Low-Carbon Comfort Management for Smart Buildings

Jennifer Williams, Benjamin Lellouch, Sebastian Stein, Christina Vanderwel, Stephanie Gauthier
University of Southampton, United Kingdom
{J.Williams, brl1u18, S.Stein, C.M.Vanderwel, S.Gauthier}@soton.ac.uk

***Abstract*—We present critical research challenges for the development of smart building management systems (BMS) to achieve low-carbon comfort. To date, work in this area has focused on optimising single-scope aspects of building resources, such as energy usage or thermal comfort, but there is a recent shift toward BMS design that could simultaneously address many aspects of building resources and comfort dimensions for occupants, such as air quality, temperature, humidity, audible noise levels, and related automated safety features. In this paper, we discuss four research directions highlighting current challenges in this domain that present opportunities for research: (A) data limitations for machine learning, (B) multiple definitions of comfort, (C) BMS usability and interfaces, and (D) safety and security of automated BMS decision-making. Addressing these challenges will enable the development of advanced human-centred energy-saving buildings that meet the needs of occupants.**

***Index Terms*—Agent-based modeling, human in the loop, green buildings, energy management, environmental monitoring**

I. Introduction

Smart buildings are buildings that have been designed to meet the human and energy needs of their occupants through technical infrastructure, functionality, and outcomes such as less energy waste and greater satisfaction from the occupants who use the building [1]. Smart buildings have the potential to not only reduce the consumption of resources (e.g., reducing costs, conserving gas, electricity, and water, reducing time and effort spent planning or managing these), but also to vastly improve the quality of life for individuals. Buildings of the future may be designed from the ground up or existing buildings may be retrofitted with certain capabilities that contribute to net-zero carbon goals and increased comfort. Historical data can be collected to provide information about the indoor environment of buildings from sensors that monitor and detect changes of the environment. This data is valuable because it enables researchers to build models of environments using agent-based machine learning techniques like reinforcement learning (RL). These models, such as Gnu-RL [2], can run efficient simulations for the environment variables without needing to interact directly with the real-world environment. For example, a model like Gnu-RL has a shorter training period (4 simulated *days*) compared to model-free algorithms (47.5 simulated *years*). Models of indoor environments can then be used to create a type of automated system called a building management system (BMS) [3], [4]. The machine learning aspect of a BMS allows it to adapt and respond to multiple types of human behaviour in the building.

Until recently, BMS design has focused on specific and narrowly-defined niche problems such as optimising energy consumption, or optimising thermal comfort. In fact, the idea of smart building management systems is starting to shift toward a more all-encompassing approach [5], [6] that incorporates multiple variables simultaneously, such as occupancy/activity modeling, resource consumption (e.g., lighting, temperature, humidity, noise, etc), and especially the optimisation of energy-related resources to reduce carbon footprint. Air quality is particularly important due to airborne transmission of diseases like the SARS-CoV-2 virus and seasonal flu [7], [8]. We envision smart buildings that also provide relevant comfort information to their occupants in a manner that helps occupants make informed choices.

A BMS can be designed with varying degrees of automated decision-making [3]. It is an open problem to determine how decision-making can be optimally distributed between a BMS and humans in the loop. Some of the decision-making may depend on the building type, the role of occupants, and how various comfort-related variables are integrated into the system. In this paper, we present a vision to some of the challenges that the research community needs to address in order to design a BMS for smart buildings. This paper provides clarity on the numerous research activities that will contribute to effective, relevant, usable, and safe systems.

II. Research Directions and Challenges

This paper sets the agenda for four key research directions and discusses the related challenges for each direction: (A) overcoming data limitations, (B) integrating multiple definitions of comfort, (C) designing usable interfaces, and (D) mitigating safety concerns. Many of these challenges overlap to some degree. In this section, we illuminate the most pressing issues for each challenge and propose some directions for future research.

A. Overcoming Data Limitations

Challenge. ***Incorporating comfort metrics into building management systems and energy usage forecasting requires innovative transfer learning solutions as well as the creation***

This work is supported by the UK Engineering and Physical Sciences Research Council (EPSRC) through the Trustworthy Autonomous Systems Hub (EP/V00784X/1) and a Turing AI Acceleration Fellowship on Citizen-Centric AI Systems (EP/V022067/1).

of novel datasets (real or synthetic) that can be adapted to multiple building types.

There are many different types of buildings that can benefit from a decision-making system that learns to optimise both comfort and resources. These buildings include schools, hospitals, libraries, offices, museums, and private homes, among others. Each building has different needs [3]–[5]. Likewise, datasets that are meant to model variables from one type of building may not be directly applicable to another type. For example, it is a difficult challenge to develop a BMS system for a school building that is fundamentally based on data collected or simulated from an office building, because there are different underlying assumptions in each case [5]. More research is needed for transfer learning, which adapts data or algorithms from one type building for the purpose of modeling other types. The work of [5] identified four data-related areas that must be considered in terms of transfer learning for BMS design: building loads, occupancy/activity, building dynamics, and energy systems. A comprehensive solution to the problem of transfer learning could survey how these four categories map between building types. This would effectively create a typology of building features, allowing researchers to navigate the transfer learning problem more effectively. Data must also be representative of the building users and address potential biases (e.g., if most of the occupants in a given dataset belong to specific demographics).

There is a need for novel datasets (both synthetic and real-world). Each existing dataset is in some sense "incomplete", meaning that no single dataset has all of the information to build a reliable and safe BMS that is optimised for low-carbon comfort. Even new datasets, such as AlphaBuilding [9] and AlphaBuilding-MedOffice[1], do not always have all of the variables needed for modeling. For example, AlphaBuilding includes simulated occupancy data, but the dataset is not suitable for RL algorithm development due to lack of interaction modeling in the data. On the other hand, AlphaBuilding-MedOffice and Gnu-RL [2] do not include occupancy information, but are suitable for RL algorithms. Yet, they are limited to medium-sized office buildings. Another challenge to be solved is determining the amount of historical data necessary (e.g., weeks, months or years) for a BMS to be effectively deployed to different environments. It is also not known how sensor placement and sensor calibration might affect BMS algorithm performance once deployed.

Real-world data is expensive and time-consuming to collect. Depending on the variables being monitored this may involve installing new sensors (such as monitoring air quality, humidity, ventilation, and indoor/outdoor air temperature) and may require obtaining consent from occupants (when monitoring occupancy/activity levels). Additionally, it may be difficult to obtain parallel financial cost information related to consumption of resources. Real data is ultimately dependent upon which sensors are installed in a building and their placement. BMS design is limited to the variables that have corresponding data. Some sensors, such as CO_2 are relatively inexpensive to install, while others (e.g., video) that monitor occupancy or activity levels are potentially more invasive and may require extra data processing. Ideal solutions to addressing the overall problem of data scarcity could meld together limited real-world data with estimated or simulated data to create hybrid datasets or to fill in small gaps where information from a particular sensor is unavailable.

[1]https://github.com/WalterZWang/AlphaBuilding-MedOffice

B. Integrating Multiple Definitions of Comfort

Challenge. ***Different types of buildings, occupants, and usage scenarios require the consideration of new, multi-dimensional comfort metrics. Similarly, these metrics may need to be optimised at different levels of granularity, either for groups of people or for areas of a building.***

Conceptually, the idea of *comfort* is embodied by many different variables. There is currently no single comfort metric definition that can be optimised for smart buildings, yet this is necessary for developing advanced BMS for smart buildings. Comfort metrics are typically taken to be a collection of metrics, with each one addressing a need separately [10]. In this research direction, we propose integrating multiple definitions of comfort to identify ones that can be modeled at coarse and fine granularity. For example, air quality could be optimised and monitored at a building-level, whereas temperature and ventilation could be optimised at a smaller scale. Some of the comfort dimensions to consider are: temperature, ventilation, humidity, air quality, noise levels, lighting, feelings of perceived privacy, and water temperature in taps.

Previous work has established comfort standards, such as the American Society of Heating, Refrigerating, and Air Conditioning Engineers (ASHRAE). The ASHRAE Standard 55 defines thermal comfort in buildings and includes two approaches to assess thermal comfort described as follows: (1) the predictive approach (with PMV and PPD indices) and (2) the adaptive approach (where indoor operative temperature is related to mean outdoor air temperature). The acceptable threshold is defined by 80% or more of occupants finding their area thermally acceptable. There is also an ASHRAE Global Thermal Comfort Database II, which is a large collection of instances that combine objective measurements of indoor temperatures with subjective or perceived comfort in each case [11]. The database can be queried to obtain comfort settings based on building type, occupancy, and subjective or objective evaluations of comfort. Another metric is EN 16798-1:2019, and this is a European standard that defines requirements for indoor environmental domains, such as indoor air quality, acoustics and lighting in regards to building design. This standard focuses on the design input parameters for the design of buildings and systems, including occupants' thermal comfort [12].

There are many interrelated ASHRAE standards for building comfort. Additional research in this area should seek to incorporate multiple comfort standards and possible interactions between environmental domains into RL-based machine learning approaches to building management. This will allow for modeling the interplay between multiple types of comfort alongside multiple types of resource consumption [22]. Recent work from [13] has attempted to create a BMS that uses ASHRAE 55 to govern automatic control of natural and artificial ventilation in a large building in Moscow, Russia. In that simulation, individuals had direct control over the opening or closing of windows, and the BMS adjusted accordingly based on the individual decisions.

Apart from defining the elements that make up *comfort*, it is also necessary to better understand which of these elements can be optimised for individuals or groups. The answer may depend on which comfort variable is under consideration, as well as the type of building or usage/occupancy of a building. For example, in a school, it might be preferable for some classrooms to have different settings from each other at different times of day after taking into consideration the preferences of classroom occupants as a group [14]. On the other hand, in an office building that is composed of many single-occupancy offices, it may be more appropriate to allow each office room to be adjusted based on individual preferences. Another way to frame the problem involves thinking of buildings in terms of zones or regions, where each zone has particular settings. Setting up zones that have particular comfort settings could be facilitated by certain design aspects of a building, which may have the effect of further reducing resource consumption.

C. Designing Usable Interfaces

Challenge. ***To maximise utility, a BMS should provide services considering the human in the loop. The system should provide information to occupants as well as incorporate feedback from occupants (or facilities management). Therefore, there is a need to develop information sharing techniques between humans and the smart system.***

A recent review on human-building interfaces and their relationship to human behavior, energy use and occupant comfort [21] has highlighted critical aspects of interfaces design including *"the ease and access of control, interface/control placement, poor interface/control design, lack of understanding, and social-behavioral dynamics"*.

We envision the development of a BMS approach for smart buildings that simultaneously optimises multiple aspects of comfort and utilisation of resources. It is important to also consider how the BMS system can be designed to be user-friendly. For example, it is not currently known how to effectively exploit data visualisation for all of the data received through sensors. New research should address how to develop interfaces and visualisations capable of addressing accessibility needs, indicating an overall building status, identifying areas where comfort is not being met, identifying areas where comfort is achieved at a high cost of resources, and similar summaries that may be relevant to a building facilities manager or facilities department. This type of interface can be achieved through additional research that explores how data visualisation and user interface design affect human decision-making, which in turn affects BMS decision-making.

In recent years, many office buildings have started utilising hot-desks as a way to preserve office real-estate by allowing workers to use any available desk. In a hot-desk scenario, it may not be possible to create comfort conditions for each individual because office occupants change constantly or hot-desks are implemented as part of an open-office plan [15]. Even a simple averaging among occupants may nullify perceived indoor environment conditions or require tracking of who is inside of a building and where.

Occupants who are allowed direct influence of the comfort settings of a building are *active consumers* of environmental comfort, whereas some occupants (e.g., visitors or hot-desk employees) may be viewed as *passive consumers*. Other passive consumers of environmental conditions can include visitors to a museum, or visitors to a library. In the case where occupants are passive consumers of comfort, it may be reasonable to explore how "zones" can be used within the building design. Occupants could be given information about the settings of each zone upon entry to the building and then the occupants can decide which areas of the building they would like to be in (or avoid). A visualization of the building status for passive consumers could include an interactive map of the building with information about comfort settings (and current conditions) in different areas. Occupancy could easily be incorporated into such a visualization, for example when visiting a museum to decide which areas and exhibits are busiest, noisiest, or have the most lighting.

Occupant roles with respect to a building can potentially impact the design of a BMS. In some cases, human-level decisions may be left to a facilities manager who decides the target settings and provides this as input to the BMS. In other cases, the BMS can use techniques from machine learning to learn preferences of occupants by monitoring how the occupants respond to changes. The BMS can then use this to automatically adapt and optimise its operation. The latter case would require occupants to provide direct or indirect feedback. For example, if the BMS is trying to learn a temperature and it is too warm, the occupant may constantly open a window. It would be left to the BMS to determine that this behavior is related to a preferred temperature and not poor indoor air quality (e.g., perceived 'stuffiness' or smell/odour). Given the volume of variables that will be modeled and optimised simultaneously, it may be intractable for a BMS to learn all settings. We propose investigation of how occupants would best co-create the BMS and learning conditions.

D. Mitigating Safety Concerns

Challenge. ***Smart building management systems require investigation into safety issues that could arise from differential levels of autonomous decision making, including threats that can be exploited maliciously, and provide means to mitigate those risks to occupants.***

Safety and automated decision-making are two important aspects of smart buildings. Not all safety concerns are a direct consequence of malicious activity. In the first instance, sensor data networks need to be kept updated and protected, so that a BMS can receive accurate data inputs. If the data at the sensor-level is compromised, then this can impact safety in many ways, in addition to compromising comfort and resources. Bad sensor input could be the result of a simple sensor failure. New comprehensive approaches are needed to detect when sensors are starting to fail, as well as new strategies for sensor repair and replacement.

The second type of safety consideration involves how a BMS reacts to inputs. For example, if a BMS detects a fire or other emergent hazard, what is the response? Recent work from [16] has begun to explore conflict mediation for smart buildings. The first step is to identify a conflict (e.g., a race condition) and then offer a mediation action that suggests an action for a human to perform in order to fix the conflict (e.g., turn off a sensor for a short period of time). Similarly, work from [17] explored conflict mediation at the human level in so-called "thermostat wars". In that human in the loop scenario, a room occupant was decidedly uncomfortable and requested a vote to increase temperature. Several occupants voted 'yes' while some voted 'no'. The system determined which action to take and recommended some additional actions for occupants who voted 'no'. Future work involving cooperative behavior could explore the use of mechanism design [21], where incentives influence occupant decisions.

The third type of safety consideration is whether the BMS can be maliciously exploited by a bad actor. There is increasing awareness among researchers that smart buildings have vulnerabilities that could be exploited with relative ease [19]. Such exploitation may involve, for example, disabling air quality inputs to a BMS. Cyber-security work for this aspect of smart building safety is already underway with the creation of simulated datasets, such as the HVAC Attack Database [20] (Heating, Ventilation, and Air Conditioning). While this is a useful database, a variety of smart building vulnerabilities should be explored in greater depth (e.g., total loss of electricity, automatically locking doors, emergency services) to understand safety risks.

III. Discussion

Integrating comfort ideals into a BMS in a way that can also optimise resource consumption is an ambitious goal, but one that has far-reaching positive impacts for society. In this paper, we have proposed four critical pathways for developing research that can lead to successful smart, comfortable, and green buildings. First, we described how data holds a critical role for algorithm development. We discussed the limitations of existing datasets and proposed increased research and innovation for transfer learning. Second, we highlighted current and previous work on the development of comfort metrics and discussed how some metrics (e.g., air quality) are well-suited to be optimised at a holistic building-level while other metrics are suitable to be optimised for individuals, depending on the building type and usage. Third, we discussed how information and user interfaces contribute to the usability of a BMS depending on occupant roles. Finally, we presented three types of safety concerns from physical sensor failures to malicious cyber attacks and propose additional research in this area to explore all aspects of safety.

The four challenges that we presented are intertwined in terms of the variables that must be considered for comfort and resources. We specifically emphasise that a smart BMS does not need to be a fully-automated system but should include human in the loop decision-making. Developing trust in a BMS that is designed for low-carbon comfort requires high usability and access to information as well as successful mitigation of safety concerns.

References

[1] Jia, Ruoxi, Baihong Jin, Ming Jin, Yuxun Zhou, Ioannis C. Konstantakopoulos, Han Zou, Joyce Kim et al. "Design automation for smart building systems." Proc. of the IEEE 106, no. 9 (2018): 1680-1699.

[2] Chen, Bingqing, Zicheng Cai, and Mario Bergés. "Gnu-RL: A precocial reinforcement learning solution for building hvac control using a differentiable mpc policy." Proc. of 6th ACM Intl Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation. 2019.

[3] Gul, Mehreen S., and Sandhya Patidar. "Understanding the energy consumption and occupancy of a multi-purpose academic building." Energy and Buildings 87 (2015): 155-165.

[4] Han, Mengjie, Ross May, Xingxing Zhang, Xinru Wang, Song Pan, Da Yan, Yuan Jin, and Liguo Xu. "A review of reinforcement learning methodologies for controlling occupant comfort in buildings." Sustainable Cities and Society 51 (2019): 101748.

[5] Pinto, Giuseppe, Zhe Wang, Abhishek Roy, Tianzhen Hong, and Alfonso Capozzoli. "Transfer learning for smart buildings: A critical review of algorithms, applications, and future perspectives." Advances in Applied Energy (2022): 100084.

[6] Yoon, Young Ran, Ye Rin Lee, Sun Ho Kim, Jeong Won Kim, and Hyeun Jun Moon. "A non-intrusive data-driven model for detailed occupants' activities classification in residential buildings using environmental and energy usage data." Energy and Buildings 256 (2022): 111699.

[7] McNeill, V. Faye. "Airborne transmission of SARS-CoV-2: evidence and implications for engineering controls." Annual Review of Chemical and Biomolecular Engineering 13 (2022).

[8] Corsi, Richard, Shelly L. Miller, Marissa VanRy, Linsey C. Marr, Leslie R. Cadet, Nira R. Pollock, David Michaels et al. "Designing infectious disease resilience into school buildings through improvements to ventilation and air cleaning." The Lancet Covid-19 Commission, Task Force on Safe Work. (2021).

[9] Li, Han, Zhe Wang, and Tianzhen Hong. "A synthetic building operation dataset." Scientific Data 8.1 (2021): 1-13.

[10] Atzeri, Anna Maria, Francesca Cappelletti, Athanasios Tzempelikos, and Andrea Gasparella. "Comfort metrics for an integrated evaluation of buildings performance." Energy and Buildings 127 (2016): 411-424.

[11] Ličina, Veronika Földváry, Toby Cheung, Hui Zhang, Richard De Dear, Thomas Parkinson, Edward Arens, Chungyoon Chun et al. "Development of the ASHRAE global thermal comfort database II." Building and Environment 142 (2018): 502-512.

[12] EN, CEN Standard. "16798-1. Energy performance of buildings—Ventilation for buildings—Part 1: Indoor environmental input parameters for design and assessment of energy performance of buildings addressing indoor air quality." Thermal Environment, Lighting and Acoustics-Module M1-6.(16798-1) (2019).

[13] Simmonds, Peter. "Using ASHRAE Standard 55 Adaptive Comfort Method for Practical Applications." In CLIMA 2022 Conference. 2022.

[14] McNeill, V. Faye, Richard Corsi, J. Alex Huffman, Cathleen King, Robert Klein, Michael Lamore, Shelly L. Miller et al. "Room-level ventilation in schools and universities." Atmospheric Environment: X (2022): 100152.

[15] Cooper, Peter Benjamin, Konstantinos Maraslis, Theo Tryfonas, and George Oikonomou. "An intelligent hot-desking model harnessing the power of occupancy sensing data." Facilities (2017).

[16] Liu, Renju, Ziqi Wang, Luis Garcia, and Mani Srivastava. "Remediot: Remedial actions for internet-of-things conflicts." Proc. of 6th ACM Intl Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, pp. 101-110. 2019.

[17] von Frankenberg, Nadine. "Towards Resolving Thermal Comfort Conflicts in Shared Spaces." In Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems, pp. 1-5. 2021.

[18] Wang, Tao, Yunjian Xu, Chathura Withanage, Lan Lan, Selin Damla Ahipaşaoğlu, and Costas A. Courcoubetis. "A fair and budget-balanced incentive mechanism for energy management in buildings." IEEE Transactions on Smart Grid 9, no. 4 (2016): 3143-3153.

[19] dos Santos, Daniel Ricardo, Mario Dagrada, and Elisa Costante. "Leveraging operational technology and the Internet of things to attack smart buildings." Journal of Computer Virology and Hacking Techniques 17, no. 1 (2021): 1-20.

[20] Elnour, Mariam, Nader Meskin, Khaled Khan, and Raj Jain. "HVAC system attack detection dataset." Data in Brief 37 (2021): 107166.

[21] Day, J. et al. "A review of select human-building interfaces and their relationship to human behavior, energy use and occupant comfort." Building and Environment, Vol. 178. 2020.

[22] Schweiker, M. et al. "Review of multi-domain approaches to indoor environmental perception and behaviour" Building and Environment, Vol. 176. 2020.