

# Data-Driven Methods for Building Control

## – A Review and Promising Future Directions

Emilio T. Maddalena, Yingzhao Lian and Colin N. Jones

Automatic Control Laboratory, École Polytechnique Fédérale de Lausanne,  
Route Cantonale, 1015 Lausanne, Switzerland

{emilio.maddalena,yingzhao.lian,colin.jones}@epfl.ch

### Abstract

A review of the heating, ventilation and air-conditioning control problem for buildings is presented, and particular emphasis is given to its distinguishing features. Next, we not only examine how data-driven algorithms have been exploited to tackle the main challenges present in this area, but also point to promising future investigations both from theoretical and from practical viewpoints. Rule based control, reinforcement learning, model predictive control (MPC), and learning MPC techniques are compared on the basis of four attributes that we expect an ideal solution to possess. Finally, on-line learning MPC with guarantees is recognized as an approach with high potential that needs to be further investigated by researchers. Such a solution is likely to be accepted by practitioners since it meets the industry expectations of reduced deployment time and costs.

*Keywords:* Heating, ventilation and air-conditioning (HVAC); building control; model predictive control (MPC); machine learning; reinforcement learning.

### 1 Introduction

Large heating, ventilation and air-conditioning (HVAC) facilities need to be constantly supervised by technicians, which are hired by companies to ‘adjust settings’, ‘ensure optimal performance’ and sometimes even operate some subsystems manually [Stluka et al., 2018, Royapoor et al., 2018]. It is reasonable to assume that a well designed automatic controller would perform better than a human in such tasks, especially when highly complex objectives are considered such as time-varying electricity prices for instance. Advanced control techniques are particularly required when consumers provide ancillary services to the power grid, adjusting their demand in real-time [Gils, 2014, Maasoumy et al., 2014c]. Residential and commercial buildings have been recognized as suitable candidates for this position due to their significant energy consumption allied with temporal flexibility, mainly associated with HVAC [Motegi et al., 2007, Beil et al., 2016]. However, despite efforts from researchers to demonstrate the effectiveness of several rigorous control methodologies for the different HVAC hierarchy levels (e.g. [Arguello-Serrano and Velez-Reyes, 1999, Huaguang and Cai, 2002, Shin et al., 2002, Anderson et al., 2008, Gondhalekar et al., 2013, Wang and Hu, 2018]), industrial solutions still rely on PID regulators for low-level devices, and rule-based schemes combined with finite-state machines for the high-level supervisory system [Salsbury, 2005, Wang and Ma, 2008, Stluka et al., 2018].

Modern buildings and in particular office buildings are equipped with a plethora of sensors, usually managed by a central building automation system (BAS): temperature probes, occupancy detectors, luminosity and air-pressure transducers, power meters, among many others. Some authors state that low resolution information and missing samples are common problems faced in HVAC data collection [Reijula et al., 2013, Afram and Janabi-Sharifi, 2014b, Žáčková et al., 2014, Wijayasekara and Manic, 2015], yet the abundance of measuring instruments has rather been regarded as an opportunity to develop data-driven solutions to existing problems in the area. Forecasting weather and thermal loads, obtaining non-parametric models of zones, operating chillers at optimal conditions, estimating occupants’ thermal comfort, and controlling blinding devices using reinforcement learning (RL) are some of the applications of machine learning in this domain [Cai et al., 2019, Xuan et al., 2019, Østergård et al., 2018, Park et al., 2018,

Atthajariyakul and Leephakpreeda, 2005, Park et al., 2019]. Another important aspect to be considered is that, by automating several time-consuming tasks present in the design of model predictive controllers (MPC), machine learning can also contribute to the popularization of this flexible and powerful framework [Aswani et al., 2012a, Ma et al., 2012b, Cigler et al., 2013].

Recent reviews in the area of building HVAC systems have addressed the available modelling approaches for individual equipment and thermal zones [Whalley and Abdul-Ameer, 2011, Afram and Janabi-Sharifi, 2014a, Afroz et al., 2018]; the applicability of MPC strategies in the field [Serale et al., 2018], and an industry perspective of how relevant classical and modern control are to the field according to engineers of different nationalities [Royapoor et al., 2018]. Additional surveys have investigated the problem of predicting energy consumption [Zhao and Magoulès, 2012, Wei et al., 2018] and implementing energy savings strategies from a system architecture (static) point of view such as exploiting heat recovery and thermal storage [Chua et al., 2013, Vakiloroyaya et al., 2014]. Fuzzy modelling and control, genetic algorithms and artificial neural networks were covered in [Singh et al., 2006, Kalogirou, 2009], and RL was briefly discussed in [Wang and Ma, 2008]. The last documents however targeted people with little knowledge on machine learning and considerable familiarity with HVAC plants and, moreover, were published ten or more years ago.

*Contribution:* In this paper, HVAC plants are analyzed from the perspective of a control engineer and an up-to-date outlook of how data-driven techniques are being used in the building control context is offered. In addition, we point to obstacles that still have to be overcome and unexplored areas, hence suggesting ideas for promising future investigations. Particular emphasis will be given to deploying advanced control techniques in already existing buildings while minimizing retrofit. Although important for the field, the subject of fault-detection and diagnosis is not included in the present survey; the reader is referred to [Katipamula and Brambley, 2005a, Katipamula and Brambley, 2005b, Namburu et al., 2007, Beghi et al., 2016, Shahnazari et al., 2019] for an overview on this matter. As for the terminology, we say a given method is ‘data-driven’ simply if it receives a data-set as an input and yields some product such as a diagnosis, a mathematical model, or a control law. Although the words ‘adaptive’ and ‘learning’ are sometimes used by the control community as synonyms (see [Fu, 1970, Sutton et al., 1992, Gaudio et al., 2019] for more information about the difference), we use the terms ‘adaptive’ and ‘*on-line* learning’ interchangeably especially when referring to MPC. Both expressions refer to techniques where mathematical models are updated during operation.

## 2 The HVAC control problem

### 2.1 The goals of HVAC systems

The ultimate goal of any HVAC facility is to regulate temperatures in different parts of a building<sup>1</sup> while simultaneously controlling additional variables related to its air quality. When rooms are occupied by people, the first *desideratum* is referred to as thermal comfort, whose definition according to the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) [ANSI/ASHRAE, 2010] is

“The condition of mind that expresses satisfaction with the thermal environment and is assessed by subjective evaluation.”

The two most common methods used to quantify thermal comfort are the predicted mean vote (PMV) and the predicted percentage of dissatisfied (PPD) [ISO, 2005]. These quantities depend on factors such as clothing insulation, a person’s metabolic rate, relative air velocity and, naturally, the ambient air temperature. In order to circumvent the subjectiveness of both the PMV and PPD indexes, the majority of authors as well as practitioners opt for simply defining temperature envelopes to be respected in each environment [Mantovani and Ferrarini, 2015, Ma et al., 2015, Ben-Nakhi and Mahmoud, 2017]. The reader is referred to [Garnier et al., 2014, Garnier et al., 2015, Gupta and Kar, 2018] for exceptions of this tendency. Besides temperature, carbon dioxide levels (CO<sub>2</sub>), air humidity and pressure, and the presence of particles must also be controlled, especially in the

<sup>1</sup>The acronym HVAC is occasionally also employed in the field of automotive and aeronautical air conditioning, but these subjects are outside the scope of this document.



the system flexibility by accumulating hot or cold water, sometimes in the form of ice [Zhang et al., 2007, Ma et al., 2012b]. Moreover, humidifiers based on water evaporators may be placed at the end of the handling units to compensate for the dehumidification caused by previous stages. Instead of an electrical resistance, the supply AHU may feature a heating coil in which warm water coming from a boiler or a solar collector circulates [Hsieh et al., 2017, Rathod et al., 2019]. In locations where weather conditions vary significantly over the year, it is common to also find heat pumps (HP) in the facilities; modern equipment are able to operate in both cooling and warming modes as needed, and are called reversible chillers or reversible heat pumps [Stabat and Marchio, 2009]. As opposed to the forced-air HVAC system of Figure 1, which is more popular in large commercial buildings, shopping malls and airports, there are radiant and variable-refrigerant-flow (VRF) plants. In radiant cooling and radiant heating, currents circulate in electric cables or fluid flows through pipes that either belong to heat exchange devices or are directly installed inside building elements such as floors and walls [Široký et al., 2011, Andersen et al., 2014, Zheng et al., 2017]. On the other hand, VRFs are composed of an outdoor unit responsible for rejecting the internal heat, and indoor units that receive a refrigerant flow (liquid or gas) and chill the internal air through pure recirculation [Aynur et al., 2009]. Although simple, these approaches clearly do not allow for controlling the ambient humidity, nor CO<sub>2</sub> levels.

If a set of buildings is considered instead of a single one, heat may be produced in a centralized fashion giving rise to the concept of district-heating (DH) networks<sup>2</sup>. In such case, heat production typically takes place in cogeneration plants and is then distributed to individual buildings through buried insulated pipes. The technologies behind DH networks have evolved over the years and nowadays DH is particularly present in Nordic and Eastern European countries [Lund et al., 2014, Werner, 2017, Darivianakis et al., 2017].

Depending on the system architecture, the energy consumption distribution in the system may vary; still, chillers tend to have higher nominal power when compared to the remaining equipment (see e.g. [Fong et al., 2006, Lecamwasam et al., 2012]). Pumps located in the chilled water and condenser water circuits also contribute to the overall demand and, within the AHUs, electric heaters and motors driving the blowers as well. Instead of operating the AHU fans always at maximum speed and modifying the air flux through dampers, VAV control can be achieved by varying the blowers speed. This practice leads to energy savings and is possible if electronic motor drives are employed.

## 2.3 Distinguishing characteristics

We summarize herein the main features (F) of the building climate control problem as:

- F1 Large sampling periods: With the exception of laboratory experiments in scaled HVAC plants such as [Xi et al., 2007, Obando et al., 2014], most supervisory controllers gather data at every 15 min or longer periods. In the DR scenario however, several nested control loops exist, each one operating with substantially different sampling frequencies (see [Gorecki et al., 2017]);
- F2 Low-quality measurements: It has been acknowledged by several authors that sensors tend to be poorly calibrated and have low-resolution [Afram and Janabi-Sharifi, 2014b, Wijayasekara and Manic, 2015]. At the same time, it should be highlighted that some variables are acquired at a much faster rate than the supervisory system control frequency, thus allowing for the application of filtering techniques or simply downsampling;
- F3 Operations-critical: The occupancy schedule cannot usually be modified for the sake of conducting identification tests. Moreover, strongly exciting input signals are impracticable while the building is occupied since comfort bounds might be violated. Identification could therefore be conducted during unoccupied periods, e.g. nighttime. This procedure however would not capture important effects caused by occupants in the environment (see subsection 3.1);
- F4 Difficult modelling: Especially true for large buildings, obtaining a representative model of its climate dynamics for reliable multi-step ahead predictions that can be used in different weather conditions is challenging [Privara et al., 2013, Žáčková et al., 2014, Váňa et al., 2014, Killian and Kozek, 2016]. This complication is due to the usual lack of precise information about the numerous construction parameters, preventing

---

<sup>2</sup>The analogous concept of district-cooling networks also exists.

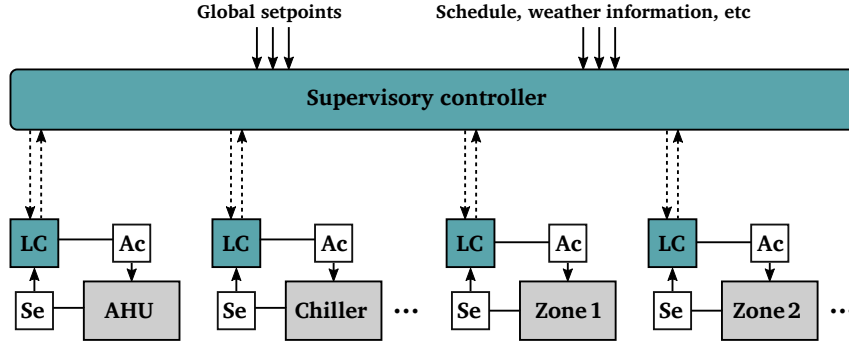


Figure 2: The control architecture of a building HVAC system: supervisory controller, low-level device controllers (LC), actuators (Ac) and sensors (Se).

white-box models from being practical; the effects of external disturbances, e.g. the conduction and radiation solar gains; and continuous internal gains variations caused by changing occupancy rates.

In addition, depending on the electricity contract being applied and the plant architecture, the objective of the control system may range from a simple minimum effort formulation to a more complex reference tracking problem [Qureshi and Jones, 2018].

## 2.4 The orthodox and the idealistic solution

The classical approach to the building HVAC control problem relies on a hierarchical structure as shown in Figure 2. Low-level controllers generally implemented in programmable logic controller (PLC) hardware handle individual devices through field sensors and actuators. An example of such a control loop would be the air temperature regulation inside the input AHU with absolute temperature sensors and actuation in the three-way valve that changes the chilled water flow through the cooling coils. In the upper layer, a high-level supervisory system manages the low-level operation by sending appropriate reference signals, e.g. speed references for the water pumps and positions for the dampers. Communication between the two layers is usually carried out through either the Modbus or BACnet (Building Automation and Control network) protocols. Typically PID controllers or in simpler cases open-loop commands are employed at the device level<sup>3</sup>, while rule based controllers (RBC) incorporating memory in the form of a state machine are used for the management system, which can include information regarding the building’s schedule and current weather. It is however well known that tuning PID parameters and decision rules is a daunting task especially when one seeks to optimize performance. Plus, in the orthodox setting with RBC, the energy demand can only be adjusted coarsely in real-time, not allowing for following an explicit DR reference for example.

Based on system description given in this section and the highlighted challenges, we propose the following list of attributes (A) that an ideal solution for the high-level supervision problem would be endowed with:

- A1 Lifetime adaptability: In view of F5, the ideal technique must require only a rough mathematical description of the dynamics in the initialization phase – in case it is model based – and refine it over time during operation. If the chosen technique is model free, actions or policies are to be updated as more data is available to improve performance;
- A2 Safety<sup>4</sup>: The process of improving performance or the model on-line should be carried out with guarantees (either deterministic or probabilistic) of safety because of F4. The chosen identification technique should be preceded by a proper data pre-processing stage due to F2, and should possess good generalization capabilities given that aggressive excitation is generally not permitted;
- A3 Incorporation of exogenous signal forecasts and complex objectives: In view of the contributions of both internal and external uncontrollable factors to the building

<sup>3</sup>State-of-the-art control strategies may be used as LCs in the case of heat pumps and compressors in general [Wallace et al., 2012, Wallace et al., 2015].

<sup>4</sup>In the present context, safety translates to constraint satisfaction, i.e. compliance with the regulations concerning the controlled climate variables. It does not refer to the ‘safety-critical’ or ‘life-critical’ concepts.

dynamics and costs – namely occupancy levels, outside temperature and time-varying energy prices – exploiting reliable predictions becomes indispensable to maximize performance;

- A4 Scalability: In case optimization is present in the loop, the mathematical program is to be formulated in a scalable and numerically stable manner seeing that, even though the available computational time is large (F1), the number of variables is also considerable. In the RL case, the learning process should not take a prohibitive amount of time to converge even when the state-action space dimension is large.

### 3 Data-driven techniques in building control

In this section a variety of works coming both from academia and industry are analyzed and clustered, indicating what the main applications of data-driven techniques in the HVAC context are. A distinction is made between control methodologies that process all information off-line and those that carry on the learning process during operation.

#### 3.1 Occupancy and occupant behavior predictions

It is possible to classify the impacts of occupants on the indoor climate as passive, which includes contributing with metabolic heat gains and increasing CO<sub>2</sub> levels; and active, which comprises acting on windows, doors, shading devices, and adjusting local zones’ temperature setpoints as shown in Figure 3. Therefore, occupancy pattern prediction should be combined with an estimate of the occupants behavior to correctly assess their effects on the HVAC control loop. An informative study on the energy savings potential associated with such information was carried out in [Oldewurtel et al., 2013].

The more elementary problem of predicting if either a zone is occupied or unoccupied was studied in [Mahdavi and Tahmasebi, 2015], where eight offices located in a university building were analyzed and three different prediction models were compared. The authors however state that, even though the two probabilistic models and the deterministic one were carefully trained with reliable data, the predictive performance of all models was only modest. A more successful investigation was carried out in [Peng et al., 2018] for 11 different zones: eight single-person offices, two multi-person offices and one meeting room. By using data from motion sensors collected every 10 minutes, occupants presence or absence and presence duration patterns were learned respectively with the k-means and k-nearest neighbors methods. Additionally, an RBC strategy was employed to turn off cooling when spaces were vacant. Most of the studies in this area are nevertheless focused on estimating the actual *number* of occupants, instead of the presence/absence of them [Capozzoli et al., 2017, Dong et al., 2018, Li and Dong, 2018]. Stochastic agent-based models [Liao and Barooah, 2010, Liao et al., 2012] and Markov chains in different flavors [Page et al., 2008, Ai et al., 2014, Flett and Kelly, 2016] are the most common design choices. Finally, the authors of [Chen et al., 2017] considered a midway problem, estimating levels of occupancy (zero, low, medium and high) with a deep convolutional bidirectional neural network in a classification setting.

Occupant behavior prediction is a large research area by itself [Yana et al., 2015], but the following works are considered more relevant to estimating internal HVAC disturbances. Haldi described the fundamental means by which an occupant could affect the building’s thermal evolution and proposed the use of Bernoulli, discrete-time Markov, and continuous-time Markov processes to predict future occurrences of such actions [Haldi and Robinson, 2009, Haldi, 2010]. Later, efforts were made to differentiate more subtle details for instance not only forecasting window states simply as open or closed, but also opening angles [Schweiker et al., 2012]. A review on occupants control over natural ventilation is given in [Roetzel et al., 2010, Barthelmes et al., 2017], while more general aspects of behavior prediction including acting on blinds were examined in [Lee and Malkawi, 2014].

Different techniques for occupancy and occupant behavior estimation require different variables to be measured from the environment. In some cases simple motion detectors, CO<sub>2</sub> level and luminosity sensors are employed [Candanedo and Feldheim, 2016], whereas other studies explore cameras and the connection of users’ smartphones to existing WiFi networks [Balaji et al., 2013]. Clearly, the best solution to a particular building lies in balancing predictive performance and retrofitting, as acquiring and installing numerous additional sensors in every zone is not usually desirable in practice. If no occupancy sensors are employed and the contribution of the internal heat gain to the overall dynamics is significant, the system

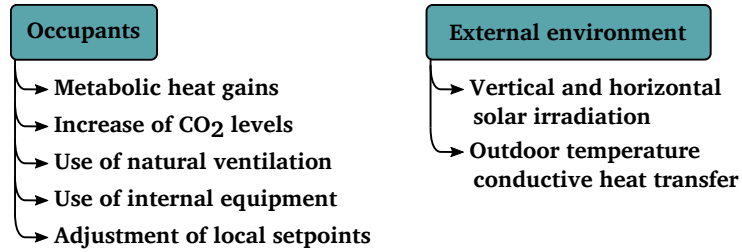


Figure 3: A list of the main internal and external factors that influence the indoor climate.

identification stage may be compromised. This is worsened by the usual correlation between the building occupancy and time. The reader is referred to [Aswani et al., 2012b] for a semi-parametric way of dealing with this problem.

### 3.2 Thermal loads and energy demand forecasting

Accurate predictions of a building’s thermal load and HVAC energy consumption are important respectively for making effective use of thermal storage systems and determining the day-ahead baseline power consumption in the DR scenario [Henze et al., 2004, Fabietti et al., 2018a]. In [Rahman and Smith, 2018] for example, deep recurrent neural networks were employed to estimate future heating demands of an academic campus building and, subsequently, and help design a TES system. The prediction techniques are categorized according to their time-horizon, e.g. hours, days, weeks [Djukanovic et al., 1993, Zhou et al., 2008, Lam et al., 2010], and according to their nature, either time-series or multi-variate regression. The most common machine learning models in the regression case are neural networks [Ferrano and Wong, 1990, Zhao et al., 2016], decision trees [Johansson et al., 2017], comparisons among them can be found in [Penya et al., 2011, Zhang et al., 2015]. Some works address the broader problem of predicting the future energy consumption of a building as a whole, encompassing HVAC loads, lightning, plugs, elevators, etc. This was the question investigated in [Zeng et al., 2019], where a total of six different edifices including shopping malls and hotels were studied. The authors stressed the importance of data pre-processing to achieve reasonable results since they faced missing samples – constituting 9.62% of all total information for a particular office edifice – and poor calibration of the energy meters. Focusing exclusively on cooling and heating loads, twelve buildings with distinct geometries but the same volume were examined through simulations in [Chou and Bui, 2014]. The geometrical properties of the constructions were used as inputs and both thermal loads as outputs, and favourable numerical results were obtained for an ensembled model of ANN in conjunction with SVM. See also [Yun et al., 2012, Bacher et al., 2013] for similar studies. Another problem highly relevant in practice is power disaggregation, where the contribution of many energy consuming devices are mixed due to physical sensing constraints, but have to be decomposed usually into equipment-level power information [Yan et al., 2012, Ji et al., 2015, Rahman et al., 2018].

### 3.3 Batch learning and building climate control

Already in [Sousa et al., 1997], a Takagi-Sugeno fuzzy model was developed from a set of input-output measurements of a simple single-zone heating problem. A nonlinear MPC (NMPC) controller was then designed for the single-input single-output system and successful real-time experiments were reported. Although simple, the case study became influential, possibly due to the attention given both to the theoretical foundations and implementation details. Later, multiple Takagi-Sugeno fuzzy models were employed to tackle the predictive control of a more realistic AHU setup [He et al., 2005]. The simultaneous control of temperature and humidity in a single test-chamber was studied [Xi et al., 2007]. There, two support vector regression models were trained to capture the individual dynamics and NMPC was adopted to control both the speed of a fan (VAV strategy) and the position of a three-way valve. Reference tracking experimental results were carried out to demonstrate the performance of the methodology.

In more recent years, case studies have become more and more complex, tackling realistic residential and commercial buildings. In [Ma et al., 2012b], the design of an MPC controller was considered for the optimal operation of an HVAC system composed of two series chillers, a large TES tank and a cooling tower. Historical data were used to build

a linear model of the University of California Merced campus thermal load, as well as time-varying bounds for its values. The authors in [Vána et al., 2014] describe the control of a large office building in Belgium equipped with radiant heating (pipes conducting hot water). By defining a resistance-capacitance model of the system, using several heat fluxes as inputs and zones’ temperatures as outputs, a linear state-space model was identified with two months recorded data. MPC was applied as the high-level supervisory layer and a low-level controller was used to act on the two-way valves, guaranteeing the correct mass flow rate and supply water temperature to achieve specified reference heat fluxes. [Aswani et al., 2012b] exploited semi-parametric regression to capture air-conditioning and VAV effects in two distinct environments located in the University of California Berkeley, whereas gray-box modelling of a considerably complete residential HVAC system in Ontario, Canada was tackled in [Afram and Janabi-Sharifi, 2015]. Subsequently, the later authors also used artificial neural networks (ANNs) for modelling and NMPC control of the same house [Afram et al., 2017], but instead of specifying heat fluxes, the supervisory system generated temperature references at each sampling period for local PID controllers. A comprehensive study of the design and deployment of an MPC controller in a Swiss office building was given in [Sturzenegger et al., 2016], which included an interesting cost-benefit analysis of the development, hardware acquisition and installation phases of the project – an aspect usually neglected in most academic works.

Besides the diverse implementations of MPC, some research has also been done on alternative supervisory systems for efficient HVAC operation. In [Pallonetto et al., 2019] for example, two rule-based approaches for DR services in a typical Irish house. The RBC technique which incorporated prediction modules was shown to outperform the standard decision tree in several metrics. Four zones of the Adelaide airport in Australia were analyzed in [Huang et al., 2015] with ANNs trained to perform multi-step-ahead prediction. A rule-based on/off scheme was then employed to minimize energy consumption based on the ANNs prediction and additional occupancy information. Similar investigations combining an ANN predictor and RBC can be found in [Moon and Kim, 2010, Moon, 2012, Moon et al., 2013]. Directly optimizing the low-level controllers setpoints in a static fashion is a simplified way of minimizing costs, which is particularly adopted when designing a dynamical model of the plant is considered infeasible [Lu et al., 2005, Kusiak et al., 2014, Asad et al., 2017]. Having numerical optimization in-the-loop may be impractical when the available hardware has limited processing power; this is unfortunately the case in many small and medium-sized buildings<sup>5</sup>. The authors in [Domahidi et al., 2014] proposed the design of a hybrid MPC controller in the first step, and later approximating it either with SVMs or the Adaptive Boosting technique to extract a set of simple decision rules, which could be easily implemented even in PLCs. In a similar manner, [May-Ostendorp et al., 2011, Le et al., 2014] studied simplification strategies but with particular focus on window openings and blind positioning. Finally, [Drgoňa et al., 2018] reported in details the learning process of an MPC controller for a residential heating system in Belgium. The feature selection process including principal component analysis dimensionality reduction was discussed, and decision trees and time delay neural networks were compared against the original MPC controller as well as PID and RBC schemes.

### 3.4 On-line learning and building climate control

Continuous model and/or controller adaptation is a desirable feature in building HVAC control. In an article published in 1985, the Johnson Controls engineer Clay G. Nesler stated that [Nesler, 1985]

“To be accepted by the HVAC control industry, the self-tuning controller must:

1. Implement the PI control algorithm.
2. Not require any prior knowledge of the process (e.g., dead time).
3. Be sufficiently robust to operate without supervision.”

While nowadays more sophisticated techniques are to some extent accepted by practitioners [Royapoor et al., 2018], the requirement for little human intervention, including avoiding building intricate models, is still a strong *desideratum*. This can be easily interpreted as a demand for on-line safe adaptation and learning.

<sup>5</sup>There has been an increase in cloud-based monitoring solutions offered by HVAC companies in recent years, a trend that could overcome local hardware limitations.

Classical indirect and direct adaptive control techniques, i.e. respectively with and without explicit plant identification, were investigated in [Nesler, 1985, Chen et al., 1990, Salsbury, 2002, Bai et al., 2008], but experiments were carried out either in single rooms or in AHU subsystems. An adaptive neuro-fuzzy technique and self-tuning PID controllers based on fuzzy rules were studied in [Jian and Wenjian, 2000, Soyguder et al., 2009], however only simulation results were presented. By exploiting the strict passivity property of building thermodynamics [Mukherjee et al., 2012], iterative learning control (ILC) was applied in simulations to steer the states of six rooms into a temperature-humidity comfort ellipse [Minakais et al., 2017]. Later, the same authors conducted experimental results in a scaled six-room building, but focusing exclusively on the temperature problem [Minakais et al., 2019]. A similar approach was also adopted in [Lautenschlager and Lichtenberg, 2016], where ILC adjusts the output of an MPC controller to account for modeling errors.

On-line model refinement of building thermodynamics is still a matter of current investigation. The main issue being the operational-critical nature of the plant (F4), which often hinders the system identification process [Agbi et al., 2012, Aswani et al., 2012b]. Building predictive control employing linear time-varying models was examined in [P olka et al., 2016] and compared with two alternatives. The authors conclude that the proposed approach enjoys good predicting capabilities while preserving convexity of the optimization problem. It is moreover claimed that the theoretical guarantees of the formulation are those derived in [Falcone et al., 2008]. Model augmentation and the use of either an extended Kalman filter or an unscented Kalman filter for simultaneous parameter and state estimation was proposed in [Radecki and Hencsey, 2013, Maasoumy et al., 2014b, Maasoumy et al., 2014a]. Platt, on the other hand, started with a grey-box model and updated its parameter values with the recursive use of genetic algorithms for short-term predictions [Platt et al., 2010]. Even though the numerical results presented in the paper were promising, validation was performed in a single zone system and, if the technique were to be combined with MPC, the outcome would be a bilevel optimization problem. An alternative approach to on-line learning was taken in the so called NEUROBAT project and subsequent works [Morel et al., 2001, Lindel f et al., 2015], in which two ANNs were employed to learn the closed-loop building model and forecast the weather, being updated constantly to improve predictions over time. Control was accomplished by means of supervisory MPC and the experiments, performed over an extensive period of time in several test sites, yielded positive energy savings in all cases. Learning-based MPC (LMPC) is a predictive control strategy with solid theoretical foundations and independent of the machine learning tool selected by the user [Aswani et al., 2013]. This technique was applied to the air-conditioning control of a laboratory in order to reduce transient and steady-state energy consumption, with an on-line model update but fixed uncertainty description [Aswani et al., 2012a]. Finally, by applying results from the linear set-membership and robust MPC, the authors of [Tanaskovic et al., 2017] reduced the set of feasible thermal models at every time step, hence accounting for uncertainties, but minimizing conservativeness. Promising simulation results were reported.

An additional class of methods that has been explored in the field of building control is reinforcement learning [Henze and Schoenmann, 2003, Yu and Dexter, 2010, de Gracia et al., 2015, Park et al., 2019]. Two distinct variations of Q-learning, namely tabular and batch with an ANN function approximator, were used in [Yang et al., 2015] to control in simulations several elements of a low exergy building such as photovoltaic-thermal panels and a heat pump. In [Dalamagkidis et al., 2007] for example, a model-free formulation was adopted to control a reversible HP (off, low, medium and high heating, or cooling), a fan to provide ventilation (off, low or high) and the position of a window (closed, slightly open, open, wide open). Regulation of the PMV index was considered and simulation results showed that the RL strategy achieved performance similar to on/off and fuzzy PD controller, but only after a training period of four years. See also [Chen et al., 2018] for a similar study, where the same problem was observed. Deep reinforcement learning for VAV HVAC control was analyzed in [Wei et al., 2017] with an effective reduction of the temperature comfort zone violation after training. Nonetheless, the presented Q-learning scheme only started to converge after approximately 75 months for a simple 1-zone case, and 90 months for a 4-zones case, causing the proposal to be infeasible in practice. In a recent study [Kazmi et al., 2019] discusses how this lengthy process can be accelerated if more than one agent is learning the same model or policy, and information can be exchanged among them. A large case study was reported with 32 residential houses implementing the RL algorithm for thermostatically controlled loads. In [Claessens et al., 2018], a direct implementation of model-free RL to control of DH networks is reported with promising low learning times.

## 4 A perspective and future directions

### 4.1 Interfacing with already deployed industrial controllers

Clearly, buildings that feature operating HVAC systems are already equipped with some form of hierarchical control structure as discussed in previous sections. Technicians that manage PLCs as well as the supervisory system aim at ensuring a smooth and continuous operation of the overall plant, and hence it is critical for them to trust and feel comfortable with the control system as a whole [Henze, 2013]. According to the industrial survey [Stluka et al., 2018]

“The operations staff are not motivated to try new controls and optimization software, because operational reliability is their number one goal”.

Therefore, keeping already functional blocks in their place is certainly desirable in practice, and more complex control strategies should be able to interface with these components rather than replacing them. The central question is then how to deal with such black-boxes that execute proprietary firmware. Obtaining informative models of the closed-loop system is potentially challenging in this context. For example, the implementation of even simple PID controllers on industrial hardware is far from straightforward: dozens of parameters are involved, anti-windup schemes and low-pass filters are added to improve performance and avoid abrupt variations, etc [O’Connor, 2000, Johnson Controls, 2010, Siemens, 2014]. Even in smart thermostats for residential applications, manufacturers claim to implement ‘self-learning algorithms with PID response’ [Siemens, 2018]. Some preliminary research in the direction of working with already deployed controllers can be found for instance in the following three works.

In [Ferreira et al., 2012], four rooms of a university building with existing variable refrigerant flow and indoor heat exchangers were considered, each one equipped with a local Mitsubishi thermostat. A pseudo random binary signal (PRBS) constrained to  $[18, 27]^{\circ}\text{C}$  was employed as the temperature setpoint during 23 days for identification purposes, and the dynamics were captured by a radial basis function neural network. More recently, a considerably complex residential HVAC system was studied, and grey-box models were developed for each subsystem; nevertheless, simplistic linear approximation were adopted for the discontinuous low-level on/off controllers [Afram and Janabi-Sharifi, 2017]. A much higher level problem was tackled by [Nghiem and Jones, 2017] in the context of DR, where Gaussian process models for buildings whose BAS implement arbitrary demand adjustment strategies were proposed.

### 4.2 How far are we from the ideal solution?

Comparing each one of the works listed in Section 3.3 and Section 3.4 among themselves would be arduous, especially since they target different HVAC subsystems or even start from a different set of initial assumptions. Moreover, some of the publications do not propose new methods and tools, but rather apply existing ones to the problems at hand. In view of these facts, we instead group only a small but representative subset of them in Table 1, to conduct a more detailed analysis based on the attributes A1–A4 explained in Section 2.4. In the aforementioned table, the acronym MPC is used for any predictive control strategy whose model is not updated on-line, LMPC in case it is, and N/A means not applicable. The symbol ‘-’ indicates if the method does not possess the associated property, and ‘+’ in case it does; in the A4 column the more ‘+’ symbols are given to a method, the lower its computational complexity.

**The reinforcement learning approach:** Its main advantage is, by definition, the ability to continuously adapt over time (A1). One can moreover argue that the strict passivity property of the building thermodynamics are ‘forgiving’ when bad actions are taken by the learning agent. Nevertheless, to the best of the authors’ knowledge, no work published in the HVAC area as of now guaranteed any type of safety during this learning process (A2), an issue that compromises the applicability of this class of methods given their long training periods. It is thus necessary to explore the existing literature about safe reinforcement learning and its intersection with control theory if designers want to minimize thermal discomfort during the initial phase of RL [Perkins and Barto, 2002]. This could be accomplished by a suitable modification of either the optimization criterion or the exploration process itself. As an example, instead of making use of heuristic exploratory methods that are blind to the risk of actions, the initial learning phase may be guided by the available prior knowledge of the

problem [Garcia and Fernández, 2015]. The challenging task would be then to effectively incorporate the specifics of the HVAC problem into this framework. From the computational complexity viewpoint (A4), both [Dalamagkidis et al., 2007] and [Yang et al., 2015] were given bad scores due to their inherent curse of dimensionality. Time-varying energy prices and exogenous DR signals could be incorporated in the RL reward function for instance, yet this could easily lead to convergence problems of the whole algorithm. The difficulty of dealing with continuous domains should also be highlighted since most of the theory has been developed for discrete states and controls. Apart from windows states and blinds, the majority of the remaining quantities within the building context are continuous.

**Rule-based methods:** They miss the attribute of continuous adaptation (A1, A2) and most of the times are also designed without any formal guarantee of constraint satisfaction. Implementation is extremely straightforward (A4) and can take into account occupancy and weather forecasts (A3) when constructing the rules in the form of decision parameters. If complex objectives are not considered and interpretability of the controller is prioritized over performance or energy consumption minimization, then RBC becomes appealing. This is mainly due to the ease of being modified by operators based on their expert knowledge and experience, justifying why RBC is widely implemented in commercial products. The proposal of extracting rules from complex hybrid MPC controllers found in [Domahidi et al., 2014] and similar papers is regarded as promising, but would require experimental validation in a realistic large-scale building to confirm its sensitivity to modelling errors, noisy measurements and imprecise weather forecast.

**Predictive control techniques:** As opposed to RL, both continuous and discrete control actions can be simultaneously considered in a single MPC formulation – at the expense of additional computation complexity when compared to the case of exclusive continuous control. Constraints such as acceptable temperature envelopes and humidity levels for each zone can be explicitly taken into account in a straightforward fashion. Furthermore, internal and external disturbance forecasts can be easily exploited (A3). As long as the resulting optimization problem is convex, we can state that MPC possesses good scalability properties (A4), especially e.g. in [Tanaskovic et al., 2017] whose technique was translated into a linear program (LP). Clearly this is not the case of [Lindelöf et al., 2015] since the obtained ANN dynamic model is nonlinear. In practice however, non-convex problems can be tackled using heuristic numerical optimization at the risk of not having any convergence guarantees. As highlighted in [Sturzenegger et al., 2016], monetary savings owing to the adoption of MPC may not outbalance the efforts associated with model development and controller commissioning, especially when the number of actuators is small compared to the number of zones. Nevertheless, the growing presence of distributed renewable sources in the electric network is causing substantial changes in the energy market due to their production uncertainty and low inertia. This in turn triggers time-dependent or demand-dependent electricity tariffs and a general rise of energy prices. We believe that the most systematic method to deal with these factors while minimizing operating is MPC rather than the previous alternatives (see [Fabietti et al., 2018b] for an example of a realistic DR thermal control study).

The time-consuming task of creating thermodynamic models is expected to become less problematic in the future with the advance of on-line learning MPC techniques in the field (A1), especially if associated with a quantification of the model uncertainty and guarantees of constraint satisfaction (A2). The few works found in the literature that combine those two properties have reported promising results, but experimental validation on more complex environments featuring multiple zones and TES subsystems for instance are still needed to confirm their suitability for real-time HVAC control. One of the limitations of the LMPC explored by [Aswani et al., 2013] is that newly collected data-points are only used to indirectly improve the cost function, but not the prediction model on which constraints are imposed. On the other hand, [Tanaskovic et al., 2017] does reduce conservativeness on-line, yet restrict its scope to linear systems. Future investigations can address these issues while standing on solid theoretical basis coming from the areas of adaptive linear MPC [Zhang et al., 2017, Lorenzen et al., 2019, Tanaskovic et al., 2019] and safe-learning [Schreiter et al., 2015, Rosolia and Borrelli, 2017, Hertneck et al., 2018].

## 5 Conclusions

The field of building HVAC research has been evolving substantially during the last decades, migrating from the investigation of individual subsystems and devices to the supervisory control of multiple loops and simultaneous regulation of several climate variables. Case

Table 1: A comparison of several representative building control strategies based on A1 (lifetime adaptability), A2 (safety during learning), A3 (incorporation of exogenous signals and complex objectives), A4 (numerical scalability), and validation method.

Reference	Class	A1	A2	A3	A4	Validation	Note
[Dadamagkidis et al., 2007]	RL	+	-	-	-	Simulations	2 independent zones, performance compared with a Fuzzy-PD controller and RBC
[Yang et al., 2015]	RL	+	-	+	-	Simulations	900-m <sup>2</sup> building with radiant floor heating and thermal storage, performance compared with RBC
[Domahidi et al., 2014]	RBC	-	N/A	+	++	Simulations	6 different complete buildings, performance compared with hybrid MPC
[Oldewurtel et al., 2012]	MPC	-	N/A	+	+	Simulations	18 complete buildings equipped with different HVAC systems, performance compared with RBC, MPC and stochastic MPC
[Váña et al., 2014]	MPC	-	N/A	+	+	Experiments	1 complete building with a radiant floor system, performance compared with the existing control scheme
[Fabietti et al., 2018b]	MPC	-	N/A	+	+	Experiments	Single zone with a commercial heat pump and an emulated electrical storage system, complex objective in a DR scenario
[Maasoumy et al., 2014b]	LMPC	+	-	+	+	Experiments	Single zone with a ground-source heat pump, performance compared among MPC, robust MPC and RBC
[Lindelöf et al., 2015]	LMPC	+	-	+	-	Experiments	10 buildings with either radiant floor systems or external radiators, performance compared with a reference controller
[Aswani et al., 2013]	LMPC	+	+	+	+	Experiments	Single zone with an air-conditioning system, performance compared with MPC and ON/OFF controllers
[Tanaskovic et al., 2017]	LMPC	+	+	-	+	Simulations	Single zone with different actuation methods, performance compared with MPC and naive <sup>d</sup> adaptive MPC

<sup>d</sup>Updating the model at every time step, but neglecting the effects of uncertainty.

studies have become increasingly more complete and, thus, also more complex. In this context, data-driven methods have been used mainly for predicting internal and external disturbances acting on the system, and creating models for the dynamics off-line. The later task reveals itself challenging because of the low-quality measurements normally available, the operations-critical nature of the HVAC problem that does not allow for strong excitation during occupied hours, and the effects of unmeasured internal heat gains. Since there are various possible HVAC system configurations and the commercial and residential settings fairly differ, it is unlikely that one single modelling approach would be suitable for all the problem instances.

In view of the A1–A4 attributes defined for the idealistic high-level supervisory control system, we consider the use of learning model predictive control – especially combined with of semi-parametric or non-parametric regression for identification – as a promising alternative that needs to be further explored by researchers. If the associated uncertainty is quantified during the on-line learning process, conservativeness could be reduced over time while ensuring thermal comfort. Finally, we see the necessity of staying in the high-level layer of the control hierarchy, possibly interfacing with existing standard proprietary systems. Combining all the aforementioned aspects in a single technique would certainly not only be interesting from an academic viewpoint, but also an anticipated solution by practitioners.

## Acknowledgment

This work has received support from the Swiss National Science Foundation under the RISK project (Risk Aware Data-Driven Demand Response, grant number 200021 175627).

## Declaration of competing interest

None declared.

## References

- [Afram and Janabi-Sharifi, 2014a] Afram, A. and Janabi-Sharifi, F. (2014a). Review of modeling methods for HVAC systems. *Applied Thermal Engineering*, 67(1-2):507–519.
- [Afram and Janabi-Sharifi, 2014b] Afram, A. and Janabi-Sharifi, F. (2014b). Theory and applications of HVAC control systems—a review of model predictive control (MPC). *Building and Environment*, 72:343–355.
- [Afram and Janabi-Sharifi, 2015] Afram, A. and Janabi-Sharifi, F. (2015). Gray-box modeling and validation of residential hvac system for control system design. *Applied Energy*, 137:134–150.
- [Afram and Janabi-Sharifi, 2017] Afram, A. and Janabi-Sharifi, F. (2017). Supervisory model predictive controller (MPC) for residential HVAC systems: Implementation and experimentation on archetype sustainable house in Toronto. *Energy and Buildings*, 154:268–282.
- [Afram et al., 2017] Afram, A., Janabi-Sharifi, F., Fung, A. S., and Raahemifar, K. (2017). Artificial neural network (ANN) based model predictive control (MPC) and optimization of HVAC systems: A state of the art review and case study of a residential HVAC system. *Energy and Buildings*, 141:96–113.
- [Afroz et al., 2018] Afroz, Z., Shafiullah, G., Urmee, T., and Higgins, G. (2018). Modeling techniques used in building HVAC control systems: A review. *Renewable and Sustainable Energy Reviews*, 83:64–84.
- [Agbi et al., 2012] Agbi, C., Zhen, S., and Bruce, K. (2012). Parameter identifiability for multi-zone building models. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pages 6951–6956.
- [Ai et al., 2014] Ai, B., Fan, Z., and Gao, R. X. (2014). Occupancy estimation for smart buildings by an auto-regressive hidden markov model. In *Proc. of the American Control Conference (ACC)*, pages 2234–2239.

- [Andersen et al., 2014] Andersen, P. D., Jiménez, M. J., Madsen, H., and Rode, C. (2014). Characterization of heat dynamics of an arctic low-energy house with floor heating. *Building Simulation*, 7(6):595–614.
- [Anderson et al., 2008] Anderson, M., Buehner, M., Young, P., Hittle, D., Anderson, C., Tu, J., and Hodgson, D. (2008). MIMO robust control for HVAC systems. *IEEE Transactions on Control Systems Technology*, 16(3):475–483.
- [ANSI/ASHRAE, 2010] ANSI/ASHRAE (2010). Standard 55-2010 – Thermal environmental conditions for human occupancy. Standard, American Society of Heating, Refrigerating and Air-Conditioning Engineers, Atlanta, GA.
- [Arguello-Serrano and Velez-Reyes, 1999] Arguello-Serrano, B. and Velez-Reyes, M. (1999). Nonlinear control of a heating, ventilating, and air conditioning system with thermal load estimation. *IEEE Transactions on Control Systems Technology*, 7(1):56–63.
- [Asad et al., 2017] Asad, H. S., Yuen, R. K. K., and Huang, G. (2017). Multiplexed real-time optimization of HVAC systems with enhanced control stability. *Applied energy*, 187:640–651.
- [Aswani et al., 2013] Aswani, A., Gonzalez, H., Sastry, S. S., and Tomlin, C. (2013). Provably safe and robust learning-based model predictive control. *Automatica*, 49(5):1216–1226.
- [Aswani et al., 2012a] Aswani, A., Master, N., Taneja, J., Culler, D., and Tomlin, C. (2012a). Reducing transient and steady state electricity consumption in HVAC using learning-based model-predictive control. *Proceedings of the IEEE*, 100(1):240–253.
- [Aswani et al., 2012b] Aswani, A., Master, N., Taneja, J., Smith, V., Krioukov, A., Culler, D., and Tomlin, C. (2012b). Identifying models of HVAC systems using semiparametric regression. In *Proc. of the American Control Conference (ACC)*, pages 3675–3680.
- [Atthajariyakul and Leephakpreeda, 2005] Atthajariyakul, S. and Leephakpreeda, T. (2005). Neural computing thermal comfort index for HVAC systems. *Energy Conversion and Management*, 46(15-16):2553–2565.
- [Aynur et al., 2009] Aynur, T. N., Hwang, Y., and Radermacher, R. (2009). Simulation comparison of VAV and VRF air conditioning systems in an existing building for the cooling season. *Energy and Buildings*, 41(11):1143–1150.
- [Bacher et al., 2013] Bacher, P., Madsen, H., Nielsen, H. A., and Perers, B. (2013). Short-term heat load forecasting for single family houses. *Energy and Buildings*, 65:101–112.
- [Bai et al., 2008] Bai, J., Wang, S., and Zhang, X. (2008). Development of an adaptive Smith predictor-based self-tuning PI controller for an HVAC system in a test room. *Energy and Buildings*, 40(12):2244–2252.
- [Balaji et al., 2013] Balaji, B., Xu, J., Nwokafor, A., Gupta, R., and Agarwal, Y. (2013). Sentinel: Occupancy based HVAC actuation using existing WiFi infrastructure within commercial buildings. In *Proc. of the ACM Conference on Embedded Networked Sensor Systems*, pages 2234–2239.
- [Balocco et al., 2014] Balocco, C., Petrone, G., Cammarata, G., Vitali, P., Albertini, R., and Pasquarella, C. (2014). Indoor air quality in a real operating theatre under effective use conditions. *Journal of Biomedical Science and Engineering*, 7(11):866.
- [Barthelmes et al., 2017] Barthelmes, V. M., Heo, Y., Fabi, V., and Corgnati, S. P. (2017). Exploration of the bayesian network framework for modelling window control behaviour. *Building and Environment*, 126:318–330.
- [Beghi et al., 2016] Beghi, A., Brignoli, R., Cecchinato, L., Menegazzo, G., Rampazzo, M., and Simmini, F. (2016). Data-driven fault detection and diagnosis for HVAC water chillers. *Control Engineering Practice*, 53:79–91.
- [Beil et al., 2016] Beil, I., Hiskens, I., and Backhaus, S. (2016). Frequency regulation from commercial building HVAC demand response. *Proceedings of the IEEE*, 104(4):745–757.

- [Ben-Nakhi and Mahmoud, 2017] Ben-Nakhi, A. E. and Mahmoud, M. (2017). Application of building-dynamics-based control strategies to improve air-conditioning performance in educational buildings. *Advances in Building Energy Research*, 11(2):153–179.
- [Cai et al., 2019] Cai, M., Pipattanasomporn, M., and Rahman, S. (2019). Day-ahead building-level load forecasts using deep learning vs. traditional time-series techniques. *Applied Energy*, 236:1078–1088.
- [Candanedo and Feldheim, 2016] Candanedo, L. M. and Feldheim, V. (2016). Accurate occupancy detection of an office room from light, temperature, humidity and co2 measurements using statistical learning models. *Energy and Buildings*, 112:28–39.
- [Capozzoli et al., 2017] Capozzoli, A., Piscitelli, M. S., Gorrino, A., Ballarini, I., and Corrado, V. (2017). CHECK data analytics for occupancy pattern learning to reduce the energy consumption of HVAC systems in office buildings. *Sustainable Cities and Society*, 35:191–208.
- [Chen et al., 1990] Chen, Y., Lee, K., and Wepfer, W. (1990). Adaptive robust control scheme applied to a single-zone HVAC system. In *Proc. of the American Control Conference (ACC)*, pages 1076–1081.
- [Chen et al., 2018] Chen, Y., Norford, L. K., Samuelson, H. W., and Malkawi, A. (2018). Optimal control of HVAC and window systems for natural ventilation through reinforcement learning. *Energy and Buildings*, 169:195–205.
- [Chen et al., 2017] Chen, Z., Zhao, R., Zhu, Q., Masood, M. K., Soh, Y. C., and Mao, K. (2017). Building occupancy estimation with environmental sensors via CDBLSTM. *IEEE Transactions on Industrial Electronics*, 64(12):9549–9559.
- [Chou and Bui, 2014] Chou, J.-S. and Bui, D.-K. (2014). Modeling heating and cooling loads by artificial intelligence for energy-efficient building design. *Energy and Buildings*, 82:437–446.
- [Chua et al., 2013] Chua, K., Chou, S., Yang, W., and Yan, J. (2013). Achieving better energy-efficient air conditioning—a review of technologies and strategies. *Applied Energy*, 104:87–104.
- [Cigler et al., 2013] Cigler, J., Siroky, J., Korda, M., and Jones, C. (2013). On the selection of the most appropriate MPC problem formulation for buildings. In *Proc. of the REHVA World Congress CLIMA*, pages 1–11.
- [Claessens et al., 2018] Claessens, B. J., Vanhoudt, D., Desmedt, J., and Ruelens, F. (2018). Model-free control of thermostatically controlled loads connected to a district heating network. *Energy and Buildings*, 159:1–10.
- [Dalamagkidis et al., 2007] Dalamagkidis, K., Kolokotsa, D., Kalaitzakis, K., and Stavrakakis, G. S. (2007). Reinforcement learning for energy conservation and comfort in buildings. *Building and environment*, 42(7):2686–2698.
- [Darivianakis et al., 2017] Darivianakis, G., Georghiou, A., Smith, R. S., and Lygeros, J. (2017). The power of diversity: data-driven robust predictive control for energy-efficient buildings and districts. *IEEE Transactions on Control Systems Technology*, (99):1–14.
- [de Gracia et al., 2015] de Gracia, A., Fernández, C., Castell, A., Mateu, C., and Cabeza, L. F. (2015). Control of a PCM ventilated facade using reinforcement learning techniques. *Energy and Buildings*, 106:234–242.
- [Djukanovic et al., 1993] Djukanovic, M., Babic, B., Sobajic, D. J., and Pao, Y.-H. (1993). Unsupervised/supervised learning concept for 24-hour load forecasting. In *IEE Proceedings C (Generation, Transmission and Distribution)*, volume 140, pages 311–318.
- [Domahidi et al., 2014] Domahidi, A., Ullmann, F., Morari, M., and Jones, C. N. (2014). Learning decision rules for energy efficient building control. *Journal of Process Control*, 24(6):763–772.
- [Dong et al., 2018] Dong, J., Winstead, C., Nutaro, J., and Kuruganti, T. (2018). Occupancy-based HVAC control with short-term occupancy prediction algorithms for energy-efficient buildings. *Energies*, 11(9):2427.

- [Drgoňa et al., 2018] Drgoňa, J., Picard, D., Kvasnica, M., and Helsen, L. (2018). Approximate model predictive building control via machine learning. *Applied Energy*, 218:199–216.
- [Fabietti et al., 2018a] Fabietti, L., Gorecki, T. T., Qureshi, F. A., Bitlislioglu, A., Lymperopoulos, I., and Jones, C. N. (2018a). Experimental implementation of frequency regulation services using commercial buildings. *IEEE Transactions on Smart Grid*, 9(3):1657–1666.
- [Fabietti et al., 2018b] Fabietti, L., Qureshi, F. A., Gorecki, T. T., Salzmann, C., and Jones, C. N. (2018b). Multi-time scale coordination of complementary resources for the provision of ancillary services. *Applied Energy*, 229:1164–1180.
- [Falcone et al., 2008] Falcone, P., Borrelli, F., Tseng, H. E., Asgari, J., and Hrovat, D. (2008). Linear time-varying model predictive control and its application to active steering systems: Stability analysis and experimental validation. *International Journal of Robust and Nonlinear Control: IFAC-Affiliated Journal*, 18(8):862–875.
- [Ferrano and Wong, 1990] Ferrano, F. J. and Wong, K. (1990). Prediction of thermal storage loads using a neural network. In *Proc. of the Annual meeting of the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE)*, pages 723–726.
- [Ferreira et al., 2012] Ferreira, P., Ruano, A., Silva, S., and Conceicao, E. (2012). Neural networks based predictive control for thermal comfort and energy savings in public buildings. *Energy and Buildings*, 55:238–251.
- [Flett and Kelly, 2016] Flett, G. and Kelly, N. (2016). An occupant-differentiated, higher-order markov chain method for prediction of domestic occupancy. *Energy and Buildings*, 125:219–230.
- [Fong et al., 2006] Fong, K. F., Hanby, V. I., and Chow, T.-T. (2006). HVAC system optimization for energy management by evolutionary programming. *Energy and Buildings*, 38(3):220–231.
- [Fu, 1970] Fu, K.-S. (1970). Learning control systems—Review and outlook. *IEEE Transactions on Automatic Control*, 15(2):210–221.
- [Garcia and Fernández, 2015] Garcia, J. and Fernández, F. (2015). A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1):1437–1480.
- [Garnier et al., 2014] Garnier, A., Eynard, J., Caussanel, M., and Grieu, S. (2014). Low computational cost technique for predictive management of thermal comfort in non-residential buildings. *Journal of Process Control*, 24(6):750–762.
- [Garnier et al., 2015] Garnier, A., Eynard, J., Caussanel, M., and Grieu, S. (2015). Predictive control of multizone heating, ventilation and air-conditioning systems in non-residential buildings. *Applied Soft Computing*, 37:847–862.
- [Gaudio et al., 2019] Gaudio, J. E., Gibson, T. E., Annaswamy, A. M., Bolender, M. A., and Lavretsky, E. (2019). Connections between adaptive control and optimization in machine learning. *arXiv preprint arXiv:1904.05856*.
- [Gils, 2014] Gils, H. C. (2014). Assessment of the theoretical demand response potential in Europe. *Energy*, 67:1–18.
- [Gondhalekar et al., 2013] Gondhalekar, R., Oldewurtel, F., and Jones, C. N. (2013). Least-restrictive robust periodic model predictive control applied to room temperature regulation. *Automatica*, 49(9):2760–2766.
- [Gorecki et al., 2017] Gorecki, T. T., Fabietti, L., Qureshi, F. A., and Jones, C. N. (2017). Experimental demonstration of buildings providing frequency regulation services in the Swiss market. *Energy and Buildings*, 144:229–240.
- [Gupta and Kar, 2018] Gupta, S. K. and Kar, K. (2018). Human-in-the-loop thermal management for smart buildings. In *Intelligent Building Control Systems*, pages 191–217. Springer.
- [Haldi, 2010] Haldi, F. (2010). *Towards a Unified Model of Occupants’ Behaviour and Comfort for Building Energy Simulation*. PhD thesis, École Polytechnique Fédérale de Lausanne.

- [Haldi and Robinson, 2009] Haldi, F. and Robinson, D. (2009). Interactions with window openings by office occupants. *Building and Environment*, 44(12):2378–2395.
- [He et al., 2005] He, M., Cai, W.-J., and Li, S.-Y. (2005). Multiple fuzzy model-based temperature predictive control for HVAC systems. *Information sciences*, 169(1-2):155–174.
- [Henze, 2013] Henze, G. P. (2013). Model predictive control for buildings: a quantum leap? *Journal of Building Performance Simulation*, 6:157–158.
- [Henze et al., 2004] Henze, G. P., Felsmann, C., and Knabe, G. (2004). Evaluation of optimal control for active and passive building thermal storage. *International Journal of Thermal Sciences*, 43(2):173–183.
- [Henze and Schoenmann, 2003] Henze, G. P. and Schoenmann, J. (2003). Evaluation of reinforcement learning control for thermal energy storage systems. *HVAC&R Research*, 9(3):259–275.
- [Hertneck et al., 2018] Hertneck, M., Köhler, J., Trimpe, S., and Allgöwer, F. (2018). Learning an approximate model predictive controller with guarantees. *IEEE Control Systems Letters*, 2(3):543–548.
- [Hsieh et al., 2017] Hsieh, S., Omu, A., and Orehounig, K. (2017). Comparison of solar thermal systems with storage: From building to neighbourhood scale. *Energy and Buildings*, 152:359–372.
- [Huaguang and Cai, 2002] Huaguang, Z. and Cai, L. (2002). Decentralized nonlinear adaptive control of an HVAC system. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 32(4):493–498.
- [Huang et al., 2015] Huang, H., Chen, L., and Hu, E. (2015). A neural network-based multi-zone modelling approach for predictive control system design in commercial buildings. *Energy and Buildings*, 97:86–97.
- [ISO, 2005] ISO (2005). ISO 7730 – ergonomics of the thermal environment – Analytical determination and interpretation of thermal comfort using calculation of the PMV and PPD indices and local thermal comfort criteria. Standard, International Organization for Standardization, Geneva, GE.
- [Ji et al., 2015] Ji, Y., Xu, P., and Ye, Y. (2015). HVAC terminal hourly end-use disaggregation in commercial buildings with fourier series model. *Energy and Buildings*, 97:33–46.
- [Jian and Wenjian, 2000] Jian, W. and Wenjian, C. (2000). Development of an adaptive neuro-fuzzy method for supply air pressure control in HVAC system. In *Proc. of the International Conference on Systems, Man and Cybernetics*, volume 5, pages 3806–3809.
- [Johansson et al., 2017] Johansson, C., Bergkvist, M., Geysen, D., De Somer, O., Lavesson, N., and Vanhoudt, D. (2017). Operational demand forecasting in district heating systems using ensembles of online machine learning algorithms. *Energy Procedia*, 116:208–216.
- [Johnson Controls, 2010] Johnson Controls (2010). Hot water system - Boiler controller application note.
- [Kalogirou, 2009] Kalogirou, S. A. (2009). Artificial neural networks and genetic algorithms in energy applications in buildings. *Advances in Building Energy Research*, 3:83–120.
- [Katipamula and Brambley, 2005a] Katipamula, S. and Brambley, M. R. (2005a). Methods for fault detection, diagnostics, and prognostics for building systems—A review, part I. *HVAC&R Research*, 11(1):3–25.
- [Katipamula and Brambley, 2005b] Katipamula, S. and Brambley, M. R. (2005b). Methods for fault detection, diagnostics, and prognostics for building systems—A review, part II. *HVAC&R Research*, 11(2):169–187.
- [Kazmi et al., 2019] Kazmi, H., Suykens, J., Balint, A., and Driesen, J. (2019). Multi-agent reinforcement learning for modeling and control of thermostatically controlled loads. *Applied Energy*, 238:1022–1035.

- [Killian and Kozek, 2016] Killian, M. and Kozek, M. (2016). Ten questions concerning model predictive control for energy efficient buildings. *Building and Environment*, 105:403–412.
- [Kusiak et al., 2014] Kusiak, A., Xu, G., and Zhang, Z. (2014). Minimization of energy consumption in HVAC systems with data-driven models and an interior-point method. *Energy conversion and management*, 85:146–153.
- [Lam et al., 2010] Lam, J. C., Wan, K. K., Liu, D., and Tsang, C. (2010). Multiple regression models for energy use in air-conditioned office buildings in different climates. *Energy Conversion and Management*, 51(12):2692–2697.
- [Lautenschlager and Lichtenberg, 2016] Lautenschlager, B. and Lichtenberg, G. (2016). Rever data-driven iterative learning for model predictive control of heating systems. *IFAC Proceedings Volumes*, 1(1):1–1.
- [Le et al., 2014] Le, K., Bourdais, R., and Guéguen, H. (2014). From hybrid model predictive control to logical control for shading system: A support vector machine approach. *Energy and Buildings*, 84:352–359.
- [Lecamwasam et al., 2012] Lecamwasam, L., Wilson, J., and Chokolich, D. (2012). *Guide to best practice maintenance & operation of HVAC systems for energy efficiency*. Department of Climate Change and Energy Efficiency.
- [Lee and Malkawi, 2014] Lee, Y. S. and Malkawi, A. M. (2014). Simulating multiple occupant behaviors in buildings: An agent-based modeling approach. *Energy and Buildings*, 69:407–416.
- [Li and Dong, 2018] Li, Z. and Dong, B. (2018). Short term predictions of occupancy in commercial buildings—Performance analysis for stochastic models and machine learning approaches. *Energy and Buildings*, 158:268–281.
- [Liao and Barooah, 2010] Liao, C. and Barooah, P. (2010). An integrated approach to occupancy modeling and estimation in commercial buildings. In *Proc. of the American Control Conference (ACC)*, pages 3130–3135.
- [Liao et al., 2012] Liao, C., Lin, Y., and Barooah, P. (2012). Agent-based and graphical modelling of building occupancy. *Journal of Building Performance Simulation*, 5(1):5–25.
- [Lindelöf et al., 2015] Lindelöf, D., Afshari, H., Alisafaei, M., Biswas, J., Caban, M., Moccellini, X., and Viaene, J. (2015). Field tests of an adaptive, model-predictive heating controller for residential buildings. *Energy and Buildings*, 99:292–302.
- [Lorenzen et al., 2019] Lorenzen, M., Cannon, M., and Allgöwer, F. (2019). Robust MPC with recursive model update. *Automatica*, 103:461–471.
- [Lu et al., 2005] Lu, L., Cai, W., Chai, Y. S., and Xie, L. (2005). Global optimization for overall HVAC systems—part I problem formulation and analysis. *Energy conversion and management*, 46(7-8):999–1014.
- [Luc et al., 2018] Luc, K. M., Heller, A., and Rode, C. (2018). Energy demand flexibility in buildings and district heating systems—a literature review. *Advances in Building Energy Research*, pages 1–23.
- [Lund et al., 2014] Lund, H., Werner, S., Wiltshire, R., Svendsen, S., Thorsen, J. E., Hvelplund, F., and Mathiesen, B. V. (2014). 4th generation district heating (4GDH): Integrating smart thermal grids into future sustainable energy systems. *Energy*, 68:1–11.
- [Ma et al., 2012a] Ma, Y., Borrelli, F., Hancey, B., Coffey, B., Benghea, S., and Haves, P. (2012a). Model predictive control for the operation of building cooling systems. *IEEE Transactions on Control Systems Technology*, 20(3):796–803.
- [Ma et al., 2012b] Ma, Y., Kelman, A., Daly, A., and Borrelli, F. (2012b). Predictive control for energy efficient buildings with thermal storage: Modeling, stimulation, and experiments. *IEEE Control Systems Magazine*, 32(1):44–64.
- [Ma et al., 2015] Ma, Y., Matuško, J., and Borrelli, F. (2015). Stochastic model predictive control for building HVAC systems: Complexity and conservatism. *IEEE Transactions on Control Systems Technology*, 23(1):101–116.

- [Maasoumy et al., 2014a] Maasoumy, M., Moridian, B., Razmara, M., Shahbakhti, M., and Sangiovanni-Vincentelli, A. (2014a). Online simultaneous state estimation and parameter adaptation for building predictive control. In *Proc. ASME Dynamic System and Control Conference (DSCC)*, pages 1–10.
- [Maasoumy et al., 2014b] Maasoumy, M., Razmara, M., Shahbakhti, M., and Vincentelli, A. S. (2014b). Handling model uncertainty in model predictive control for energy efficient buildings. *Energy and Buildings*, 77:377–392.
- [Maasoumy et al., 2014c] Maasoumy, M., Rosenberg, C., Sangiovanni-Vincentelli, A., and Callaway, D. S. (2014c). Model predictive control approach to online computation of demand-side flexibility of commercial buildings HVAC systems for supply following. In *Proc. of the American Control Conference (ACC)*, pages 1082–1089.
- [Mahdavi and Tahmasebi, 2015] Mahdavi, A. and Tahmasebi, F. (2015). Predicting people’s presence in buildings: An empirically based model performance analysis. *Energy and Buildings*, 86:349–355.
- [Mantovani and Ferrarini, 2015] Mantovani, G. and Ferrarini, L. (2015). Temperature control of a commercial building with model predictive control techniques. *IEEE Transactions on Industrial Electronics*, 62(4):2651–2660.
- [May-Ostendorp et al., 2011] May-Ostendorp, P., Henze, G. P., Corbin, C. D., Rajagopalan, B., and Felsmann, C. (2011). Model-predictive control of mixed-mode buildings with rule extraction. *Building and Environment*, 46(2):428–437.
- [McDowall, 2007] McDowall, R. (2007). *Fundamentals of HVAC systems*. Academic Press.
- [Minakais et al., 2019] Minakais, M., Mishra, S., and Wen, J. T. (2019). Database-driven iterative learning for building temperature control. *IEEE Transactions on Automation Science and Engineering*, 1(1):1–11.
- [Minakais et al., 2017] Minakais, M., Okaeme, C. C., Mishra, S., and Wen, J. T. (2017). Iterative learning control for coupled temperature and humidity in buildings. *IFAC Proceedings Volumes*, 50(1):13420–13425.
- [Moon, 2012] Moon, J. W. (2012). Performance of ANN-based predictive and adaptive thermal-control methods for disturbances in and around residential buildings. *Building and Environment*, 48:15–26.
- [Moon and Kim, 2010] Moon, J. W. and Kim, J.-J. (2010). ANN-based thermal control models for residential buildings. *Building and Environment*, 45(7):1612–1625.
- [Moon et al., 2013] Moon, J. W., Yoon, S.-H., and Kim, S. (2013). Development of an artificial neural network model based thermal control logic for double skin envelopes in winter. *Building and Environment*, 61:149–159.
- [Morel et al., 2001] Morel, N., Bauer, M., El-Khoury, M., and Krauss, J. (2001). Neurobat, a predictive and adaptive heating control system using artificial neural networks. *International Journal of Solar Energy*, 21(2-3):161–201.
- [Motegi et al., 2007] Motegi, N., Piette, M. A., Watson, D. S., Kiliccote, S., and Xu, P. (2007). Introduction to commercial building control strategies and techniques for demand response. Technical report, Demand Response Research Center, Berkeley, CA.
- [Mukherjee et al., 2012] Mukherjee, S., Mishra, S., and Wen, J. T. (2012). Building temperature control: A passivity-based approach. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pages 6902–6907.
- [Namburu et al., 2007] Namburu, S. M., Azam, M. S., Luo, J., Choi, K., and Pattipati, K. R. (2007). Data-driven modeling, fault diagnosis and optimal sensor selection for HVAC chillers. *IEEE Transactions on Automation Science and Engineering*, 4(3):469–473.
- [Nesler, 1985] Nesler, C. G. (1985). Experiences in applying adaptive control to thermal processes in buildings. In *Proc. of the American Control Conference (ACC)*, pages 1535–1540.

- [Nghiem and Jones, 2017] Nghiem, T. X. and Jones, C. N. (2017). Data-driven demand response modeling and control of buildings with gaussian processes. In *Proc. of the American Control Conference (ACC)*, pages 2919–2924.
- [Obando et al., 2014] Obando, G., Pantoja, A., and Quijano, N. (2014). Building temperature control based on population dynamics. *IEEE Transactions on Control Systems Technology*, 22(1):404–412.
- [O’Connor, 2000] O’Connor, D. (2000). A process control primer. Technical report, Honeywell, Freeport, IL.
- [Oldewurtel et al., 2012] Oldewurtel, F., Parisio, A., Jones, C. N., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B., and Morari, M. (2012). Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings*, 45:15–27.
- [Oldewurtel et al., 2013] Oldewurtel, F., Sturzenegger, D., and Morari, M. (2013). Importance of occupancy information for building climate control. *Applied energy*, 101:521–532.
- [Østergård et al., 2018] Østergård, T., Jensen, R. L., and Maagaard, S. E. (2018). A comparison of six metamodeling techniques applied to building performance simulations. *Applied Energy*, 211:89–103.
- [Page et al., 2008] Page, J., Robinson, D., Morel, N., and Scartezzini, J.-L. (2008). A generalised stochastic model for the simulation of occupant presence. *Energy and Buildings*, 40(2):83–98.
- [Pallonetto et al., 2019] Pallonetto, F., De Rosa, M., Milano, F., and Finn, D. P. (2019). Demand response algorithms for smart-grid ready residential buildings using machine learning models. *Applied Energy*, 239:1265–1282.
- [Park et al., 2019] Park, J. Y., Dougherty, T., Fritz, H., and Nagy, Z. (2019). LightLearn: An adaptive and occupant centered controller for lighting based on reinforcement learning. *Building and Environment*, 147:397–414.
- [Park et al., 2018] Park, S., Ahn, K. U., Hwang, S., Choi, S., and Park, C. S. (2018). Machine learning vs. hybrid machine learning model for optimal operation of a chiller. *Science and Technology for the Built Environment*, pages 1–12.
- [Pčolka et al., 2016] Pčolka, M., Žáčeková, E., Robinett, R., Čelikovský, S., and Šebek, M. (2016). Bridging the gap between the linear and nonlinear predictive control: Adaptations for efficient building climate control. *Control Engineering Practice*, 53:124–138.
- [Peng et al., 2018] Peng, Y., Rysanek, A., Nagy, Z., and Schlüter, A. (2018). Using machine learning techniques for occupancy-prediction-based cooling control in office buildings. *Applied energy*, 211:1343–1358.
- [Penya et al., 2011] Penya, Y. K., Borges, C. E., Agote, D., and Fernández, I. (2011). Short-term load forecasting in air-conditioned non-residential buildings. In *Proc. of the International Symposium on Industrial Electronics (ISIE)*, pages 1359–1364.
- [Perkins and Barto, 2002] Perkins, T. J. and Barto, A. G. (2002). Lyapunov design for safe reinforcement learning. *Journal of Machine Learning Research*, 3(Dec):803–832.
- [Platt et al., 2010] Platt, G., Li, J., Li, R., Poulton, G., James, G., and Wall, J. (2010). Adaptive HVAC zone modeling for sustainable buildings. *Energy and Buildings*, 42(4):412–421.
- [Privara et al., 2013] Privara, S., Cigler, J., Váňa, Z., Oldewurtel, F., Sagerschnig, C., and Žáčeková, E. (2013). Building modeling as a crucial part for building predictive control. *Energy and Buildings*, 56:8–22.
- [Qureshi and Jones, 2018] Qureshi, F. A. and Jones, C. N. (2018). Hierarchical control of building HVAC system for ancillary services provision. *Energy and Buildings*, 169:216–227.
- [Radecki and Hencsey, 2013] Radecki, P. and Hencsey, B. (2013). Online thermal estimation, control, and self-excitation of buildings. In *52nd IEEE Conference on Decision and Control*, pages 4802–4807.

- [Rahman and Smith, 2018] Rahman, A. and Smith, A. D. (2018). Predicting heating demand and sizing a stratified thermal storage tank using deep learning algorithms. *Applied Energy*, 228:108–121.
- [Rahman et al., 2018] Rahman, I., Kuzlu, M., and Rahman, S. (2018). Power disaggregation of combined HVAC loads using supervised machine learning algorithms. *Energy and Buildings*, 172:57–66.
- [Rathod et al., 2019] Rathod, N., La Bella, A., Puleo, G., Scattolini, R., Rossetti, A., and Sandroni, C. (2019). Modelling and predictive control of a solar cooling plant with flexible configuration. *Journal of Process Control*, 76:74–86.
- [Reijula et al., 2013] Reijula, J., Holopainen, R., Kähkönen, E., Reijula, K., and Tommelein, I. D. (2013). Intelligent HVAC systems in hospitals. *Intelligent Buildings International*, 5(2):101–119.
- [Roetzel et al., 2010] Roetzel, A., Tsangrassoulis, A., Dietrich, U., and Busching, S. (2010). A review of occupant control on natural ventilation. *Renewable and Sustainable Energy Reviews*, 14(3):1001–1013.
- [Rosolia and Borrelli, 2017] Rosolia, U. and Borrelli, F. (2017). Learning model predictive control for iterative tasks. a data-driven control framework. *IEEE Transactions on Automatic Control*, 63(7):1883–1896.
- [Royapoor et al., 2018] Royapoor, M., Antony, A., and Roskilly, T. (2018). A review of building climate and plant controls, and a survey of industry perspectives. *Energy and Buildings*, 158:453–465.
- [Salsbury, 2002] Salsbury, T. I. (2002). A new pulse modulation adaptive controller (PMAC) applied to HVAC systems. *Control Engineering Practice*, 10(12):1357–1370.
- [Salsbury, 2005] Salsbury, T. I. (2005). A survey of control technologies in the building automation industry. *IFAC Proceedings Volumes*, 38(1):90–100.
- [Schreiter et al., 2015] Schreiter, J., Nguyen-Tuong, D., Eberts, M., Bischoff, B., Markert, H., and Toussaint, M. (2015). Safe exploration for active learning with gaussian processes. In *Joint European conference on machine learning and knowledge discovery in databases*, pages 133–149.
- [Schweiker et al., 2012] Schweiker, M., Haldi, F., Shukuya, M., and Robinson, D. (2012). Verification of stochastic models of window opening behaviour for residential buildings. *Journal of Building Performance Simulation*, 5:55–74.
- [Serale et al., 2018] Serale, G., Fiorentini, M., Capozzoli, A., Bernardini, D., and Bemporad, A. (2018). Model predictive control (MPC) for enhancing building and HVAC system energy efficiency: Problem formulation, applications and opportunities. *Energies*, 11(3):631.
- [Shahnazari et al., 2019] Shahnazari, H., Mhaskar, P., House, J. M., and Salsbury, T. I. (2019). Modeling and fault diagnosis design for HVAC systems using recurrent neural networks. *Computers & Chemical Engineering*.
- [Shin et al., 2002] Shin, Y., Chang, Y. S., and Kim, Y. (2002). Controller design for a real-time air handling unit. *Control Engineering Practice*, 10(5):511–518.
- [Siemens, 2014] Siemens (2014). PID control with PID\_compact. Available at [www.infopl.net/files/descargas/siemens/infoPLC\\_net\\_100746401\\_S7-1200\\_PID\\_Compact\\_DOKU\\_v1d0\\_en.pdf](http://www.infopl.net/files/descargas/siemens/infoPLC_net_100746401_S7-1200_PID_Compact_DOKU_v1d0_en.pdf).
- [Siemens, 2018] Siemens (2018). Smart thermostat – RDS 110. Available at [www.downloads.siemens.com/download-center/Download.aspx?pos=download&fct=getasset&id1=A6V10807602](http://www.downloads.siemens.com/download-center/Download.aspx?pos=download&fct=getasset&id1=A6V10807602).
- [Singh et al., 2006] Singh, J., Singh, N., and Sharma, J. (2006). Fuzzy modeling and control of HVAC systems—A review. *Journal of Scientific & Industrial Research*, 65(6):470–477.
- [Široký et al., 2011] Široký, J., Oldewurtel, F., Cigler, J., and Prívara, S. (2011). Experimental analysis of model predictive control for an energy efficient building heating system. *Applied Energy*, 88(9):3079–3087.

- [Sousa et al., 1997] Sousa, J., Babuška, R., and Verbruggen, H. (1997). Fuzzy predictive control applied to an air-conditioning system. *Control engineering practice*, 5(10):1395–1406.
- [Soyguder et al., 2009] Soyguder, S., Karakose, M., and Alli, H. (2009). Design and simulation of self-tuning PID-type fuzzy adaptive control for an expert HVAC system. *Expert systems with applications*, 36(3):4566–4573.
- [Stabat and Marchio, 2009] Stabat, P. and Marchio, D. (2009). Opportunities for reversible chillers in office buildings in Europe. *Building Simulation*, 2(2):95–108.
- [Stluka et al., 2018] Stluka, P., Parthasarathy, G., Gabel, S., and Samad, T. (2018). Architectures and algorithms for building automation—An industry view. In *Intelligent Building Control Systems*, pages 11–43. Springer.
- [Stromback, 2017] Stromback, J. (2017). Explicit demand response in Europe – Mapping the markets. Technical report, Smart Energy Demand Coalition.
- [Sturzenegger et al., 2016] Sturzenegger, D., Gyalistras, D., Morari, M., and Smith, R. S. (2016). Model predictive climate control of a swiss office building: Implementation, results, and cost-benefit analysis. *IEEE Transactions on Control Systems Technology*, 24(1):1–12.
- [Sutton et al., 1992] Sutton, R. S., Barto, A. G., and Williams, R. J. (1992). Reinforcement learning is direct adaptive optimal control. *IEEE Control Systems Magazine*, 12(2):19–22.
- [Systemair, 2011] Systemair (2011). Air handling units – NB units (normabloc). Available at [systemair.com/globalassets/downloads/leaflets-and-catalogues/english/nb\\_ahu.pdf](http://systemair.com/globalassets/downloads/leaflets-and-catalogues/english/nb_ahu.pdf).
- [Tanaskovic et al., 2019] Tanaskovic, M., Fagiano, L., and Gligorovski, V. (2019). Adaptive model predictive control for linear time varying MIMO systems. *Automatica*, 105:237–245.
- [Tanaskovic et al., 2017] Tanaskovic, M., Sturzenegger, D., Smith, R., and Morari, M. (2017). Robust adaptive model predictive building climate control. *IFAC-Papersonline*, 50(1):1871–1876.
- [Vakiloroyaya et al., 2014] Vakiloroyaya, V., Samali, B., Fakhar, A., and Pishghadam, K. (2014). A review of different strategies for HVAC energy saving. *Energy Conversion and Management*, 77:738–754.
- [Vána et al., 2014] Vána, Z., Cigler, J., Široký, J., Žáčeková, E., and Ferkl, L. (2014). Model-based energy efficient control applied to an office building. *Journal of Process Control*, 24(6):790–797.
- [Wallace et al., 2012] Wallace, M., Das, B., Mhaskar, P., House, J., and Salsbury, T. (2012). Offset-free model predictive control of a vapor compression cycle. *Journal of Process Control*, 22(7):1374–1386.
- [Wallace et al., 2015] Wallace, M., Mhaskar, P., House, J., and Salsbury, T. I. (2015). Offset-free model predictive control of a heat pump. *Industrial & Engineering Chemistry Research*, 54(3):994–1005.
- [Wang and Ma, 2008] Wang, S. and Ma, Z. (2008). Supervisory and optimal control of building HVAC systems: A review. *HVAC&R Research*, 14(1):3–32.
- [Wang and Hu, 2018] Wang, Z. and Hu, G. (2018). Economic MPC of nonlinear systems with nonmonotonic lyapunov functions and its application to HVAC control. *International Journal of Robust and Nonlinear Control*, 28(6):2513–2527.
- [Wei et al., 2017] Wei, T., Wang, Y., and Zhu, Q. (2017). Deep reinforcement learning for building HVAC control. In *Proceedings of the 54th Annual Design Automation Conference 2017*, page 22.
- [Wei et al., 2018] Wei, Y., Zhang, X., Shi, Y., Xia, L., Pan, S., Wu, J., Han, M., and Zhao, X. (2018). A review of data-driven approaches for prediction and classification of building energy consumption. *Renewable and Sustainable Energy Reviews*, 82:1027–1047.
- [Werner, 2017] Werner, S. (2017). International review of district heating and cooling. *Energy*, 137:617–631.

- [Whalley and Abdul-Ameer, 2011] Whalley, R. and Abdul-Ameer, A. (2011). Heating, ventilation and air conditioning system modelling. *Building and Environment*, 46(3):643–656.
- [Wijayasekara and Manic, 2015] Wijayasekara, D. and Manic, M. (2015). Data-fusion for increasing temporal resolution of building energy management system data. In *Proc. of the Annual Conference of the IEEE Industrial Electronics Society (IECON)*, pages 4550–4555.
- [Xi et al., 2007] Xi, X.-C., Poo, A.-N., and Chou, S.-K. (2007). Support vector regression model predictive control on a HVAC plant. *Control Engineering Practice*, 15(8):897–908.
- [Xuan et al., 2019] Xuan, Z., Xuehui, Z., Liequan, L., Zubing, F., Junwei, Y., and Dongmei, P. (2019). Forecasting performance comparison of two hybrid machine learning models for cooling load of a large-scale commercial building. *Journal of Building Engineering*, 21:64–73.
- [Yan et al., 2012] Yan, C., Wang, S., and Xiao, F. (2012). A simplified energy performance assessment method for existing buildings based on energy bill disaggregation. *Energy and Buildings*, 55:563–574.
- [Yana et al., 2015] Yana, D., O’Brien, W., Hongc, T., Feng, X., Gunay, H. B., Tahmasebi, F., and Mahdavi, A. (2015). CHECK simulating multiple occupant behaviors in buildings: An agent-based modeling approach. *Energy and Buildings*, 107:264–278.
- [Yang et al., 2015] Yang, L., Nagy, Z., Goffin, P., and Schlueter, A. (2015). Reinforcement learning for optimal control of low exergy buildings. *Applied Energy*, 156:577–586.
- [Yoon et al., 2014] Yoon, J. H., Baldick, R., and Novoselac, A. (2014). Dynamic demand response controller based on real-time retail price for residential buildings. *IEEE Transactions on Smart Grid*, 5(1):121–129.
- [Yu and Dexter, 2010] Yu, Z. and Dexter, A. (2010). Online tuning of a supervisory fuzzy controller for low-energy building system using reinforcement learning. *Control Engineering Practice*, 18(5):532–539.
- [Yun et al., 2012] Yun, K., Luck, R., Mago, P. J., and Cho, H. (2012). Building hourly thermal load prediction using an indexed ARX model. *Energy and Buildings*, 54:225–233.
- [Žáčková et al., 2014] Žáčková, E., Váňa, Z., and Cigler, J. (2014). Towards the real-life implementation of MPC for an office building: Identification issues. *Applied Energy*, 135:53–62.
- [Zeng et al., 2019] Zeng, A., Liu, S., and Yu, Y. (2019). Comparative study of data driven methods in building electricity use prediction. *Energy and Buildings*.
- [Zhang et al., 2017] Zhang, X., Kamgarpour, M., Georghiou, A., Goulart, P., and Lygeros, J. (2017). Robust optimal control with adjustable uncertainty sets. *Automatica*, 75:249–259.
- [Zhang et al., 2015] Zhang, Y., O’Neill, Z., Dong, B., and Augenbroe, G. (2015). Comparisons of inverse modeling approaches for predicting building energy performance. *Building and Environment*, 86:177–190.
- [Zhang et al., 2007] Zhang, Y., Zhou, G., Lin, K., Zhang, Q., and Di, H. (2007). Application of latent heat thermal energy storage in buildings: State-of-the-art and outlook. *Building and Environment*, 42(6):2197–2209.
- [Zhao et al., 2016] Zhao, D., Zhong, M., Zhang, X., and Su, X. (2016). Energy consumption predicting model of VRV (variable refrigerant volume) system in office buildings based on data mining. *Energy*, 102:660–668.
- [Zhao and Magoulès, 2012] Zhao, H. and Magoulès, F. (2012). A review on the prediction of building energy consumption. *Renewable and Sustainable Energy Reviews*, 16:3586–3592.
- [Zheng et al., 2017] Zheng, X., Han, Y., Zhang, H., Zheng, W., and Kong, D. (2017). Numerical study on impact of non-heating surface temperature on the heat output of radiant floor heating system. *Energy and Buildings*, 155:198–206.

[Zhou et al., 2008] Zhou, Q., Wang, S., Xu, X., and Xiao, F. (2008). A grey-box model of next-day building thermal load prediction for energy-efficient control. *International Journal of Energy Research*, 32(15):1418–1431.