Machine Learning to improve indoor climate and building energy performance

Christoph Hopmann

c.hopmann@hotmail.com

Instituto Superior Técnico, Universidade de Lisboa, Portugal

October 2017

ABSTRACT

Using model predictive control for energy management systems is becoming more popular. These systems control the building performance based on a forecast of the control based on a building model. Uncertainties in building operation are a major issue in the use of model predictive control, as they decrease prognosis precision. This includes user's varying thermal requirements. The aim of this paper is to develop a comfort temperature predictor for individually controllable offices based on the weather conditions using machine learning. For this purpose, the fundamentals of indoor comfort and machine learning are presented. Recent trends in research on machine learning in buildings are reviewed. The methodology for developing and testing the comfort temperature predictor is explained. Afterwards the results of forecasting the comfort temperature and the energetic impact of the predictor are described.

The results show that supervised learning artificial neural networks and Gaussian Process Regression tools can predict comfort temperatures based on weather conditions with sufficient precision, better than the currently common temperature setpoints. Roughly one year of data is required to reach this performance. Real-time learning with reinforcement learning using artificial neural network value function approximation and sample reuse needs 50 days of learning to reach good precision. Supervised learning can reduce the heating load and reduce overheating, while there is no positive impact on the cooling load with either tool. The paper closes with a review of the tool and an outlook towards its improvements and applications

Key-words: building energy systems, thermal comfort, machine learning, predictive control

1. Introduction

The built environment plays a significant role in modern societies, as humans in western countries spend 90 % of their time indoors [1]. Indoor comfort can have an important impact on someone's well-being, not only short-term while occupying a specific building, but also in the long-run. Indoor comfort is a very individual state. Consequently, a trend towards more individual control options in buildings can be observed. This does entail an increase in control uncertainties, as individual behaviour is very difficult to predict [2,3].

Over the last years machine learning has found its way into a broad field of applications. In buildings, its main application has been in load prediction, control optimization, and occupation and comfort prediction. Commercial implementation however is still scarce [4]. Uncertainties in building operation, including varying thermal requirements by users, are a major issue in the use of model predictive control, as they decrease prognosis precision. The basic concept of using predictions to effectively control HVAC systems gets undermined by spontaneous changes in user requirements. The Predicted Mean Vote (PMV) developed by *Fanger* can be used to predict the comfort demand of groups. It is however not suited to predict individual demands, and some of the data required is difficult to measure [3]. The aim of this paper is to develop a comfort temperature predictor for individually controllable offices based on the environmental weather conditions using machine learning. The feasibility of different supervised and reinforcement learning tools is tested concerning precision and amount of data required, and the impact on the energy consumption is evaluated.

2. Fundamentals

2.1 Indoor Thermal Comfort

Thermal comfort is achieved when a human does not feel the need to change the state of the thermal environment or to adapt to it. This condition is reached when the heat balance of a human is neutral, i.e. the heat produced by metabolism is equal to the heat dissipated to the environment [5]. The perception of the thermal environment is mainly influenced by six primary factors: air temperature, velocity and relative humidity, temperature of surrounding surfaces, clothes and a person's activity. Additionally, there are further physical, physiological and other factors, that influence one's thermal comfort, as depicted in Figure **1** [6]. *Fanger* developed a thermal comfort model based on a heat balance model of the human body. Using the result of climate chamber experiments, he developed an equation that can be used to predict the thermal comfort of a group of people depending on the six main primary factors. The resulting PMV is used in several national and international standards [5–10].

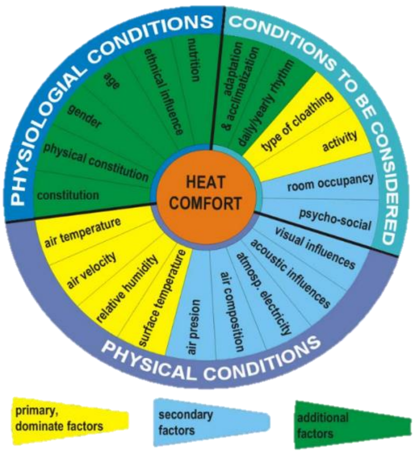

*Figure **1** Primary, secondary and additional factors influencing thermal comfort [6]*

Adaptive comfort is based on an individual's conscious and unconscious adaptation to her thermal environment in order to restore their thermal comfort. As a consequence of adaptation, an individual's comfort temperature may change. The main influencing factor found is the outdoor air temperature, both present and the running mean temperature of a week [5,11–14].

2.2 Machine Learning

The goal of machine learning is to build computer systems that automatically improve with experience[4]. A machine is thus given the ability to learn without being explicitly programmed [15].

Machine learning algorithms can be classified into three main learning types: supervised learning, unsupervised learning and reinforcement learning. There are also hybrids between the learning types, like semi-supervised learning. Figure **2** gives a graphical representation of these types including their typical output data type, the used method and an application example.

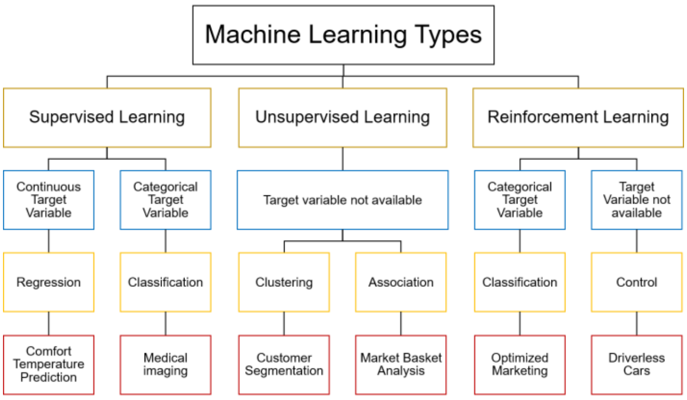

*Figure **2** Types of Machine Learning, adopted from [4]*

In supervised learning, the learning system is fed with inputs and the corresponding desired outputs, the so called labelled data or training data. Depending on the type of label the learning approach differs. For continuous output data, one uses regression, while for categorical outputs, classification is used [4,15–17]. Reinforcement learning evolves through interaction with its environment. Through feedback from the environment the learning machine receives an evaluation about a chosen action. Using a mix of exploiting existing knowledge and exploring the unknown the agent tries to maximize its reward. This process is suitable in situations where no learning data is available or updates occur in short intervals [4,17–19].

3. Literature Review

Previous work related to machine learning in building energy systems shows that the integration of advanced control systems into existing HVAC systems is feasible [20], with all advanced control systems outperforming classic controllers [18,21–25]. However, in some cases, it proved to be difficult to properly estimate the real performance of classic controllers, since only few of them are configured correctly [23,26]. It may well be that less advanced methods are more effective than the most recent tools, especially in cases where correlations are simple [27]. It can be observed that the implementation of the control systems is difficult: a proper implementation requires a high level of expertise and the demand on computational power for real time application is still high, and some of the methods used require extensive sensor networks [22,24,25,28–32]. The learning process is another issue hindering the spread, as the learning periods are still fairly long, making big sets of learning data a necessity. Complete learning under operation is not yet an option, as the impact on the system performance would be too high [18,22,25,28,30,31,33–35].

4. Methodology

Using supervised and reinforcement learning in *Matlab*, a tool to predict the comfort temperature of individuals dependant on weather conditions shall be developed. Data from a German office building is used to train and test the tools. The setpoint predictions provided by the tools are used in a building simulation of a simple office model to evaluate its impact on the energy consumption.

The PMV's main parameters are used as an orientation to develop the comfort temperature predictor. The running mean temperature is assumed as a replacement for the clothing value and the metabolic rate used in the PMV, with a minimum set to 10 °C [5]. Relative humidity measurements from within the room cannot be directly used. As indoor air temperatures usually lie within a limited range of temperatures it is fair to assume a relatively direct link to the outdoor vapour pressure if no humidification is used. In moderate climates, there may also be a direct link between outdoor and indoor relative humidity, which is therefore also considered as an alternative to the vapour pressure. Radiation data is rarely measured. As an approximation, the use of cloudiness, supplied by a nearby public weather station, as an indicator for the incoming irradiation and as an indicator for (perceived) likelihood of rain and the respective choice of clothes is suggested. Alternatively, daily sunshine hours are considered. The air velocity is not implemented. It is assumed that in general air speeds will not be high enough to cause a user discomfort. From intuition, the daily maximum and minimum temperature may also have an influence on the comfort temperature and are thus considered for testing.

All temperature setpoints used have been collected at the central building of the savings bank *Kreissparkasse Göppingen* in Germany (48,7 ° N, 9,6 ° E). Hourly temperature setpoints are available for a total of 976 days. The data derives from 35 rooms with individual temperature control spread over the first to the fourth floor of the building. Weather data collected directly on site is small and only includes the outdoor temperatures and relative humidity. As it does not contain all data deemed necessary the use of onsite data is discarded. The German weather service (DWD) provides historical weather data from weather stations all around Germany through the Climate Data Centre (CDC). The closest station to Göppingen that provides the required data is located at Stuttgart-Echterdingen (48,7 ° N, 9,2 ° E DWD-station ID 4931) [36].

Since complete weather data sets close enough to the building are only available as daily averages, the temperature setpoints are averaged over a day. A day is only considered suitable for use if it includes at least eight hourly temperature setpoints, else the day is excluded from the dataset.

The Mathworks' *Matlab 2017a* contains two toolboxes suited for machine learning applications: the Statistics and Machine Learning toolbox and the Neural Network toolbox. The Machine Learning toolbox contains several tools that enable the set-up of regression learning tools based on input and output parameters. It comes with a total of 19 standard tools which can be divided into linear regression, decision trees, support vector machines, bagged and boosted trees, and Gaussian process regression (GPR). As a first step, all tools are run with standard conditions to estimate their overall feasibility. The Neural Network toolbox enables a quick set-up of artificial neural networks (ANN). Different learning methods can be applied that vary in computational effort and performance. The Levenberg-Marquardt backpropagation is used since it is a good compromise between learning speed and performance. One hidden layer is used, with the number of neurons being varied between nine and 60 in increments of three to find the best, non-overfitting number of neurons. The performance of all tools is analysed using the RMS compared to a testing data sample. To decide on which tools and which weather parameters to focus on, learning is performed for a sample of eight rooms, with two rooms from each floor. The two rooms are chosen randomly among all rooms on the floor. The three best performing regression tools and the three best performing neural networks are chosen according to their median, their lower and upper quartile, and their minima and maxima. They are then trained with the remaining rooms and reanalysed in comparison to a baseline scenario of 21 °C heating setpoint and 23 °C cooling setpoint. Out of the six fully analysed tools the top three are chosen to vary the amount of training data fed into them. For this purpose, the whole data set is still fed into the trainer, but the amount used for training is varied between 5 % and 80 % in steps of 5 %, with the remainder being equally split into use for validation and testing.

Matlab has no standard tools for reinforcement learning so it is necessary to develop a specific algorithm for this purpose. Since a continuous output is required from a set of up to six continuous inputs a lookup table for the value function is deemed impractical. Instead, a function approximation using artificial neural networks is performed. The neural network is based on the results from the supervised learning approach, as it is already optimized in size for the problem at hand. The error between the prediction and the desired setpoint is used as a penalty function and is directly used to improve the value function through gradient descent based error backpropagation for the ANN. All nodes are sigmoids.

A year worth of data, i.e. 365 data points, are used for training, while the remainder is used for testing. A discounted learning rate is tested for reducing overfitting, while the reuse of samples is tested to decrease the amount of data needed for the targeted performance. The discount rate for learning is varied between 0,5 % to 2,5 % and the number of sample reuses is varied between one and ten repetitions. After individual testing, both methods are combined to evaluate their combined performance improvement. Three target values are analysed: Time until RMS below 2 °C, time until RMS below 1,5 °C and minimum RMS.

A model of a single office is used to evaluate the energetic performance of a supervised and a reinforcement learning tool. The exterior constructions are according to the German energy saving regulation (*Energieeinsparverordnung*, EnEv) [37]. The office is 4,0 m deep and 2,5 m wide, with a height of 2,8 m. The window is 2,1 m wide and 1,7 m high. The three internal loads are one occupant, electric equipment and lighting. The occupant performs work at the desk, leading to an activity level of 115 W. The office is equipped with a laptop and two additional screens accounting for a load of 120 W. Lighting is supplied by LEDs, adding up to 14 W. The office is occupied from 8 am to 12 pm and from 1 pm to 6 pm on weekdays and there is no occupation during weekends. The electric equipment and lights are turned on during these periods as well. Air conditioning is modelled using the *EnergyPlus* Ideal Loads Air System template, which supplies conditioned air under idealised conditions. The system’s heating and cooling capacity is limited to 400 W to represent the limitations of state-of-the-art active ceilings used for individual temperature control in offices. The HVAC system is available two hours before and after occupation starts.

5. Results

Table **1** shows all available weather parameters and a respective code number that is used for quick referencing and in graphics. Code number 1 is internally used for the time stamp and thus not used for external coding.

*Table **1** Weather Parameters*

Weather Parameter	Code Number
Average outdoor air temperature	2
Vapour pressure	3
Cloud cover	4
Relative humidity	5
Sunshine hours	6
Maximum outdoor air temperature	7
Minimum outdoor air temperature	8
Running mean air temperature	9

5.1 Neural Networks

The top three combinations gathered from training with eight sample rooms are trained with the remaining 27 rooms. The resulting boxplots of the RMS are shown in Figure **3**, including a boxplot for the reference scenario. All neural networks have a median RMS of 1,12 °C and a maximum of 1,93 °C, 1,99 °C and 1,95 °C from left to right. All three outperform the reference scenario with its median RMS of 2,11 °C and a maximum of 2,75 °C.

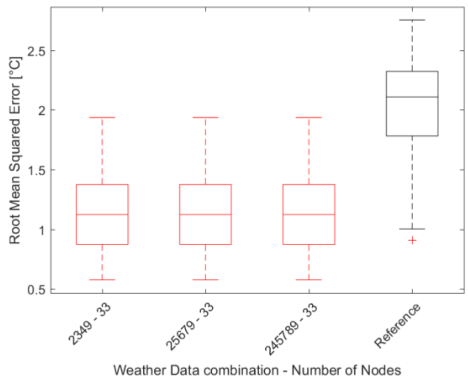

*Figure **3** Neural Network Top 3 Total RMS for all rooms, compared to reference scenario*

5.1.1 Required training data

It is reasonable to perform an analysis of the required amount of data to reach the target value. For the parameter combination 2-4-5-7-8-9 the size of the training set is varied from 5 % to 80 % of the available data set. Figure 4 shows the result.

The median shows a logarithmic decay, with a median RMS of less than 2 °C being achieved with 15 % of the available data being used for training, equalling 131 data points. At roughly 30 %, or 261 data points, the medians go below 1.5 °C. The decay of the medians starts to become small at around 50 % to 60 %. As a conclusion, at 15 % the use of neural networks becomes generally possible, and it starts to become feasible at 40 %, equivalent to 348 data points. It can also be concluded that using an overall larger data set improves the functionality of a neural network to predict indoor temperature setpoints.

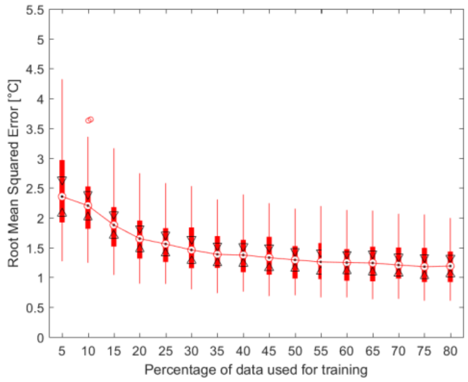

Figure 4 Neural Network 245789 - 33 Total RMS over training data size

5.2 Regression Learners

Figure **5** shows the RMS for three best-performing tool-parameter combinations in comparison to the reference scenario. The median for GPR Exponential (18) is 1,15 °C and the maximum 1,98 °C, for GPR Matern 5/2 (17) the values are 1,16 °C and 2 °C respectively and for GPR Rational Quadratic (19) 1,15 °C and 2 °C respectively. The medians of all tools are below the threshold value of 2 °C and the maximum RMS are roughly equal to the threshold. All analysed tools clearly outperform the reference scenario. Overall, the use of regression learning tools is viable with the amount of data available.

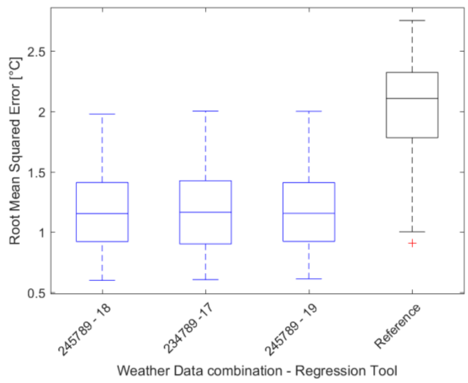

*Figure **5** Regression Tool Top 3 RMS for all rooms, compared to reference scenario*

5.3 Reinforcement Learning

Since the supervised learning tools comply with the conditions of a RMS below 2 °C and outperform the baseline scenario, the performance of reinforcement learning tools is evaluated.

Supervised learning methods are often used for function approximation of the value function in reinforcement learning [38]. ANNs have shown good potential for comfort temperature learning in the supervised learning approach, so that their application in reinforcement learning is only logical. The use of the right error function is essential for a good performance of reinforcement learning tools. The root mean squared error, the mean squared error and a linear error are tested. The RMS-error function outperforms the half-MSE for the used learning rates which are generally standard in machine learning, and seems slightly more stable than the simple error. The RMS-error function is used for further improvements on the reinforcement learning tool.

An iteration through learning rates of 0,01, 0,03 and 0,05 and discount rates of 0,995, 0,985 and 0,975 and their combinations shows an optimal behaviour for the combination of a learning rate of 0,03 and a discount rate of 0,975. The result for this combination is depicted in Figure **6**. With five sample reuses, one obtains a good compromise between learning speed, stability and very low overfitting.

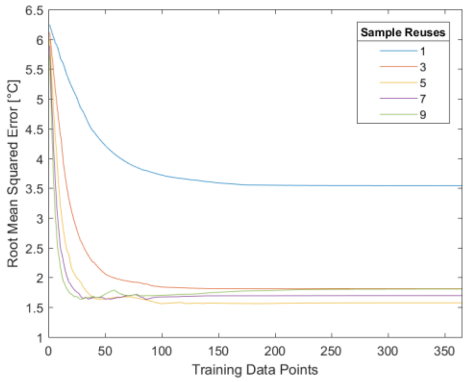

*Figure **6** Reinforcement Learning: Discounted Sample Reuse - Learning Rate 0,03, Discount Rate 0,975*

5.4 Energetic Evaluation

The energetic performance of the tools is separated into heating and cooling. The reinforcement learning tool is able to outperform the baseline scenario for heating, while the performance of the supervised leaning is on par with the baseline scenario. Given the differences between the setpoints used in the building simulation and the actual user inputs it could be concluded that the reinforcement learning tool would require more energy in a real application than the simulation implies, as the setpoint temperatures are on average too low compared to actual user input. The baseline scenario on the other hand has too high setpoints, implying overheating. A very well-tuned supervised artificial neural network may be able to reduce overheating and therefore reduce the heating load. The cooling load for the machine learning tools gathered from the building simulations is approximately twice as high as the load for the baseline scenarios. However, the cooling setpoints of the machine learning tools match the user inputs decently well, while the baseline scenario was roughly 2 °C too warm in the transition seasons spring and autumn, and 3,5 °C too warm in summer, compared to user inputs. This implies a drastically higher energy consumption in a real application for the baseline scenario as compared to the simulation. As there is no over-conditioning when using the baseline scenario, it can be concluded that the machine learning tools cannot provide any energetic benefit for cooling. Especially when using activated building components which are thermally slow, they should help to increase user comfort due to the increased precision in predicting comfort temperatures.

6. Conclusion, Outlook

In summary, both supervised and reinforcement learning tools could predict individual comfort temperatures based on weather parameters with the desired precision and their predictions being closer to user desires than the currently common heating and cooling setpoint temperatures. Artificial neural networks may be able to reduce overheating, thus decreasing the overall heating load, while machine learning tools in general might help to increase indoor comfort during the cooling period.

However, this potentially comes at the cost of increasing the cooling load, depending on user behaviour.

A deeper analysis of the topic seems worthwhile, considering the impact setpoint prediction may have on the flexibility and efficiency of building energy management systems, especially in combination with other advanced control techniques such as occupancy prediction and (model) predictive control. Further tests need to be performed using a wider range of data, from varying locations to be able to give a statement about the general applicability. As the data basis for this work was coming from a bank the comfort temperatures throughout the year, while varying, lay within a certain range, especially since the choice of clothing is highly limited. Predicting comfort temperatures in workplaces with a larger possibility of clothing options may prove to be more challenging, and would thus prove to be a good benchmark for the tools. Apart from more theoretical analyses real life implementations are an important next step. They would enable insights into actual changes in energy consumption and, more importantly, create user feedback. As shown by *Yang* and *Newman* user satisfaction and willingness to interact with learning machines is an important factor towards their performance [39]. An advanced implementation would later also test the interaction of comfort prediction, occupation prediction and model predictive control. Implementing the tools within an existing framework would require an interface that enables docking into it. Considering that more and more companies are moving into the market, a common standard would be useful [40].

References

[1] European Commision. Indoor air pollution: new EU research reveals higher risks than previously thought. Brussels, Belgium; 2003.

[2] Gorvett Z. The never-ending battle over the best office temperature. BBC 2016, 20 June 2016; Available from: http://www.bbc.com/capital/story/20160617-the-never-ending-battle-over-the-best-office-temperature. [March 03, 2017].

[3] van Hoof J. Forty years of Fanger's model of thermal comfort: comfort for all? Indoor Air 2008;18(3):182–201.

[4] Ramasubramanian K, Singh A. Machine Learning Using R. Berkeley, CA: Apress; Springer Science+Business Media; 2017.

[5] Holopainen R, Tuomaala P, Hernandez P, Häkkinen T, Piira K, Piippo J. Comfort assessment in the context of sustainable buildings: Comparison of simplified and detailed human thermal sensation methods. Building and Environment 2014;71:60–70.

[6] Wagner A. Thermal comfort and air quality requirements - impact on design strategies and HVAC concepts. Karlsruher Institut für Technologie; 2014.

[7] Rietschel H. Raumklimatechnik. 16th ed. Berlin, Heidelberg: Springer; 2008.

[8] American Societey of Heating, Refigerating and Air-Conditioning Engineers. Physiological principles and thermal comfort(19938.1-8.32). Atlanta: ASHRAE.

[9] Deutsches Institut für Normung. Indoor environmental input parameters for design and assessment of energy performance of buildings addressing indoor air quality, thermal environment, lighting and acoustics(EN 15251:2012-12). Berlin: Beuth Verlag GmbH; 2012.

[10] ISO International Organization for Standardization. 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(ISO 7730:2006-05).

[11] Nicol JF, Humphreys MA. Adaptive thermal comfort and sustainable thermal standards for buildings. Energy and Buildings 2002;34(6):563–72.

[12] O'Brien W, Gunay HB. The contextual factors contributing to occupants' adaptive comfort behaviors in offices – A review and proposed modeling framework. Building and Environment 2014;77:77–87.

[13] Gao PX, Keshav S. SPOT: A Smart Personalised Office Thermal Control System. UWSpace 2013.

[14] Nicol F, Humphreys M. Derivation of the adaptive equations for thermal comfort in free-running buildings in European standard EN15251. Building and Environment 2010;45(1):11–7.

[15] Ng A. Machine Learning; 2017.

[16] Paluszek M, Thomas S. MATLAB Machine Learning. 1st ed. Berkeley, CA: Apress; Springer Science+Business Media LLC; 2017.

[17] Kwok JT, Zhou Z-H, Xu L. Machine Learning. Springer Handbook of Computational Intelligence 2015:495–522.

[18] Dalamagkidis K, Kolokotsa D, Kalaitzakis K, Stavrakakis GS. Reinforcement learning for energy conservation and comfort in buildings. Building and Environment 2007;42(7):2686–98.

[19] Kubat M. An Introduction to Machine Learning. Cham: Springer; Springer International Publishing; 2015.

[20] Ferreira PM, Ruano AE, Silva S, Conceição E. Neural networks based predictive control for thermal comfort and energy savings in public buildings. Energy and Buildings 2012;55:238–51.

[21] Aswani A, Master N, Taneja J, Krioukov A, Culler D, Tomlin C (eds.). Energy-Efficient Building HVAC Control Using Hybrid System LVMPC; 2012.

[22] Ruano A, Pesteh S, Silva S, Duarte H, Mestre G, Ferreira PM et al. PVM-based intelligent predictive control of HVAC systems. IFAC-PapersOnLine 2016;49(5):371–6.

[23] Ferreira PM, Silva SM, Ruano AE (eds.). Energy Savings in HVAC Systems Using Discrete Model-Based Predictive Control; 2012.

[24] Dounis AI, Caraiscos C. Advanced control systems engineering for energy and comfort management in a building environment—A review. Renewable and Sustainable Energy Reviews 2009;13(6-7):1246–61.

[25] Erickson VL, Cerpa AE (eds.). Occupancy Based Demand Response HVAC Control Strategy; 2010.

[26] Barrett E, Linder S. Autonomous HVAC Control, A Reinforcement Learning Approach. Lecture Notes in Computer Science 2015;9286:3–19.

[27] Tso GK, Yau KK. Predicting electricity energy consumption: A comparison of regression analysis, decision tree and neural networks. Energy 2007;32(9):1761–8.

[28] Mamidi S, Chang Y-H, Maheswaran R (eds.). Improving building energy efficiency with a network of sensing, learning and prediction agents; 2012.

[29] Edwards RE, New J, Parker LE. Predicting future hourly residential electrical consumption: A machine learning case study. Energy and Buildings 2012;49:591–603.

[30] Tsanas A, Xifara A. Accurate quantitative estimation of energy performance of residential buildings using statistical machine learning tools. Energy and Buildings 2012;49:560–7.

[31] Zhao H-x, Magoulès F. A review on the prediction of building energy consumption. Renewable and Sustainable Energy Reviews 2012;16(6):3586–92.

[32] Sun Y, Huang G. Recent Developments in HVAC System Control and Building Demand Management. Curr Sustainable Renewable Energy Rep 2017;4(1):15–21.

[33] Buratti C, Vergoni M, Palladino D. Thermal Comfort Evaluation Within Non-residential Environments: Development of Artificial Neural Network by Using the Adaptive Approach Data. Energy Procedia 2015;78:2875–80.

[34] Afram A, Janabi-Sharifi F, Fung AS, Raahemifar K. 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 2017;141:96–113.

[35] Yu Z, Dexter A. Online tuning of a supervisory fuzzy controller for low-energy building system using reinforcement learning. Control Engineering Practice 2010;18(5):532–9.

[36] Deutscher Wetterdienst. CDC (Climate Data Center). Offenbach; Available from: ftp://ftp-cdc.dwd.de/pub/CDC/. [06.13.2017].

[37] Bundestag of the Federal Republic of Germany. Verordnung über energiesparenden Wärmeschutz und energiesparende Anlagentechnik bei Gebäuden: Energieeinsparverordnung (EnEv); 2002.

[38] Sutton RS, Barto AG. Reinforcement Learning: An Introduction. 2nd ed. Cambridge, Massachusetts: The MIT Press; 2012.

[39] Yang R, Newman MW (eds.). Learning from a learning thermostat: lessons for intelligent systems for the home. New York, NY, USA: ACM; 2013.

[40] Lobo S. Keynote: Künstliche Intelligenz. Classic Remise Berlin; 2017.