an AIAA and AIA position paper, <https://www.aiaa-aerospace.org/report/digital-twin-paper/> (2022).
- [99] W. Huang, H. Zheng, Architectural drawings recognition and generation through machine learning, in: *ACADIA*, 2018.
- [100] N. Nauata, K.-H. Chang, C.-Y. Cheng, G. Mori, Y. Furukawa, House-gan: Relational generative adversarial networks for graph-constrained house layout generation, in: *European Conference on Computer Vision*, Springer, 2020, pp. 162–177.
- [101] V. J. Gan, H. Wong, K. T. Tse, J. C. Cheng, I. M. Lo, C. M. Chan, Simulation-based evolutionary optimization for energy-efficient layout plan design of high-rise residential buildings, *Journal of cleaner production* 231 (2019) 1375–1388.
- [102] D. Allemang, J. Hendler, F. Gandon, Expert modeling in owl (2011).
- [103] B. Eisenhower, Z. O'Neill, S. Narayanan, V. A. Fonoberov, I. Mezić, A methodology for meta-model based optimization in building energy models, *Energy and Buildings* 47 (2012) 292–301.
- [104] F. Bre, N. Roman, V. D. Fachinotti, An efficient metamodel-based method to carry out multi-objective building performance optimizations, *Energy and buildings* 206 (2020) 109576.
- [105] X. Chen, H. Yang, K. Sun, Developing a meta-model for sensitivity analyses and prediction of building performance for passively designed high-rise residential buildings, *Applied energy* 194 (2017) 422–439.
- [106] T. Østergård, R. L. Jensen, S. E. Maagaard, A comparison of six metamodeling techniques applied to building performance simulations, *Applied Energy* 211 (2018) 89–103.
- [107] D. Coakley, P. Raftery, M. Keane, A review of methods to match building energy simulation models to measured data, *Renewable and sustainable energy reviews* 37 (2014) 123–141.
- [108] A. Chong, Y. Gu, H. Jia, Calibrating building energy simulation models: A review of the basics to guide future work, *Energy and Buildings* 253 (2021) 111533.
- [109] A. Chakrabarty, E. Maddalena, H. Qiao, C. Laughman, Scalable bayesian optimization for model calibration: Case study on coupled building and hvac dynamics, *Energy and Buildings* 253 (2021) 111460.
- [110] S. Zhan, G. Wichern, C. Laughman, A. Chong, A. Chakrabarty, Calibrating building simulation models using multi-source datasets and meta-learned bayesian optimization, *Energy and Buildings* (2022) 112278.
- [111] D. Blum, J. Arroyo, S. Huang, J. Drgoña, F. Jorissen, H. T. Walnum, Y. Chen, K. Benne, D. Vrabie, M. Wetter, et al., Building optimization testing framework (boptest) for simulation-based benchmarking of control strategies in buildings, *Journal of Building Performance Simulation* 14 (5) (2021) 586–610.
- [112] M. Botvinick, S. Ritter, J. X. Wang, Z. Kurth-Nelson, C. Blundell, D. Hassabis, Reinforcement learning, fast and slow, *Trends in cognitive sciences* 23 (5) (2019) 408–422.
- [113] Z. Zhang, A. Chong, Y. Pan, C. Zhang, K. P. Lam, Whole building energy model for hvac optimal control: A practical framework based on deep reinforcement learning, *Energy and Buildings* 199 (2019) 472–490.
- [114] S. Brandi, M. Fiorentini, A. Capozzoli, Comparison of online and offline deep reinforcement learning with model predictive control for thermal energy management, *Automation in Construction* 135 (2022) 104128.
- [115] J. L. Hensen, R. Lamberts, *Building performance simulation for design and operation*, Routledge, 2012.
- [116] J. A. Clarke, *Energy simulation in building design*, Routledge, 2007.
- [117] B. Dong, Z. O'Neill, Z. Li, A bim-enabled information infrastructure for building energy fault detection and diagnostics, *Automation in Construction* 44 (2014) 197–211.
- [118] M. A. ul Haq, M. Y. Hassan, H. Abdullah, H. A. Rahman, M. P. Abdullah, F. Hussin, D. M. Said, A review on lighting control technologies in commercial buildings, their performance and affecting factors, *Renewable and Sustainable Energy Reviews* 33 (2014) 268–279.
- [119] S. D. Panjaitan, A. Hartoyo, A lighting control system in buildings based on fuzzy logic, *TELKOMNIKA (Telecommunication Computing Electronics and Control)* 9 (3) (2011) 423–432.
- [120] K. R. Wagiman, M. N. Abdullah, M. Y. Hassan, N. H. M. Radzi, 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) (2019) 208–215.
- [121] E. S. Lee, S. E. Selkowitz, The design and evaluation of integrated envelope and lighting control strategies for commercial buildings (1994).
- [122] N. Yoon, M. A. Piette, J. M. Han, W. Wu, A. Malkawi, Optimization of window positions for wind-driven natural ventilation performance, *Energies* 13 (10) (2020) 2464.
- [123] T. Psomas, M. Fiorentini, G. Kokogiannakis, P. Heiselberg, 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 (2017) 18–30.
- [124] H. Turker, S. Bacha, 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) (2018) 10281–10292.
- [125] A. S. Al-Ogaili, T. J. T. Hashim, N. A. Rahmat, A. K. Ramasamy, M. B. Marsadek, M. Faisal, M. A. Hannan, Review on scheduling, clustering, and forecasting strategies for controlling electric vehicle charging: Challenges and recommendations, *Ieee Access* 7 (2019) 128353–128371.
- [126] E. Kontou, Y. Yin, Y.-E. Ge, Cost-effective and ecofriendly plug-in hybrid electric vehicle charging management, *Transportation Research Record* 2628 (1) (2017) 87–98.
- [127] E. M. Wanjiru, S. M. Sichilalu, X. Xia, Optimal control

- of heat pump water heater-instantaneous shower using integrated renewable-grid energy systems, *Applied Energy* 201 (2017) 332–342.
- [128] M. Starke, J. Munk, H. Zandi, T. Kuruganti, H. Buckberry, J. Hall, J. Leverette, Real-time mpc for residential building water heater systems to support the electric grid, in: 2020 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT), IEEE, 2020, pp. 1–5.
- [129] C. Passenberg, D. Meyer, J. Feldmaier, H. Shen, Optimal water heater control in smart home environments, in: 2016 IEEE International Energy Conference (ENERGYCON), IEEE, 2016, pp. 1–6.
- [130] J. Dragoña, J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, L. Helsen, All you need to know about model predictive control for buildings, *Annual Reviews in Control* 50 (2020) 190–232. doi:<https://doi.org/10.1016/j.arcontrol.2020.09.001>. URL <https://www.sciencedirect.com/science/article/pii/S1367578820300584>
- [131] A. Afram, F. Janabi-Sharifi, Theory and applications of hvac control systems—a review of model predictive control (mpc), *Building and Environment* 72 (2014) 343–355.
- [132] A. Kathirgamanathan, M. De Rosa, E. Mangina, D. P. Fynn, Data-driven predictive control for unlocking building energy flexibility: A review, *Renewable and Sustainable Energy Reviews* 135 (2021) 110120.
- [133] F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, M. Morari, Use of model predictive control and weather forecasts for energy efficient building climate control, *Energy and Buildings* 45 (2012) 15–27.
- [134] A. Parisio, L. Fabbietti, M. Molinari, D. Varagnolo, K. H. Johansson, Control of hvac systems via scenario-based explicit mpc, in: 53rd IEEE conference on decision and control, IEEE, 2014, pp. 5201–5207.
- [135] M. Klaučo, M. Kvasnica, Explicit mpc approach to pmv-based thermal comfort control, in: 53rd IEEE conference on decision and control, IEEE, 2014, pp. 4856–4861.
- [136] J. Dragoña, M. Kvasnica, M. Klaučo, M. Fikar, Explicit stochastic mpc approach to building temperature control, in: 52nd IEEE Conference on Decision and Control, 2013, pp. 6440–6445. doi:10.1109/CDC.2013.6760908.
- [137] Y. Yang, G. Hu, C. J. Spanos, Hvac energy cost optimization for a multizone building via a decentralized approach, *IEEE Transactions on Automation Science and Engineering* 17 (4) (2020) 1950–1960.
- [138] Y. Yang, S. Srinivasan, G. Hu, C. J. Spanos, Distributed control of multizone hvac systems considering indoor air quality, *IEEE Transactions on Control Systems Technology* 29 (6) (2021) 2586–2597.
- [139] X. Zhang, W. Shi, B. Yan, A. Malkawi, N. Li, Decentralized and distributed temperature control via hvac systems in energy efficient buildings, *arXiv preprint arXiv:1702.03308* (2017).
- [140] N. Radhakrishnan, Y. Su, R. Su, K. Poolla, Token based scheduling for energy management in building hvac systems, *Applied energy* 173 (2016) 67–79.
- [141] N. Radhakrishnan, S. Srinivasan, R. Su, K. Poolla, Learning-based hierarchical distributed hvac scheduling with operational constraints, *IEEE Transactions on Control Systems Technology* 26 (5) (2017) 1892–1900.
- [142] Y. Yang, G. Hu, C. J. Spanos, Stochastic optimal control of hvac system for energy-efficient buildings, *IEEE Transactions on Control Systems Technology* 30 (1) (2021) 376–383.
- [143] B. Sun, P. B. Luh, Q.-S. Jia, Z. Jiang, F. Wang, C. Song, 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) (2012) 588–602.
- [144] B. Sun, P. B. Luh, Q.-S. Jia, B. Yan, Event-based optimization within the lagrangian relaxation framework for energy savings in hvac systems, *IEEE Transactions on Automation Science and Engineering* 12 (4) (2015) 1396–1406.
- [145] L. Yu, S. Qin, M. Zhang, C. Shen, T. Jiang, X. Guan, A review of deep reinforcement learning for smart building energy management, *IEEE Internet of Things Journal* (2021).
- [146] M. Han, R. May, X. Zhang, X. Wang, S. Pan, D. Yan, Y. Jin, L. Xu, A review of reinforcement learning methodologies for controlling occupant comfort in buildings, *Sustainable Cities and Society* 51 (2019) 101748.
- [147] K. Mason, S. Grijalva, A review of reinforcement learning for autonomous building energy management, *Computers & Electrical Engineering* 78 (2019) 300–312.
- [148] Z. Wang, T. Hong, Reinforcement learning for building controls: The opportunities and challenges, *Applied Energy* 269 (2020) 115036.
- [149] T. Yang, L. Zhao, W. Li, A. Y. Zomaya, Reinforcement learning in sustainable energy and electric systems: A survey, *Annual Reviews in Control* 49 (2020) 145–163.
- [150] T. Wei, Y. Wang, Q. Zhu, Deep reinforcement learning for building hvac control, in: Proceedings of the 54th annual design automation conference 2017, 2017, pp. 1–6.
- [151] L. Yu, W. Xie, D. Xie, Y. Zou, D. Zhang, Z. Sun, L. Zhang, Y. Zhang, T. Jiang, Deep reinforcement learning for smart home energy management, *IEEE Internet of Things Journal* 7 (4) (2019) 2751–2762.
- [152] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, D. Wierstra, Continuous control with deep reinforcement learning, *arXiv preprint arXiv:1509.02971* (2015).
- [153] L. Yu, Y. Sun, Z. Xu, C. Shen, D. Yue, T. Jiang, X. Guan, Multi-agent deep reinforcement learning for hvac control in commercial buildings, *IEEE Transactions on Smart Grid* 12 (1) (2020) 407–419.
- [154] V. Hanumaiah, S. Genc, Distributed multi-agent deep reinforcement learning framework for whole-building hvac control, *arXiv preprint arXiv:2110.13450* (2021).
- [155] S. Iqbal, F. Sha, Actor-attention-critic for multi-agent reinforcement learning, in: International Conference on Machine Learning, PMLR, 2019, pp. 2961–2970.
- [156] W. Arnold, T. Srivastava, L. Spangher, U. Agwan, C. Spanos, Adapting surprise minimizing reinforcement learning techniques for transactive control, in: Proceedings of the Twelfth ACM International Conference on Future Energy Systems, e-Energy '21, Association for Computing Machinery, New York, NY, USA, 2021, p. 488–492. doi:10.1145/3447555.3466590. URL <https://doi.org/10.1145/3447555.3466590>
- [157] C. Finn, P. Abbeel, S. Levine, Model-agnostic meta-learning for fast adaptation of deep networks, *CoRR abs/1703.03400* (2017). *arXiv:1703.03400*.

- URL <http://arxiv.org/abs/1703.03400>
- [158] D. Jang, L. Spangher, T. Srivistava, M. Khattar, U. Agwan, S. Nadarajah, C. Spanos, 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, Association for Computing Machinery, New York, NY, USA, 2021, p. 131–139. doi:10.1145/3486611.3486668. URL <https://doi.org/10.1145/3486611.3486668>
- [159] D. Jang, L. Spangher, M. Khattar, U. Agwan, C. Spanos, 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, Association for Computing Machinery, New York, NY, USA, 2021, p. 483–487. doi:10.1145/3447555.3466589. URL <https://doi.org/10.1145/3447555.3466589>
- [160] Z. Xu, G. Hu, C. J. Spanos, S. Schiavon, Pmv-based event-triggered mechanism for building energy management under uncertainties, *Energy and Buildings* 152 (2017) 73–85.
- [161] J. Drgoňa, D. Picard, M. Kvasnica, L. Helsen, Approximate model predictive building control via machine learning, *Applied Energy* 218 (2018) 199–216.
- [162] J. Drgoňa, L. Helsen, D. Vrabie, Stripping off the implementation complexity of physics-based model predictive control for buildings via deep learning, *Advances in Neural Information Processing Systems* (2019).
- [163] B. Karg, S. Lucia, Deep learning-based embedded mixed-integer model predictive control, in: 2018 european control conference (ecc), IEEE, 2018, pp. 2075–2080.
- [164] A. Afram, F. Janabi-Sharifi, A. S. Fung, K. Raahemifar, 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 (2017) 96–113.
- [165] K. Ku, J. Liaw, M. Tsai, T. Liu, Automatic control system for thermal comfort based on predicted mean vote and energy saving, *IEEE Transactions on Automation Science and Engineering* 12 (1) (2014) 378–383.
- [166] P. Ferreira, A. Ruano, S. Silva, E. Conceicao, Neural networks based predictive control for thermal comfort and energy savings in public buildings, *Energy and buildings* 55 (2012) 238–251.
- [167] P. Cappers, C. Goldman, D. Kathan, Demand response in u.s. electricity markets: Empirical evidence, *Energy* 35 (4) (2010) 1526–1535, demand Response Resources: the US and International Experience. doi:<https://doi.org/10.1016/j.energy.2009.06.029>. URL <https://www.sciencedirect.com/science/article/pii/S0360544209002461>
- [168] J. Hamari, J. Koivisto, H. Sarsa, Does gamification work? – a literature review of empirical studies on gamification, in: 2014 47th Hawaii International Conference on System Sciences, 2014, pp. 3025–3034. doi:10.1109/HICSS.2014.377.
- [169] I. C. Konstantakopoulos, A. R. Barkan, S. He, T. Veeravalli, H. Liu, C. Spanos, A deep learning and gamification approach to improving human-building interaction and energy efficiency in smart infrastructure, *Applied energy* 237 (2019) 810–821.
- [170] I. C. Konstantakopoulos, H. P. Das, A. R. Barkan, S. He, T. Veeravalli, H. Liu, A. B. Manasawala, Y.-W. Lin, C. J. Spanos, Design, benchmarking and explainability analysis of a game-theoretic framework towards energy efficiency in smart infrastructure, *arXiv preprint arXiv:1910.07899* (2019).
- [171] L. Spangher, A. Tawade, A. Devonport, C. Spanos, 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, Association for Computing Machinery, New York, NY, USA, 2019, p. 14–22. doi:10.1145/3363459.3363527. URL <https://doi.org/10.1145/3363459.3363527>
- [172] H. P. Das, I. C. Konstantakopoulos, A. B. Manasawala, T. Veeravalli, H. Liu, C. J. Spanos, A novel graphical lasso based approach towards segmentation analysis in energy game-theoretic frameworks, in: 2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA), IEEE, 2019, pp. 1702–1709.
- [173] H. P. Das, I. Konstantakopoulos, A. B. Manasawala, T. Veeravalli, H. Liu, C. J. Spanos, 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), 2020.
- [174] L. Spangher, A. Gokul, M. Khattar, J. Palakapilly, A. Tawade, A. Bouyamourn, A. Devonport, C. Spanos, 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, 2020, pp. 438–444.
- [175] B. Orland, N. Ram, D. Lang, K. Houser, N. Kling, M. Coccia, Saving energy in an office environment: A serious game intervention, *Energy and Buildings* 74 (2014) 43–52.
- [176] D. Geelen, D. Keyson, S. Boess, H. Brezet, Exploring the use of a game to stimulate energy saving in households, *Journal of Design Research* 14 10 (1–2) (2012) 102–120.
- [177] M. Casals, M. Gangolells, M. Macarulla, N. Forcada, A. Fuertes, R. V. Jones, Assessing the effectiveness of gamification in reducing domestic energy consumption: Lessons learned from the energaware project, *Energy and Buildings* 210 (2020) 109753.
- [178] D. S. Watson, S. Kiliccote, N. Motegi, M. A. Piette, Strategies for demand response in commercial buildings (2006).
- [179] P. Alstone, J. Potter, M. A. Piette, P. Schwartz, M. A. Berger, L. N. Dunn, S. J. Smith, M. D. Sohn, A. Aghajanzadeh, S. Stensson, et al., 2025 california demand response potential study-charting california's demand response future. final report on phase 2 results, Tech. rep., Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States) (2017).
- [180] M. H. Albadi, E. F. El-Saadany, Demand response in electricity markets: An overview, in: 2007 IEEE power engineering society general meeting, IEEE, 2007, pp. 1–5.
- [181] Federal Energy Regulatory Commission, Energy primer: a handbook for energy market basics. URL https://www.ferc.gov/sites/default/files/2020-06/energy-primer-2020_0.pdf

- [182] C. Deb, F. Zhang, J. Yang, S. E. Lee, K. W. Shah, A review on time series forecasting techniques for building energy consumption, *Renewable and Sustainable Energy Reviews* 74 (2017) 902–924.
- [183] J. Jazaeri, T. Alpcan, R. Gordon, M. Brandao, T. Hoban, C. Seeling, Baseline methodologies for small scale residential demand response, in: 2016 IEEE Innovative Smart Grid Technologies-Asia (ISGT-Asia), IEEE, 2016, pp. 747–752.
- [184] L. Xu, S. Wang, R. Tang, Probabilistic load forecasting for buildings considering weather forecasting uncertainty and uncertain peak load, *Applied energy* 237 (2019) 180–195.
- [185] Y. Kim, H.-g. Son, S. Kim, Short term electricity load forecasting for institutional buildings, *Energy Reports* 5 (2019) 1270–1280.
- [186] T. X. Nghiem, C. N. Jones, Data-driven demand response modeling and control of buildings with gaussian processes, in: 2017 American Control Conference (ACC), IEEE, 2017, pp. 2919–2924.
- [187] J. L. Mathieu, P. N. Price, S. Kiliccote, M. A. Piette, Quantifying changes in building electricity use, with application to demand response, *IEEE Transactions on Smart Grid* 2 (3) (2011) 507–518.
- [188] D. Jung, V. B. Krishna, W. G. Temple, D. K. Yau, Data-driven evaluation of building demand response capacity, in: 2014 IEEE International Conference on Smart Grid Communications (SmartGridComm), IEEE, 2014, pp. 541–547.
- [189] S. Burger, J. P. Chaves-Avila, C. Batlle, I. Perez-Arriaga, The value of aggregators in electricity systems (2016).
- [190] U. Agwan, C. J. Spanos, K. Poolla, Asset participation and aggregation in incentive-based demand response programs, in: 2021 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm), IEEE, 2021, pp. 89–94.
- [191] P. Shinde, M. Amelin, A literature review of intraday electricity markets and prices, 2019 IEEE Milan PowerTech (2019) 1–6.
- [192] A. Brusafferri, M. Matteucci, P. Portolani, A. Vitali, Bayesian deep learning based method for probabilistic forecast of day-ahead electricity prices, *Applied Energy* 250 (2019) 1158–1175.
- [193] U. Agwan, Optimal prosumer aggregations: Design and modeling (2020).
- [194] W. Saad, Z. Han, H. V. Poor, Coalitional game theory for cooperative micro-grid distribution networks, in: 2011 IEEE international conference on communications workshops (ICC), IEEE, 2011, pp. 1–5.
- [195] D. Wu, X. Ma, S. Huang, T. Fu, P. Balducci, Stochastic optimal sizing of distributed energy resources for a cost-effective and resilient microgrid, *Energy* 198 (2020) 117284.
- [196] M. Quashie, C. Marnay, F. Bouffard, G. Joós, Optimal planning of microgrid power and operating reserve capacity, *Applied Energy* 210 (2018) 1229–1236.
- [197] U. Agwan, K. Poolla, C. J. Spanos, Optimal composition of prosumer aggregations, in: 2021 IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe), IEEE, 2021, pp. 1–5.
- [198] S.-J. Kim, G. B. Giannakis, An online convex optimization approach to real-time energy pricing for demand response, *IEEE Transactions on Smart Grid* 8 (6) (2016) 2784–2793.
- [199] N. Liu, X. Yu, C. Wang, C. Li, L. Ma, J. Lei, Energy-sharing model with price-based demand response for microgrids of peer-to-peer prosumers, *IEEE Transactions on Power Systems* 32 (5) (2017) 3569–3583.
- [200] L. Jia, Q. Zhao, L. Tong, Retail pricing for stochastic demand with unknown parameters: An online machine learning approach, in: 2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton), IEEE, 2013, pp. 1353–1358.
- [201] U. Agwan, L. Spangher, W. Arnold, T. Srivastava, K. Poolla, C. J. Spanos, Pricing in prosumer aggregations using reinforcement learning, in: Proceedings of the Twelfth ACM International Conference on Future Energy Systems, 2021, pp. 220–224.
- [202] B. Dong, Y. Liu, W. Mu, Z. Jiang, P. Pandey, T. Hong, B. Olesen, T. Lawrence, Z. O’Neil, C. Andrews, et al., A global building occupant behavior database, *Scientific data* 9 (1) (2022) 1–15.
- [203] G. Fierro, M. Pritoni, M. AbdelBaky, D. Lengyel, J. Leyden, A. Prakash, P. Gupta, P. Raftery, T. Pefter, G. Thomson, et al., Mortar: an open testbed for portable building analytics, *ACM Transactions on Sensor Networks (TOSN)* 16 (1) (2019) 1–31.
- [204] C. Miller, A. Kathirgamanathan, B. Picchetti, P. Arjunan, J. Y. Park, Z. Nagy, P. Raftery, B. W. Hobson, Z. Shi, F. Meggers, The building data genome project 2, energy meter data from the ASHRAE great energy predictor iii competition, *Scientific data* 7 (1) (2020) 1–13.
- [205] B. Balaji, A. Bhattacharya, G. Fierro, J. Gao, J. Gluck, D. Hong, A. Johansen, J. Koh, J. Ploennigs, Y. Agarwal, et al., Brick: Metadata schema for portable smart building applications, *Applied energy* 226 (2018) 1273–1292.