10-2024

OcAPO: Fine-grained occupancy-aware, empirically-driven PDC control in open-plan, shared workspaces

Ravi ANURADHA
University of Maryland at Baltimore

Dulaj Sanjaya WEERAKOON
sanjayawm@smu.edu.sg

Archan MISRA
Singapore Management University, archanm@smu.edu.sg

Pervasive and Mobile Computing 103 (2024) 101945

Contents lists available at ScienceDirect

Pervasive and Mobile Computing

journal homepage: www.elsevier.com/locate/pmc

OcAPO: Fine-grained occupancy-aware, empirically-driven PDC control in open-plan, shared workspaces

Anuradha Ravi [a,*], Dulaj Sanjaya Weerakoon [b], Archan Misra [b]

[a] Department of Information Systems, College of Engineering and Information Technology, University of Maryland Baltimore County, 1000 Hilltop Cir, Baltimore, 21250, MD, USA
[b] School of Computing and Information Systems, Singapore Management University, 80 Stamford Rd, Singapore, 178902, Singapore

A R T I C L E I N F O

Keywords:
HVAC control
Thermal comfort
Smart building management
Occupancy estimation

A B S T R A C T

Passive Displacement Cooling (PDC) is a relatively recent technology gaining attention as a means of significantly reducing building energy consumption overheads, especially in tropical climates. PDC eliminates the use of mechanical fans, instead using chilled-water heat exchangers to perform convective cooling. In this paper, we identify and characterize the impact of several key parameters affecting occupant comfort in a $1000\ \mathrm{m}^2$ open-floor area (consisting of multiple zones) of a ZEB (Zero Energy Building) deployed with PDC units and tackle the problem of setting the temperature setpoint of the PDC units to assure occupant thermal comfort and yet conserve energy. We tackle two key practical challenges: (a) the zone-level (i.e., occupant-experienced) temperature differs significantly, depending on occupancy levels, from that measured by the ceiling-mounted thermal sensors that drive the PDC control loop, (b) sparsely deployed sensors are unable to capture the often-significant differences in ambient temperature across neighboring zones. Using extensive real-world coarser-grained measurement data (collected over 60 days under varying occupancy conditions), (a) we first uncover the various parameters that affect the occupant-level ambient temperature, and then (b) devise a trace-based model that helps identify the optimum combination of PDC setpoints, collectively across multiple zones, while accommodating variations in the occupancy levels and weather conditions. Using this trace-based model, our *OcAPO* system can assure ambient temperature experienced by occupants within a tolerance of $0.3\ °\mathrm{C}$. In contrast, the existing approach of occupancy-agnostic, rule-based setpoint control violates this tolerance interval more than 80% of the time. However, this initial model requires unnecessary and continual database lookups and is unable to derive finer-grained setpoints, thereby potentially missing opportunities for additional energy savings. We thus collected data for another 15 days, with finer-grained setpoint control in increments of 0.2° under varying occupancy conditions in the second phase. To determine PDC setpoints efficiently, we subsequently used the empirical data to train a KNN-based regression model. Additional studies on our real-world testbed demonstrate the regressor-based *OcAPO* approach is able to assure occupant-level ambient temperature within a narrow $0.2\ °\mathrm{C}$ tolerance. We also demonstrate that the regression version of *OcAPO* can reduce the opening percentage of PDC valves (an indirect proxy for energy consumption) by 58.9% under low occupancy compared to the trace-based model.

* Corresponding author.
E-mail addresses: anuradhar@ieee.org (A. Ravi), archanm@smu.edu.sg (A. Misra).

https://doi.org/10.1016/j.pmcj.2024.101945
Received 6 December 2023; Received in revised form 27 March 2024; Accepted 15 May 2024
Available online 28 May 2024

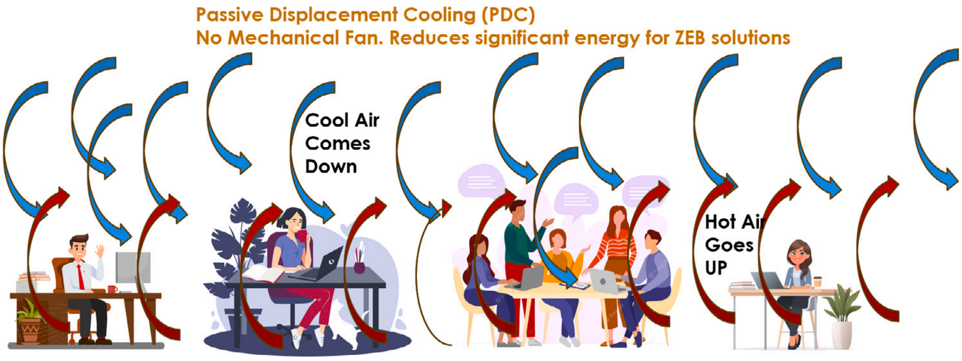

Fig. 1. Working principle of Passive Displacement Cooling (PDC).

1. Introduction

Commercial buildings generate ∼33% of the world's total energy consumption, within which HVAC represents the dominant (38%, translating to 12% of the global energy demand) consumption load [1]. Passive Displacement Cooling (PDC) technologies[1] is a relatively new, energy-efficient cooling mechanism (reducing the cooling energy overhead by ∼15%) that is gaining increasing traction, especially in tropical regions where heating support is typically unnecessary. As illustrated in Fig. 1, PDCs fundamentally dispense with the use of energy-consuming fans for forced air circulation, and instead utilize the principles of natural convective cooling, where chilled air (typically cooled by heat exchange with chilled water pipes) drifts towards the lower portions of the room while warmer air rises towards the ceiling. Combining such PDC with occupancy-aware control [2], can further lower the overall energy cost.

Early research on occupancy-aware HVAC control relied principally on a binary estimator occupancy (occupied vs. unoccupied), often using PIR [3] or CO_2 sensing [4] to estimate whether a region has *any* occupants or not, and accordingly activate or deactivate cooling functions. To overcome the natural transients (hysteresis) associated with cooling, more sophisticated methods employ occupancy prediction techniques [5,6] to control HVAC valves preemptively, initiating the cooling down of space before the anticipated arrival of human occupants. These strategies have, however, been predominantly tested in traditional *compartmentalized* commercial spaces, where individual layouts (e.g., office or meeting rooms) are effectively thermally isolated and where the binarized occupancy characterization is adequate for controlling the cooling of individual compartments.

In this work, we tackle the problem of performing efficient, occupancy-aware PDC-based cooling in more modern, open-plan workspaces where the lack of explicit walls/partitions implies a strong mutual correlation between the thermal states of nearby or neighboring spaces. Our goal is to intelligently adjust the temperature set points, which in turn activate or deactivate PDC controllers, so as to assure *occupant thermal comfort* in such open spaces where the aggregate number of occupants can exhibit significant variation. As a specific exemplar of this problem, we tackle the issue of occupancy-aware, thermal comfort preserving, PDC control of a ∼1000 m^2 open-plan collaboration space on one floor of an operational university ZEB (Zero-Energy Building). Fig. 2 illustrates the layout of our collaborative workspace, which consists of 8 zones (virtually divided into 3 areas), each individually controlled by a ceiling-mounted PDC air vent that supplies cold air, which drifts downwards, while the hot air rises. For this specific deployment, the PDC vents open up to circulate the cool air when the *nominal* room temperature, measured via ceiling-mounted temperature sensors, exceeds the zone's PDC setpoint and closes when this temperature falls below the zone's PDC setpoint. Fig. 2 illustrates this phenomenon. The ceiling-mounted temperature sensors measure the *nominal* room temperature of a given area and push the updates to the Building Management System (BMS). The PLC control logic in the BMS checks if the PDC setpoint is less than the ceiling temperature. If so, the PDC valves are opened up to supply cool air; else the PDC valves remain closed. The valves open between 10%–100% depending on the difference in the temperature of the PDC setpoint and ceiling temperature.

The collaboration space presents a couple of difficult challenges that we believe to be representative of legacy environments where the PDC system, including vents and PLC control logic, have been deployed by commercial vendors and which we now seek to augment with more sophisticated occupancy-aware control techniques. (Note that the only means of programmatic external control is via PDC setpoint alteration, as the underlying low-level PLC control logic, which also drives the valve opening percentage between 10% and 100%. is not accessible.)

- The temperature sensors, whose readings are compared against the programmed set points, are mounted on the ceiling. Both our empirical studies (detailed in Section 4.2) and prior research [7] provide evidence that the actual floor/table-level ambient temperature, experienced by the human occupants, is higher than the temperature measured at the ceiling. Moreover, we shall show (Section 4.1) that this table-vs.-ceiling discrepancy is a function of different occupancy *levels*, with the ambient temperature being much higher under high occupancy. As a consequence, we need to ensure that the set point settings are appropriately adjusted, *in an occupancy-aware fashion*, to ensure acceptable human-experienced thermal comfort.

[1] https://www.daikin.com.sg/product-series/pdv/

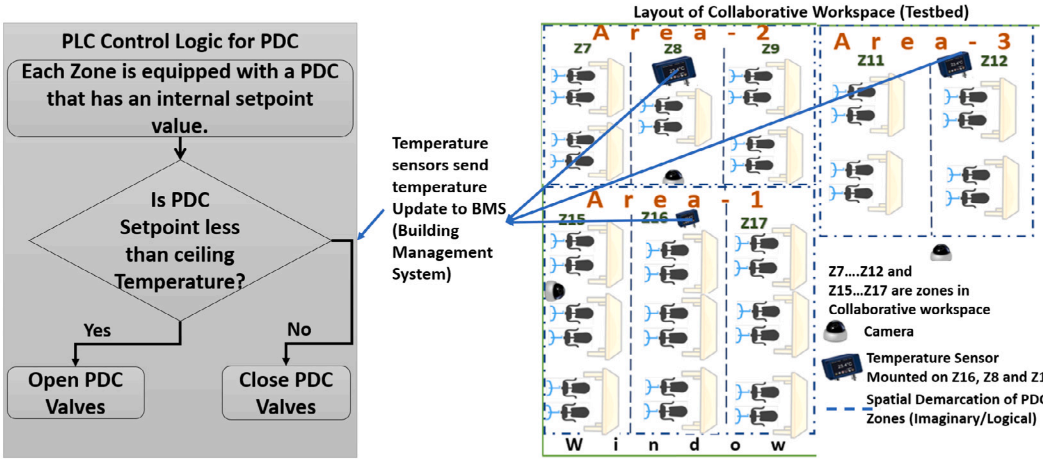

Fig. 2. Layout of collaborative workspace (Testbed) and interface with BMS.

- While the 8 zones in the exemplar workspace each have their individual PDC controller, the ceiling-level temperature sensor is deployed more sparsely, with one temperature sensor serving multiple (typically 2–3) zones. For instance, Z15, Z16, and Z17 (covering an area of 100 m^2) are together associated with a single temperature sensor, mounted on the ceiling of Z16. Given the open-plan layout of the space, this implies that the reading captured by the sensor is also a *combined* function of the possibly different temperature values in each of the constituent zones.

Any practical solution for dynamic setpoint adjustment must thus be able to (a) accommodate the varying occupancy levels of different zones, (b) estimate each zone's occupant-level temperature, and (c) control a *set* of setpoints jointly.

We present our approach, called *OcAPO*,[2] that effectively tackles these challenges. We explore two successively-staged versions of *OcAPO* that broadly operate as follows. In the initial, less-sophisticated variant, *OcAPO* utilizes extensive but coarser-grained empirical training data (captured over a wider operating range) to build a lookup table that correlates the desired zone-level temperature to appropriate temperature setpoint values under varying levels of *{occupancy levels, external weather conditions}*. In the subsequent, more efficient variant, *OcAPO* uses additional finer-grained data (captured over a smaller operating range) to train a regressor that can represent the above-mentioned correlation pattern. For operational deployment, we build a practical 15–20 min look-ahead occupancy prediction system that combines real-time, non-binary occupancy estimates (derived by executing vision-based DNNs on infrastructural camera feeds) with historical occupancy data to predict zone-level occupancy. Finally, using the occupancy and external weather measurements, *OcAPO* derive and sets temperature setpoints for multiple adjacent zones collectively so as to ensure that the *occupant-experienced ambient* stays within a desired, thermally-comfortable range.

Key Contributions: We make the following key contributions:

- Using extensive real-world measurements and observational data, we quantify both how (a) the relationship between a setpoint value and the occupant-perceived (table-level) ambient temperature is occupancy level dependent, and (b) the occupant-perceived ambient temperature differs (is usually higher than) the ceiling-level temperature sensor readings under varying levels of occupancy. These findings not only motivate the need for an occupancy-aware PDC control logic but also reveal the need to move beyond simple binary (occupied vs. unoccupied) occupancy estimates.
- We demonstrate the practical challenge of non-binary occupancy prediction (i.e., over 15–20 min lookahead intervals) for different regions of an open plan university collaboration space, where student movement patterns exhibit a high level of spontaneity. To tackle this challenge, we instead combine the diurnal occupancy trend (number of occupants in the morning/afternoon/evening) with real-time video sensing-based occupancy estimation to obtain reasonably accurate predictors for occupancy 15–20 min in advance.
- We then construct a trace-based model to find, given the weather condition and estimated per-zone occupancy, the optimum tuple of zone setpoints that minimize the cumulative deviation from a target ambient temperature (24.0 °C). This initial model utilizes data traces, collected in coarser increments of 0.5° over a wider range to minimize the amount of empirical data collection effort. The model utilizes a lookup table, where the occupancy is expressed in 4-levels (unoccupied, low, medium, high), and the weather conditions are represented by 2 °C intervals. This approach results in a relatively low error in achieving optimal thermal comfort across a range of occupancy conditions, achieving mean error of 0.2 °C under low, and 0.4 °C under high occupancy levels. While effective, this lookup table-based model requires constant access to a database and cannot perform setpoint control in finer-grained increments (say in 0.2 °C intervals).

[2] **OC**cupancy-aware, **P**DC Control for **O**pen-Plan, **S**hared Workspaces for improved thermal comfort.

- To support even fine-grained temperature control, we obtain additional data under *targeted*, fine-grained PDC setpoint settings (in 0.2 °C increments) and then develop a KNN regressor that can generate "optimal" PDC setpoint settings (jointly for all zones of a given area) given the weather temperature, occupancy level, and required zone-level temperature. The KNN regressor (trained on the data collected over a 75-day period) first identifies the "neighborhood" points for a given test data and uses the mean-value for each such cluster/neighborhood to perform real-time, per-zone setpoint control. This approach enables *OcAPO* to achieve an even tighter ambient temperature tolerance of 0.1 °C under low and 0.2 °C under high occupancy levels, respectively.
- We conduct pilot studies, lasting 3 days under real-world occupancy conditions, on the ZEB collaborative space, evaluating the efficacy of *OcAPO's* occupancy-aware dynamic setpoint control logic (using both trace-based and KNN regressor model) vs. the baseline of constant 24 °C setpoint. Our results demonstrate the advantages of dynamic, occupancy-aware setpoint adaptation: over a total deployment period of 45 hrs, the KNN regressor-based *OcAPO* model specifies setpoints distinct from rules-based setting of 24 °C for 78% of the time – i.e., a total of 702 min – and achieves a mean cumulative deviation of 0.2 °C. In contrast, the baseline, rules-based constant setpoint method exceeded the tolerance range (24 °C) 86% of the time, with a mean deviation of 1.1 °C.
- We monitor the PDC valve opening percentage (%) when deploying the fine-grained KNN and coarser-grained trace-based *OcAPO* models, vs. the baseline BMS rule-based technique. We confirm that KNN-based *OcAPO* translates to more granular valve openings (an average of 56.4%) across all zones and occupancy compared to the trace-based *OcAPO* model, which opens the valves fully at 76.2% resulting in higher energy consumption. In contrast, BMS rule-based PDC control causes the PDC valves to be 38.7% open on average: while this arguably reduces the energy overhead, it causes much higher fluctuations in ambient temperature (e.g., by 1.1 °C under high occupancy), potentially adversely affecting occupant comfort. Thus, with just an increase of 17.7% opening of PDC valves, *OcAPO* can successfully maintain the optimal thermal comfort within the tolerance range of 0.2 °C.

2. Related work

The evidence of research in demand-driven HVAC control based on occupancy and weather conditions goes back for decades [8]. However, the evolution in hardware and the subsequent establishment of programmatic interfaces to control HVAC give rise to new challenges and possibilities. Dong et al. [9] utilized the indoor temperature and CO2 levels, indoor and outdoor RH levels and occupancy state provided by a motion sensor to perform non-linear model predictive controls of HVAC parameters. PIR sensors have proven to reduce energy consumption by 18% [10] in long-term usage. MODES [11] utilizes the fused representation of an array of thermal and vibration sensors to estimate the number of occupants in a given area. Studies in a laboratory environment demonstrated energy savings of up to 77%, with 10% improvement in occupant comfort. By using a low-power joint infrared and optical camera sensor to estimate occupancy and control the HVAC, Cao et al. [12] were able to achieve 26% reduction in energy consumption compared to a schedule-driven control strategy. Researchers have also used other sensing modalities, such as WiFi [13] and vision, to estimate occupancy and perform HVAC actuation. Nagarathinam et al. [14] discuss the possibility of maintaining optimum thermal comfort by modeling the mean radiant temperature of the room at different occupancy levels, and appropriately adjusting each AHU's (Air Handling Unit) operating parameters. They demonstrate an energy savings of 15%. Simulation-based studies [15] suggest that using predicted occupancy values to control HVAC settings can provide an additional 10%–15% reduction in a building's heating and cooling load.

Some early research [16,17] has applied Markov and Regression-based models to control HVAC, taking into account the current and predicted occupancy. Researchers have developed model predictive control (MPC) strategies [6,18,19] to perform real-time HVAC adaptation while adhering to system constraints. Smarra et al. [20] use random forest-based MPC for HVAC control, achieving an impressive energy savings of 49.5%. OBSERVE [21] utilizes optical cameras at intersections to record the transitions of occupants across various rooms on a floor, in turn estimating the number of occupants in every room. A Markov chain model, incorporating such empirical transition probabilities, is used to predict room-level occupancy at different intervals and then control the HVAC settings. Their model reports 42% energy saving.

More recently, reinforcement learning (RL) based approaches have been proposed to support *data-driven* HVAC control. Various machine learning algorithms (e.g., [22–27]) have been proposed to learn the dependencies between various features affecting the thermal comfort of users. Chen et al. [28] developed an RL-based differential MPC policy for HVAC control that reduced average HVAC energy consumption by 16.7%.

"Optimal thermal comfort" is typically a subjective parameter, depending on additional variables such as personal preferences, clothing worn and level of physical activity. Past work (e.g., [29]) has studied how to perform HVAC setpoint regulation to optimize comfort-related metrics in both personal and shared spaces. Most commercial and residential buildings utilize ASHRAE's thermal 55 standards,[3] addressing the conditions for acceptable thermal comfort in buildings, to set a static setpoint temperature. Such settings do not account for the effects generated by changes in occupancy or environmental conditions. Researchers have leveraged data-driven and statistical models [30–32] to predict/estimate an occupant's thermal comfort, and thereby control the HVAC setpoint for *personalized* spaces. Research has also explored [33,34] the use of individual physiological parameters to control the thermostat in a shared space. Recently, researchers have adopted machine-learning techniques [35,36] to model/predict thermal comfort. Most of these techniques target personalizing the thermal comfort of a single occupant in a personal (rather than shared, multi-occupant) space. More recently, [37] developed a model to predict the occupant's thermal comfort based on their arrival time series

[3] https://www.ashrae.org/technical-resources/bookstore/standard-55-thermal-environmental-conditions-for-human-occupancy

data, and utilized simulation studies to adjust the HVAC setpoint in a multi-tenanted space (a mosque building). However, such a study does not (a) perform fine-grained, zone-level HVAC setpoint control in larger shared spaces, or (b) account for the temporal unpredictability in occupancy patterns (e.g., over the week), which makes medium-term occupancy prediction significantly more erroneous.

Despite this extensive body of research using state-of-the-art techniques, the actual deployment of extensive sensor instrumentation in commercial buildings to support smart HVAC control has been muted. Techniques proposed in the literature [21,38] require customized placement of sensors (optical cameras, PIRs) that are often deemed to be impractical or unattractive. Most buildings incorporate the placement of sensors (PIR, temperature, and optical cameras for surveillance) in their design. Addition/modification to the placement of these sensors will require tremendous effort. Thus, we propose *OcAPO* that utilize the sensors in the infrastructure (without the need to add/modify the placements) to estimate occupancy and yet control the PDC/HVAC. Moreover, past work principally addresses the issue of HVAC control at the granularity of an entire room. However, modern open-floor co-working spaces (such as our testbed) are divided into sub-areas (called zones) that require zone-level occupancy estimates and HVAC control. Moreover, as we shall discuss in Section 4.2, the open floor nature of spaces implies a coupling between the temperature values of nearby zones—this must be factored in while devising the appropriate HVAC/PDC control strategies.

The real-world university space that we shall experimentally study, primarily use motion and temperature sensing input as part of a *rule-based* PDC control policy. In our rule-based deployment, the PDC setpoint is set a static value of 24 °C, with the PDC valve closing when no motion is found for 30 min and re-opening whenever the motion is subsequently detected. Motion sensors, however not only often fail to detect stationary humans but also provide only binary occupancy status—this significantly limits the efficacy of the PDC systems. To overcome these shortcomings, we shall adopt a data-driven, non-binary occupancy-based PDC control logic.

3. TestBed and data collection

We first describe our target operating space and the data collection mechanism for PDC setpoint control.

3.1. TestBed — collaborative zone

Our testbed is a collaborative shared workspace of area $1000\ m^2$ used by students of our university campus. It is logically divided into 8 zones virtually represented as 3 different areas, with no physical demarcation between them (Fig. 2). A building management system (BMS) manages each zone equipped with PDC units for cool air circulation at the ceiling level. The PDC relies on the natural convection process of pushing the hot air up towards the ceiling generated by the occupants OR other heat-generating equipment/devices in the space while pushing down the cold air towards the ground level. The cold air is driven by the PDC valves, which are opened between 0%–100% based on an internal PLC-based PID (proportional integral derivative) controller. The valve behavior depends on the difference between the ceiling-mounted (room) temperature sensor reading and the setpoint of each zone: the valves open if the room temperature is above the PDC setpoint, and close if the room temperature drops below this setpoint. Because this is a real-time operational building, we are constrained to adjust the PDC setpoint between 22 °C and 25 °C and do not have any direct control over the opening/closing of PDC valves. In our testbed, PDC valves open up during the operational hours of the building (8.30 AM to 10.30 PM on weekdays). The testbed is characterized by a single ceiling-mounted temperature sensor providing *nominal* temperature readings for each area, comprising multiple zones: Area-1 consisting of Z15, Z16 and Z17 (sensor mounted on Z16), Area-2 consisting of Z7, Z8 and Z9 (sensor mounted on Z8) and Area-3 consisting of Z11 and Z12 (sensor mounted on Z12).

3.2. Data collection

Our testbed is equipped with Beckhoff PLC and is interfaced with our research servers using the OPC-UA interface (as illustrated in Fig. 12). We subscribe to the relevant OPC tags (room temperature (RTTS), PDC setpoint value (CRSP), and valve control (VALC)); the PLC notifies our server of updated readings whenever the value of one of these tags changes. The existing infrastructure does not record the fine-grained changes at the zone or occupant level. Hence, we collect fine-grained table-level (ambient) temperature data for each zone by placing two tripod-mounted Texas instrument BLE sensor tags per zone.[4] Each sensor is connected to a Raspberry Pi device and transmits temperature data once every minute via Bluetooth. We then averaged the temperature from 2 sensors of each zone to obtain the zone-level *ambient* temperature data (as illustrated in Fig. 3). We collected data intermittently over two distinct phases. The first phase comprises data collected between February–March 2022 for measuring the temperature with different occupancy levels and May–June 2022 during the holiday season when there were no occupants. During this phase, we manually controlled the setpoint of different regions to vary over a wider range (22.5–25°), but in coarser 0.5° increments. We used this first phase to evaluate and quantify the effect of various parameters, such as occupancy and weather, on occupant thermal comfort. We utilized these insights to develop a trace-based model to identify the collective PDC setpoints for the zones in a given area. The second phase involved collection of data during the period September–October 2023 (15 days). During this phase, we used the insights from our initial studies to investigate optimal setpoint settings over a tighter range (23–24.2°), but at finer granularity (increments of 0.2°). This additional data was used to build a machine-learning based regressor model for setpoint determination.

[4] https://www.ti.com/tool/TIDC-CC2650STK-SENSORTAG

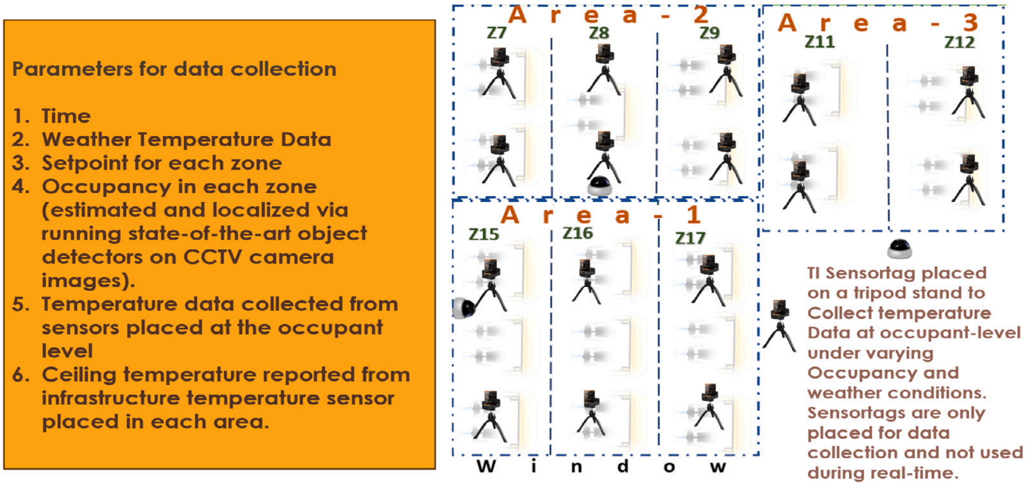

Fig. 3. Data collection setup and parameters collected.

During the first phase, we collected the room temperature data to understand (i) the effect of varying weather conditions, (ii) the effect of varying occupancy, and (iii) the effect of PDC setpoint modification on each zone and its neighboring zones. To evaluate the effect of different parameters, we vary the setpoints at different times of the day (Morning, Afternoon, and Evening). While we had the liberty of varying all the PDCs concurrently when the space was unoccupied, we had to be mindful of preserving occupant comfort whenever the space was occupied. Hence, we always kept one zone in each area at 24 °C while varying the PDC setpoint of the other zones in each area between 22.5 °C and 25 °C. For each area, we collected the temperature data for different combinations of zone setpoints and under different occupancy levels, with every zone experiencing a temperature between 22.5 – 25 °C. However, due to the study period overlapping with COVID-related restrictions and greater adoption of remote working, we did not experience high occupancy in the morning and evening times. To understand the phenomena of change in room temperature for varying weather conditions, we also crawled the temperature data for the weather station closest to the academic building from the national (Singapore) real-time weather database.[5] To estimate the occupancy, we use feeds from existing ceiling-mounted cameras to count the number of people in each zone. We use an object detector (YOLOV3) to compute bounding boxes across persons detected in the scene. We extract the object's pixel coordinates (Xmin, Ymin, Xmax, Ymax, Xcenter, Ycenter) and localize the bounding boxes to the zone's physical coordinates using a logistic regression.

Our learning from the first phase helped us identify a much narrower range within which PDC setpoint can be varied, for individual regions, without dramatically affecting occupant thermal comfort. Thus, in the second segment (Sep–Oct 2023), we adjusted the PDC setpoint at an increment of 0.2 °C, with the range varying between (a) 23.0–24.2 °C for Area-2 and Area-3, and (b) 23.6–24.2 °C for Area-1. As in the first phase, we collected the weather temperature, occupancy level, and occupant level temperature data from temperature sensors placed at the occupant level.

4. Empirical findings on PDC setpoint control

We now discuss several key findings on how occupancy and ceiling-based sensor readings affect occupant thermal comfort.

4.1. Impact of occupancy and weather conditions

With the data obtained from the first phase, we individually evaluate the impact of occupancy and weather conditions on the room temperature (measured at both the ceiling and table/occupant levels). Although it is well-established that, in principle, occupancy and weather impact the room temperature, we conduct a fine-grained analysis of how the ambient (occupant-level) temperature is affected by these factors. We first find all the dates and timeslots where the PDC setpoints remained constant, for a given area, throughout the day. We then bin the weather conditions and occupancy values by (a) quantizing occupancy into 4 distinct levels: {*No*: 0; *Low*:1–3; *Med*:4–7; *High*: >8}, and (b) quantizing external temperature into 6 bins, {22:22.0-24.0; 24:24.1-26.0; 26:26.1-28.0; 28:28.1-30.0; 30:30.1-32.0; 32:32.1-34.0}, which collectively capture the climatic variation of temperature observed. For comparative analysis, we first identify similar weather conditions across the four months (Feb–Mar and May–Jun 2022) under zero occupancy, with the observation period classified into Morning (8–13:00 h), Afternoon (13–18:00 h), and Evening (18–23:00 h).

[5] https://data.gov.sg/dataset/realtime-weather-readings

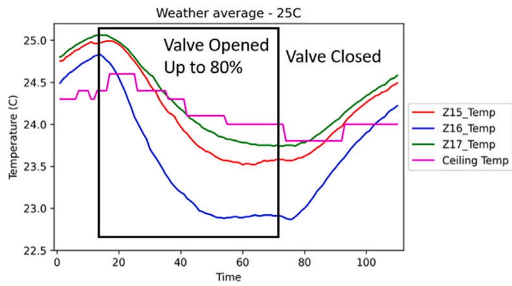

Fig. 4. Weather average temperature — 25C.

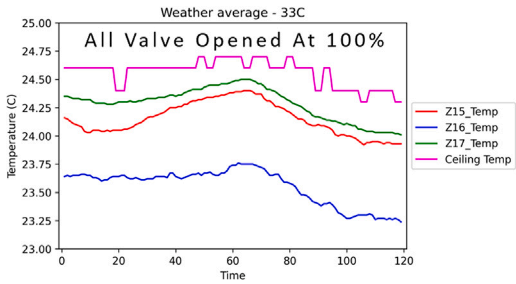

Fig. 5. Weather average temperature — 33C.

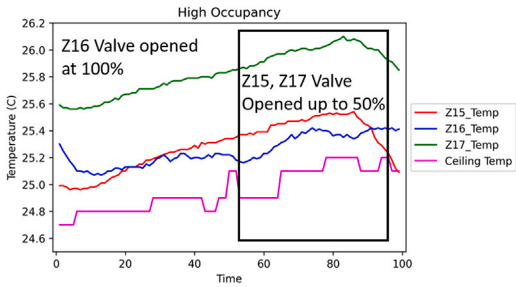

Fig. 6. Room temperature during high occupancy.

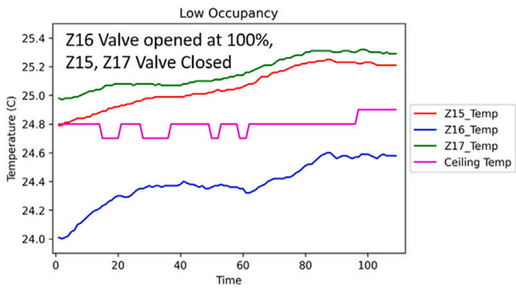

Fig. 7. Room temperature during low occupancy.

We track both the zone-level ambient (measured using BLE-equipped TI temperature sensors) and the ceiling-level temperature values. For exposition, we use a representative combined PDC setpoint value of (25, 24, 25) for Area-1 (Z15-16-17). Figs. 4 and 5 illustrate the recorded temperature for two different weather buckets, namely, ∼ 25 °C and ∼ 33 °C, with this area completely unoccupied. We observe that the average ceiling temperature recorded at 25 °C is ∼ 24.1 °C, and the valves were only open for an average duration of 70% (varying between 1%–80%). At 33 °C, the ceiling temperature rises up to 24.5 °C, and the valves remain open throughout at 100%. While the ambient temperature was significantly different for external temperatures of 24 °C and 32 °C, the differences were marginal between 22° and 28 °C. Figs. 6 and 7 illustrate the variation in room temperature for two different

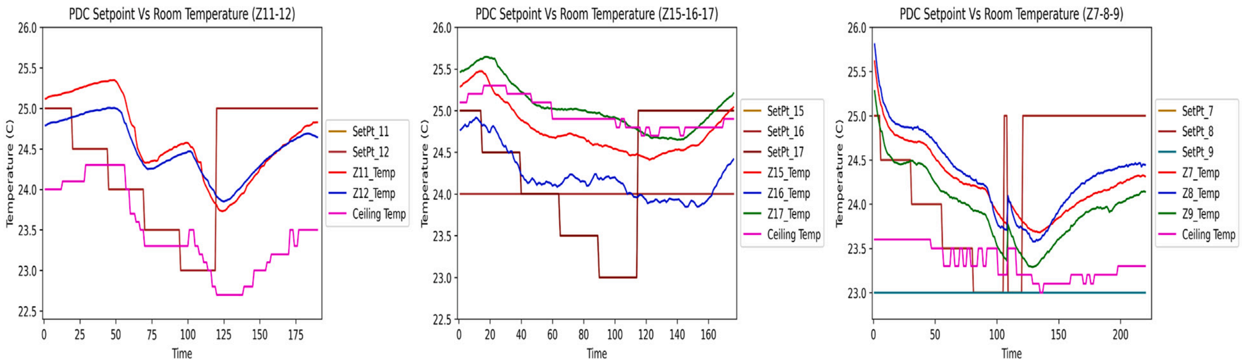

Fig. 8. Recorded room temperature for varying PDC setpoints.

occupancy conditions, using the setpoint temperature setting of (25, 24, 25) for the three zones Z15-Z16-Z17. We observe that the average ceiling temperature for low occupancy was recorded to be 25.0 °C, while for high occupancy, the ceiling temperature remained at 24.7 °C (as valves opened up due to high temperature). We also note that the ambient zone temperature recorded under low occupancy is significantly lower than that at high occupancy. The valves for Z15 and Z17 remained shut throughout the low occupancy, while for high occupancy, the valves opened up to 50% as the ceiling temperature increased up to 25.2 °C. The ambient level temperature reduces as the valves open up under high occupancy.

In our testbed, we observed high occupancy only during the middle of the day (see the trend in Fig. 9), when the external temperature is also high. The occupancy is observed to be low only in the mornings and at dawn. Tropical climates such as Singapore do not experience very low temperatures. Moreover, the building is closed during hotter summer months. Accordingly, due to insufficient observational data, we did not study the behavior of the PDC control under a combination of high occupancy and low external temperature.

Similar qualitative trends were observed in the next phase of our empirical studies. During this phase, we did observe a general increase in occupancy (across all zones), compared to the first phase. This can be attributed to the elimination of COVID restrictions and the normalization of regular workplace patterns. Compared to the initial phase, Area-1 thus observed a significantly higher occupancy of students in the second phase—we shall discuss this further in Section 6.

4.2. Impact of PDC setpoint on ceiling and zone temperature

With the objective of understanding the coupled effect of PDC setpoint values across adjacent zones, we simultaneously adjusted the PDC setpoint of two zones (Z15 & Z17 for Area-1, and Z7 & Z8 for Area-2) in decrements of 0.5 °C every 30 min, starting from 25C, and subsequently reversing the steps (in 0.5 °C increments) back to 25 °C. Because the experiment was performed in an occupied setting, we maintained a consistent temperature of 24C for Z16 in Area-1 and Z9 in Area-2. For Area-3, we varied the temperature for both zones. All the experiments were performed under high occupancy conditions. Fig. 8 illustrates the variation in room temperature (ceiling and zone level) for all three areas. We list our observations as follows:

- **Gradual change in zone temperature:** A change in PDC setpoint starts to manifest an impact on occupant-level temperature values after a gap of ∼20 min. However, it can take longer if the setpoint temperature is higher than the room temperature. For instance, as observed in Area-3 (Z11-12), the temperature drops only after the setpoint temperature was set at 24.0 °C. This is unlike the phenomena illustrated in Area-1 (Z15-16-17), where the temperature drops after 20 min immediately after the setpoint temperature is reduced from 25 °C to 24.5 °C. This is due to the ceiling level temperature recorded and compared with the PLC logic.
- **Impact on Nearby Zones:** For Area-1, we observe that the ambient temperature for zone Z16 reaches the optimum temperature of 24 °C only when we lower the PDC setpoints of the other two zones. The zone level temperature of Z15 and Z17 does not reach the optimum temperature in this experiment. However, we observe that the temperature reaches 23.5°–24 °C if the zone's setpoint is retained at a lower temperature for a longer duration. For Area-3, we observe that occupant level temperature drops to the optimum level of 24.0 °C when both zones are set at 23 °C. The data reveals the non-trivial effects of thermal coupling between zones, which motivates us to adopt a strategy of adjusting zone setpoints *collectively*, rather than individually.
- **Varying relation between ceiling and zone temperature:** The ceiling temperature for Area-2 (Z7-8-9), and Area-3 (Z11-12) is significantly lower than the occupant-level temperature. We hypothesize that these phenomena, observed consistently across multiple areas, are because (a) the density of occupants in the two areas is, on average lower by 70% compared to Area-1, and (b) the ceiling temperature sensor is mounted on one of the unoccupied zones of both areas, and thus unable to accurately capture the heat generated within the occupied zone. As developing a precise mathematical model relating ambient to ceiling-level zone temperature seems very challenging and would require significantly more data, we instead design *OcAPO* to adopt a more empirical data-driven approach.

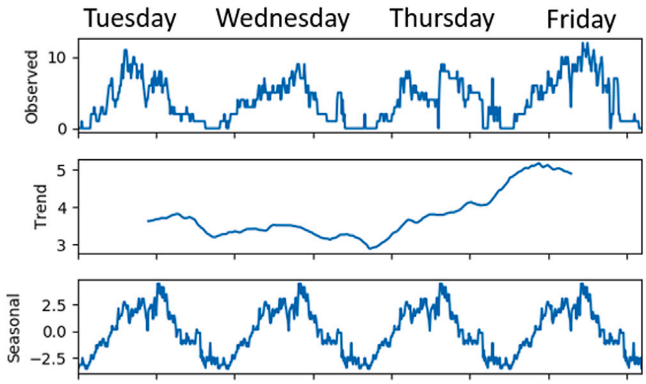

Fig. 9. Diurnal occupancy trend (1 week, Area 1).

5. Proposed PDC setpoint control

Ideally, PDC setpoint control should use a time series model that factors in the time-varying occupancy levels and the gradual evolution in the room temperature data. However, building such a supervised time series model requires a significant amount of data, capturing fine-grained variations in ambient room temperature, external weather, and occupancy.

Data-Driven Lookup: For practical use, we use the data collected in the first phase to develop a trace-based model that uses the observed testbed data to build a lookup table that captures the first-order relationship between these attributes. We then use diurnal models of occupancy and real-time occupancy estimates, obtained by applying object detection vision DNNs over camera data (images captured via CCTV cameras in the infrastructure), together with a *comfort cost function* (CCF), to determine the optimal tuple of setpoints over small (15 min) time intervals.

ML-Based KNN Regressor: With the finer-grained data collected in the second phase, we also developed an efficient KNN Regressor model (without having to lookup the database) that uses machine learning to set PDC setpoints of all the zones in an area given real-time occupancy estimates, weather conditions, and required occupant-level ambient temperature.

5.1. Short-term occupancy prediction

As per our analysis reported in Section 4.2, it takes ∼20 min to cool an area. Accordingly, it is ideal to tune the PDC setpoint every 20 min windows (or epochs) and perform one-epoch lookahead occupancy prediction (i.e., 20 min ahead of time). However, predicting such coarse-grained zone-level occupancy using historical data at 15–20 min temporal granularity is challenging due to unpredictable fluctuations both in (a) the number of occupants in each zone, (b) their arrival time and stay duration. However, we do observe a simple diurnal trend for each zone across different day segments—morning (8 am–1 pm), afternoon (1pm–3.30; 3.30–6.30) and evening (6.30–10.30 pm). For example, Fig. 9 plots the diurnal occupancy over one week for Area-1 (comprising zones Z15, Z16 and Z17). In addition, for our analysis and model building, we use the 4 distinct occupancy levels {No, Low, Medium and High} explained earlier. We see that while no clear short-term trend (defined by the divergence from a moving average value) is visible, the occupancy typically is Low-Medium in the morning, High in the afternoon, and Low in the evening. Accordingly, at the start of daily segment, we use the observed diurnal value; for all subsequent 20-min intervals in that segment, we approximate the predicted one-step lookahead occupancy based on the current epoch's real-time occupancy estimate (obtained from the camera feeds).

5.2. Lookup table for trace-based model

Using the data collected over 60 days, we populate a lookup table that helps compute the optimum combination of PDC setpoints that help maintain occupant thermal comfort (23.5 °C–24 °C). As observed, a zone's temperature decreases ∼20 min after a reduction in PDC setpoint, and vice versa. The table (illustrated in Fig. 10) consists of different combinations of PDC setpoints (SP), weather data, ceiling temperature, and ambient table-level temperature data recorded by BLE-equipped temperature sensors (S_15, S_16, S_17). As described in Section 4.1, the weather values are quantized into 6 bins; for each combination of PDC setpoints, our lookup table consists of 20 min data. While a 3-zone area conceptually contains $3^6 = 81$ unique tuples (each setpoint can have 6 values between (22.5 °C, 25 °C)), our training data was collected typically using 10 different combinations, which proved to be sufficient for effective thermal comfort.

5.3. Real-time setpoint control: TraceBased model

We utilize the lookup table entries, together with input values for current weather and predicted occupancy (which together serve as a table index), to determine the "optimal" setpoint control values. For occupancy estimation, a YoLov3 object detector [39]

Look Up Table For Area-1 (Zone-15, 16, 17)

	Input	Output			Input			Input			
Time	Weather	SP_1	SP_2	SP_3	O_1	O_2	O_3	S_15	S_16	S_17	CeilingTemp
T1	28	24.5	24	24	LOW	HIGH	HIGH	24.8	24.2	24.8	24.7
T10	28	24.5	24	24	LOW	HIGH	HIGH	24.5	23.8	24.5	24.3
T20	28	24.5	24	24	LOW	HIGH	HIGH	24.1	23.7	24.2	24.0

- **Weather** = Outside Weather Temperature in Celsius
- **SP_1** = Setpoint-1 (Zone-15); **SP_2** = Setpoint-1 (Zone-16); **SP_3** = Setpoint-1 (Zone-16);
- **O_1** = Occupancy for Zone-15; **O_2 =** Occupancy for Zone-16; **O_3 =** Occupancy for Zone-17
- **S_15, S_16, S_17** = TI Sensor Tag Zone Temperature Values for Zone15, 16 and 17

Fig. 10. Look-up table.

is executed on the camera images to obtain bounding boxes for human objects, which are then translated into real-world physical coordinates. For each matching (weather, occupancy) entry, we compute the CCF value by computing the cumulative deviation from the ideal ambient temperature (23.5 °C for occupied; 25 °C for unoccupied). Algorithm 1 details the CCF computation. Intuitively, the setpoint choice seeks to minimize the weighted sum of any deviations from the ideal ambient temperature. To emphasize user comfort, we apply a higher penalty (weight=2) if the projected temperature is higher than the desired value (i.e., if users will feel warmer); to concurrently prevent unnecessary cooling, we apply a lower penalty (weight=0.5) if the projected temperature is lower than the desired value.

Algorithm 1 Trace-Based Lookup Algorithm

Input: Weather temperature (w), and Estimated Occupancy (o) for next 15-min epoch.
for j ∈ All Combination of PDC Setpoints (Weather=w,Occup=o) **do**
  for Minutes 1..15 **do**
    calculate CCF(j)=

$$\sum_{k \in zones} (2 * (23.5 - z_k^j) * \mathrm{I}(z_k^j > 23.5) \& \mathrm{I}(o_k^j \neq Unoccupied) +$$
$$0.5 * (z_k^j - 23.5) * \mathrm{I}(z_k^j < 23.5) \& * \mathrm{I}(o_k^j \neq Unoccupied) +$$
$$0.5 * |z_k^j - 25.0| * \mathrm{I}(o_k^j = Unoccupied)$$

    where z_k^j represents the predicted ambient temperature of the k^{th} zone under the j^{th} setpoint setting & I(…) represents the Indicator function.
  end for
end for
Find the combination which has the lowest cost—i.e., $\arg\min_j CCF(j)$

5.4. KNN regression model

K-Nearest Neighbor (KNN) is a supervised machine-learning model to learn the non-linear relationship between the dependent and independent variables. KNN regression models identify the nearest neighborhood points of a data point and clusters them together; the KNN predictor uses the average value of the dependent variable across all these neighborhood points to predict its output. We now detail how we trained the KNN-based regressor model, using the data collected across both phases, and how this model is used to provide real-time estimates of preferred PDC setpoint values.

- **Preparing data for model**: To prepare the data for training the model, we extracted relevant data (occupancy level = {LOW|MED|HIGH}), weather temperature, occupant level temperature sensor values from both the first and second phases for different combinations of PDC setpoints, as specified in Section 3.2. The training dataset consists of a combination of setpoints for a specific area (Area 1/2/3) in the testbed across diverse occupancy scenarios. An exemplar set of data points for a particular PDC combination is depicted in Fig. 11. Notably, the independent variables include "occupancy for each zone, weather temperature, and MIN and MAX levels for required occupant temperature in each zone", while the dependent variables encompass "PDC setpoint for each zone". To account for occupant comfort, two values are considered for occupant level temperature, namely "minimum" (c_min) and "maximum" (c_max). This ensures that the model output stays within the specified occupant comfort range. For the minimum zone level temperature value, we average the sensor temperature value of two sensors from the zone. For the maximum, we assign an occupancy-dependent value as follows: $\{NO : 24.5\}; \{LOW : 24.2\}; \{MED : 24.0\}; \{HIGH : 24.0\}$. These settings reflect the general intuition that the maximum permitted temperature range should decrease with increased occupancy, to reflect the additional anthropogenic heat generated.

Independent Variables											OUTPUT		
Weather	Occ_15	Occ_16	Occ_17	Valve_15	Valve_16	Valve_17	S_15	S_16	S_17	Ceiling_Temp (Area-1)	SetPt_15	SetPt_16	SetPt_17
30.2	0	4	1	100	100	100	24.15855	23.67516	24.13338	24.29999924	24.2	24	24.2
30.2	0	3	1	100	100	100	24.15855	23.68774	24.12079	24.29999924	24.2	24	24.2
30.2	0	4	1	100	100	100	24.151	23.70285	24.12582	24.29999924	24.2	24	24.2
30.2	0	4	1	100	100	100	24.14345	23.70789	24.11324	24.29999924	24.2	24	24.2
30.2	0	4	1	100	100	100	24.13841	23.70537	24.09813	24.29999924	24.2	24	24.2
30.2	0	4	1	100	100	100	24.13589	23.71292	24.09058	24.29999924	24.2	24	24.2
30.2	0	4	1	100	100	100	24.13338	23.71796	24.09309	24.29999924	24.2	24	24.2
30.2	0	4	1	100	100	100	24.13086	23.72299	24.10065	24.29999924	24.2	24	24.2
30.2	0	3	1	100	100	100	24.13589	23.71544	24.10065	24.29999924	24.2	24	24.2
30.2	0	3	1	100	100	100	24.13086	23.70789	24.10568	24.29999924	24.2	24	24.2
30.2	0	2	1	100	100	100	24.12582	23.70789	24.09813	24.29999924	24.2	24	24.2
30.2	0	3	1	100	100	100	24.12079	23.70285	24.09813	24.29999924	24.2	24	24.2
30.2	0	2	1	100	100	100	24.1082	23.68271	24.10065	24.29999924	24.2	24	24.2
30.2	0	2	1	100	100	100	24.10316	23.68774	24.09813	24.29999924	24.2	24	24.2
30.2	0	3	1	100	100	100	24.09561	23.68774	24.09058	24.29999924	24.2	24	24.2
30.2	0	3	1	99	99	99	24.09561	23.67264	24.08554	24.29999924	24.2	24	24.2
30.2	0	2	5	98	100	98	24.09058	23.66257	24.06792	24.10000038	24.2	24	24.2
30.2	0	3	4	97	100	97	24.08302	23.6248	24.07043	24.10000038	24.2	24	24.2
30.2	0	3	2	96	100	96	24.08806	23.61221	24.07043	24.10000038	24.2	24	24.2
30.2	0	2	2	95	100	95	24.08554	23.60466	24.07043	24.10000038	24.2	24	24.2

- **Weather** = Outside Weather Temperature in Celsius
- **SetPt_15** = Setpoint-1 (Zone-15); **SetPt_16** = Setpoint-2 (Zone-16); **SetPt_17** = Setpoint-3 (Zone-17);
- **Occ_1** = Occupancy for Zone-15; **Occ_2 =** Occupancy for Zone-16; **Occ_3 =** Occupancy for Zone-17
- **S_15, S_16, S_17** = TI Sensor Tag Zone Temperature Values for Zone15, 16 and 17

Fig. 11. Training data for KNN regressor.

- **Model Training**: We train the KNN Regressor model using the data we collected from two phases. The regression model learns the non-linear relationship between the independent variables and the PDC setpoints for each zone. The KNN algorithm clusters the K nearby data points together and computes their average for PDC setpoint value estimation. During the training process, we found the mean square error to be the lowest when we selected K=2. Thus, the model was trained to average the PDC setpoint values for the nearest 2 neighbors.

5.5. Real-time setpoint control: KNN regression

OcAPO uses the trained KNN regression model to estimate the optimal PDC setpoint in real time. As with the trace-based model, *OcAPO* obtains real-time estimates of the number of using the YOLOv3 object detector and obtains the current weather temperature data from the Singapore weather API. As for the minimum and maximum occupant-level temperature, *OcAPO* follows a static rule-based assignment as discussed in Algorithm 2. Due to the inaccuracies of the object detector, the occupancy levels may change sporadically within the 20-minute interval. Thus, *OcAPO* performs statistical smoothing of the occupancy levels extracted per frame.

6. Experimental results

We now quantify the performance of both the proposed versions of *OcAPO* (trace-based and KNN Regressor). We test our approach in an uncontrolled (i.e., occupancy driven by natural usage) testbed setting. We compare our system to the baseline rule-driven, time-of-day based "BMS" (Building Management System) system. All three approaches were evaluated over 3 consecutive days (Wed, Thu, Fri) on separate weeks. Our evaluation focuses on (a) *OcAPO's* ability to maintain optimal thermal comfort during the deployment, (b) dynamic PDC setpoint control of *OcAPO* compared to the currently operating baseline, (c) Evaluation of PDC valve opening percentage under all three settings (BMS, *OcAPO* trace-based and *OcAPO* KNN regression, and (d) Occupancy Estimation errors.

6.1. Experimental setup

OcAPO composes of different modules as illustrated in Fig. 12. The CCTV camera images are captured at a rate of 3fps from the cameras deployed in the testbed (Collaborative zone) using a third-party interface. *OcAPO* retrieves images from 6 cameras that monitor the 8 zones of interest and runs one thread per camera (6 threads in total) to execute the object detection on the images. The bounding boxes resulting from the object detection are localized using a logistic regressor and sent via Kafka stream to the PDC Setpoint controller. The controller is composed of 3 modules. The "Occupancy Estimation" module performs statistical smoothing for each zone's occupancy. The estimated occupancy, weather information and required PDC setpoint for each zone are given as

Algorithm 2 KNN Regression Algorithm

```
Procedure KNN_Regression(KNNTrainedModel)
for Minutes 1..15 do
    Get Weather temperature (w), Estimated Occupancy (per frame) (o)
    PerFrameOccupancy = FrameNumber : EstimatedOccupancy
end for
Occ_Periodic = CALL Statistical Smoothing(PerFrameOccupancy)
ZoneThermalComfort = CALL Assign Thermal Comfort(Occ_Periodic)
PDCSetpoint_Per_Zone = KNNTrainedModel.fit(Weather, Occ_Periodic)
End Procedure
Procedure Statistical_Smoothing (PerFrameOccupancy)
for zone,Count ∈ Per-Frame-Occupancy do
    TotalCount = TotalCount + Count
    TotalFrame = Frame + 1
end for
Total_Occupancy = TotalCount/TotalFrame
if
Total_Occupancy == 0 then
    Occ_Periodic = NO
else if
Total_Occupancy > 0and <= 2 then
    Occ_Periodic = LOW
else if
Total_Occupancy > 2and <= 5 then
    Occ_Periodic = MED
else if
Total_Occupancy > 5 then
    Occ_Periodic = HIGH
end if
End Procedure
Procedure Assign Thermal Comfort (Occ_Periodic)
for Occ ∈ Occ_Periodic do
    if Occ == NO then
        ZoneThermalComfort_min = 23.7;ZoneThermalComfort_max = 24.5
    else if Occ == LOW then
        ZoneThermalComfort_min = 23.7;ZoneThermalComfort_max = 24.2
    else if Occ == MED or Occ == HIGH then
        ZoneThermalComfort_min = 23.7;ZoneThermalComfort_max = 24.0
    end if
end for
End Procedure
```

input to "find optimal PDC setpoint" for each zone. *OcAPO* interfaces with BMS via the "Kepware OPC interface". *OcAPO* acts as the OPC client and writes the PDC setpoint value to the OPC tag. The PLC controller in the BMS system compares the new PDC setpoint with the room temperature to open and close the PDC valves.

6.2. Ambient temperature control: OcAPO (trace-Based and KNN regressor) vs BMS (static setpoint)

We evaluate *OcAPO's* ability to maintain thermal comfort (more precisely, occupant-level ambient temperature) across varying occupancy levels. More specifically, we compare the temperatures recorded at each zone via the temperature sensors while executing the two proposed methodologies (Trace Based and KNN Regressor) with the BMS static setpoint. We plot the zone, ceiling temperatures, and PDC setpoint across low and high occupancy to compare the performances of the methodologies under varying occupancy levels. The observations are listed below:

- **Area-3:** The PDC setpoint regulated by *OcAPO* and BMS for Area-3 is shown in Fig. 13 for high occupancy and in Fig. 14 for low occupancy. The figures also plot the concurrently-measured ceiling and zone-level temperature values. As noted, whether the occupancy is high or low, the BMS setpoint consistently exceeds the standard ambient value temperature value of 24 °C at the zone level by an average of 1.2 °C. This is due to the discrepancy observed in the ceiling and zone-level temperature. Although the zone temperatures sometimes exceed 25 °C, the ceiling level temperature readings lie within 23.5 (23.5 °C, 24 °C). As these

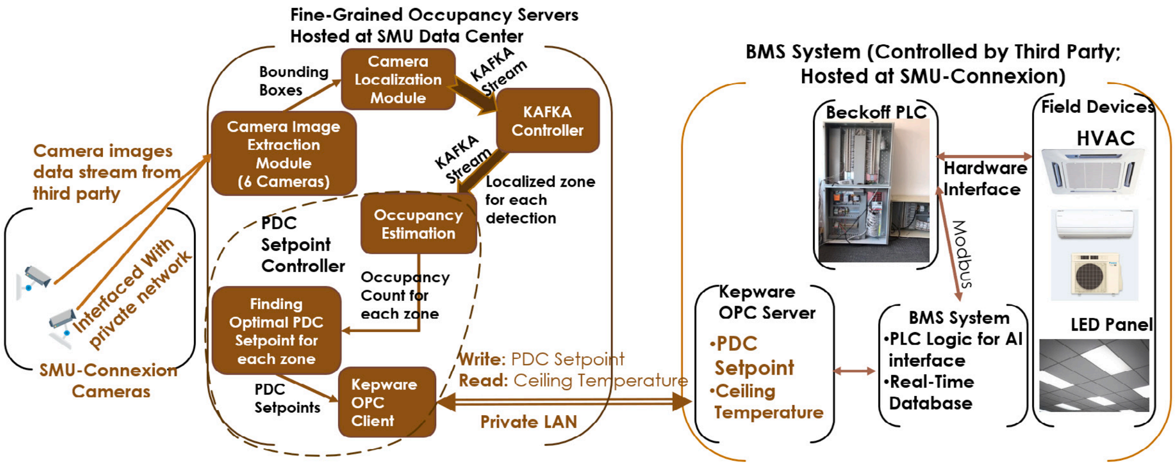

Fig. 12. Modules in *OcAPO* and interaction with BMS (Building Management System).

readings do not exceed the BMS-specified setpoint of 24 °C, the PLC logic does not open the PDC valves, thus resulting in higher ambient temperature.
On the contrary, *OcAPO's* Trace-Based approach results in an average temperature of 24.4 °C in Area-3 under low occupancy and 24.5 °C under high occupancy. The KNN regression-based variant of *OcAPO* is equally successful in maintaining the occupant level temperature at 24.3 °C under both low and high occupancy. The KNN-based approach benefits from the finer-grained observational data collected, and is able to ensure a tighter control on occupant-level ambient temperature by specifying finer-grained PDC setpoints across all occupancy conditions (No, Low and High).

- **Area-2:** Similar to Area-3, BMS deviates from the desired temperature (24 °C) by 0.9 °C in Area-2. The average zone temperature when BMS was in operation is noted to be 24.9 °C under high occupancy and 25.0 °C under low occupancy. During periods of low occupancy, the PDC valves remained closed, leading to a higher ambient temperature. In contrast, during high occupancy, the PDC valves opened, allowing the dissipation of heat generated by humans. As the concentration of heat increases, the ceiling temperature rises, prompting the opening of the PDC valves. As shown in Fig. 15, the zone temperature for BMS setting under high occupancy did drop to the optimum of 24 °C for a brief period as the PDC valves opened up due to high heat from the occupants. However, the zone temperature starts to rise again within 5 min. In contrast, as observed in Figs. 15 and 16, the Trace-Based variant of *OcAPO* maintains the zone temperature at an average of 24.1 °C under low and 24.4 °C under high occupancy. The KNN-based *OcAPO* variant achieves even tighter temperature control, with average ambient values of 24.0 °C and 24.1 °C for low and high occupancy, respectively.
- **Area-1:** As observed from Figs. 17 and 18, in Area-1, the discrepancy between the ceiling and zone level temperature is less, implying a higher degree of effectiveness for BMS-based control. The average zone temperature maintained by BMS is 24.4 °C for low occupancy, rising a bit to 24.9 °C under high occupancy. Trace-based *OcAPO* maintains the average occupant level temperature at 24.0 °C under low occupancy and 24.4 °C under high occupancy. Similar to the other zones, the KNN-based version of *OcAPO* provides finer-grained temperature control, achieving an average ambient temperature of 24.1 °C under both low and high occupancy.
- Given the dynamic nature of usage of the space, where individual occupants and groups occupied the space for widely varying time duration, we did not conduct an extensive user study of perceived comfort—i.e., we did not explicitly survey individual occupants to obtain their feedback on thermal comfort levels. Instead, signage was placed throughout the testbed, whenever *OcAPO* was deployed, encouraging the occupants to contact us via email if they experienced discomfort. Throughout the deployment phase, we received no complaints from occupants, and can thus conclude that our PDC settings did not cause noticeable discomfort.

6.3. Dynamic PDC setpoint control

We also studied the behavior of dynamic PDC setpoint control (setting varying setpoints across zones for a given area) compared to the occupancy-unaware static BMS control. In all three areas, the BMS PDC setpoint value is static at 24 °C.

- For Area 3, it is noted that Traced-based *OcAPO* sets the PDC setpoint value at 23 °C for both low and high occupancy. The lack of differentiation in setpoint settings can be ascribed to the coarser-grained nature of our training data (captured in 0.5 °C increments). *OcAPO's* KNN-based regressor, however, performs more differentiated control: for high occupancy scenarios, the average PDC setpoint is specified at 23 °C, whereas for low occupancy, the average setpoint value is specified to be 23.4 °C. While the zones are unoccupied, the trace-based version sets the temperature of PDC setpoints at 25 °C, while the KNN iteration determines the 'optimal' temperature to be 24.6 °C. This subtle change results in the maintenance of better occupant thermal, as observed in the earlier sections.

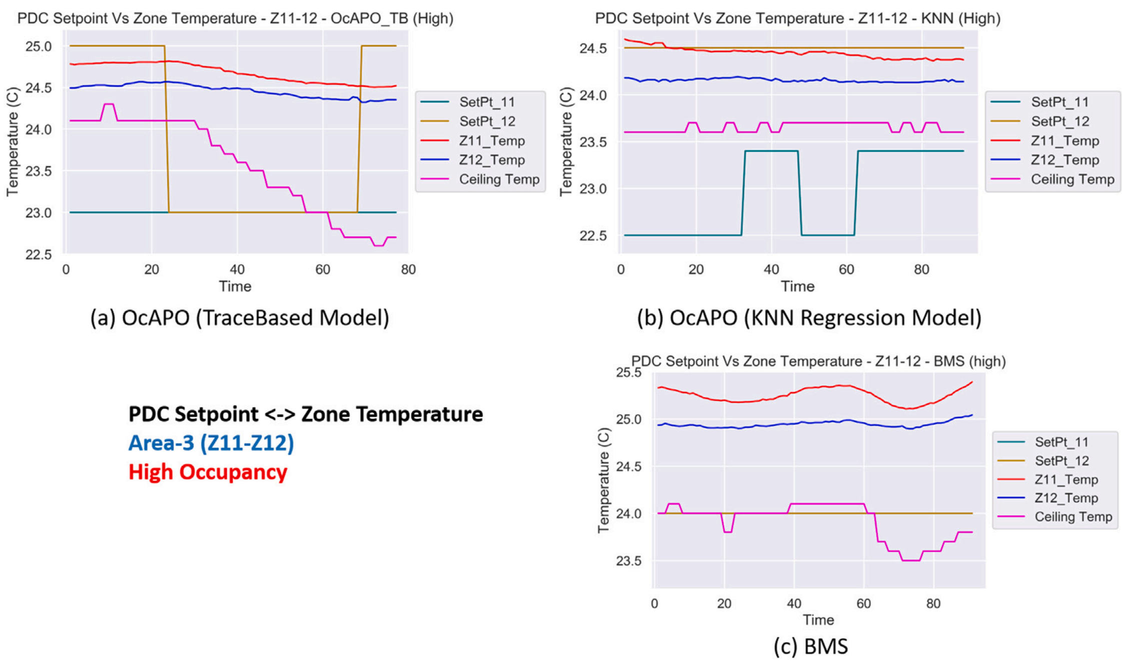

Fig. 13. PDC setPoint by BMS, OcAPO (TraceBased and KNN) Vs ceiling and zone temperature — high occupancy (Area-3).

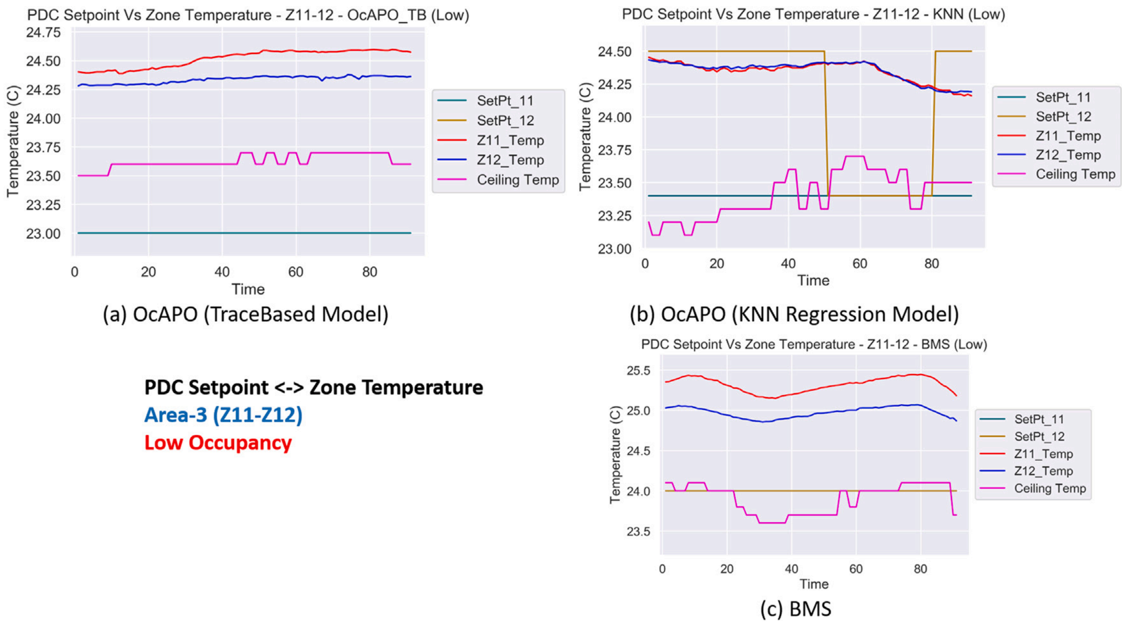

Fig. 14. PDC setPoint by BMS, OcAPO (TraceBased and KNN) Vs ceiling and zone temperature — high occupancy (Area-3).

- *OcAPO's* Trace-based version behaves similarly in Area-2. The setpoint specifications under low occupancy have Mode= 23 °C (with a maximum value of 23.5 °C); and under high occupancy, the setpoint distribution has Mode=22.5 °C (with a maximum value of 23 °C). *OcAPO's* KNN version results in a PDC setpoint distribution with Mode=23.4 °C (Max=24 °C) under high occupancy, and Mode=23.4 °C (Max=24.6 °C) under low occupancy.
- In Area-1, when all zones are occupied, KNN version of *OcAPO* deviates from the standard setpoint of 24 °C by setting the PDC setpoint 0.2 °C higher in the low occupancy (Mode: 24.2 °C) and 0.2 °C lower under high occupancy (Mode: 23.8 °C). However, if only one zone is occupied, the KNN version sets the zone at 24 °C, with the unoccupied zone PDC setpoint at 24.2 °C. The trace-based version sets the PDC setpoint at an average of 24 °C under low occupancy and 23.5 °C under high

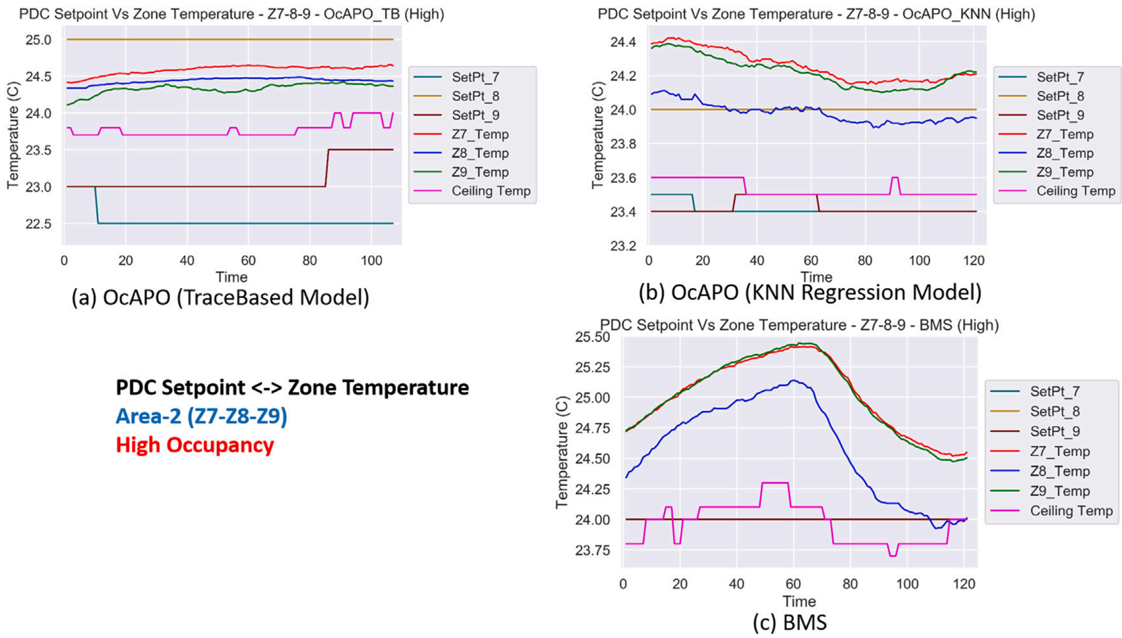

Fig. 15. PDC setPoint by BMS, OcAPO (TraceBased and KNN) Vs ceiling and zone temperature — high occupancy (Area-2).

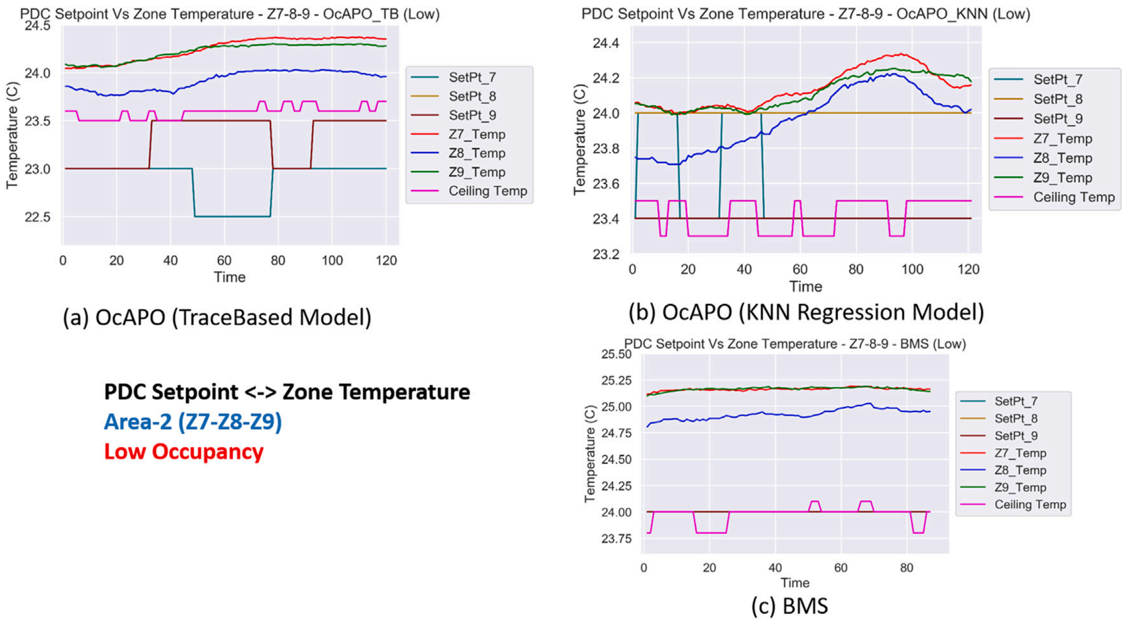

Fig. 16. PDC setPoint by BMS, OcAPO (TraceBased and KNN) Vs ceiling and zone temperature — high occupancy (Area-2).

occupancy. With additional data from the second phase, the KNN version is successful in finding slightly more-relaxed values of PDC setpoint, which in turn permits less aggressive opening of the PDC valves (thereby implicitly conserving energy).

- We observe that *OcAPO* determines the PDC setpoint 0.2 °C to 1 °C *above* the static BMS setpoint (24 °C) when the zone is not occupied and 0.6 °C to 1 °C *below* the BMS static setpoint when the zone is occupied. *OcAPO* effectively execute dynamic setpoint control to assure occupant thermal comfort.

6.4. PDC valve opening results

The changes to the PDC setpoint and room temperature trigger the PDC valves to operate (open/close). As illustrated earlier in Fig. 2, the PDC valves open between 10%–100% depending on the differences in the ceiling temperature and PDC setpoint value. We

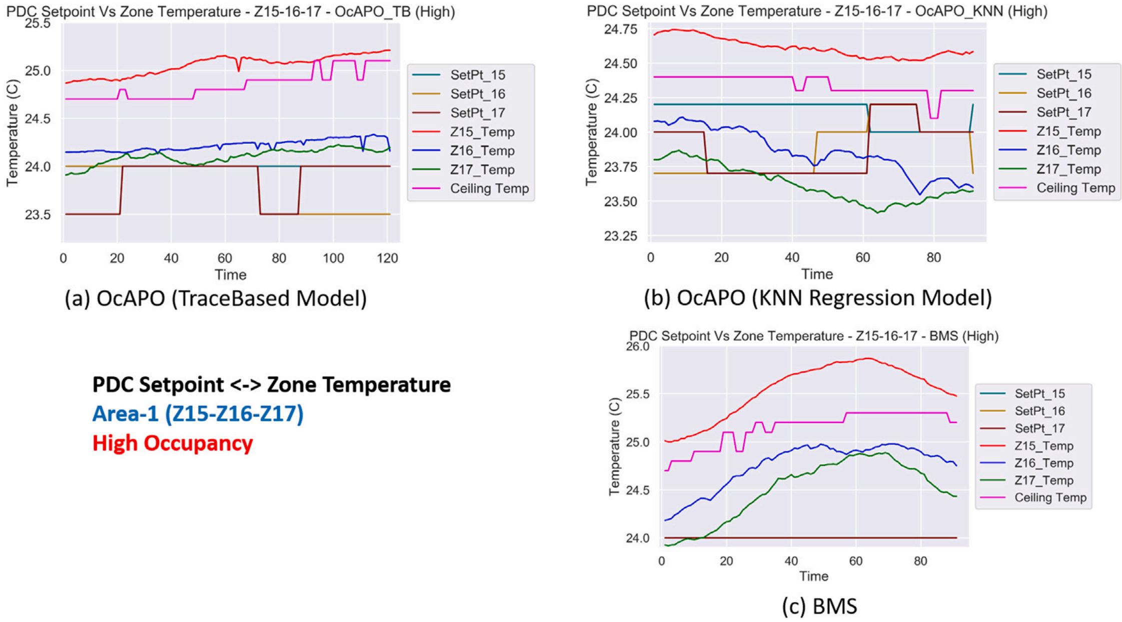

Fig. 17. PDC setPoint by BMS, OcAPO (TraceBased and KNN) Vs ceiling and zone temperature — high occupancy (Area-1).

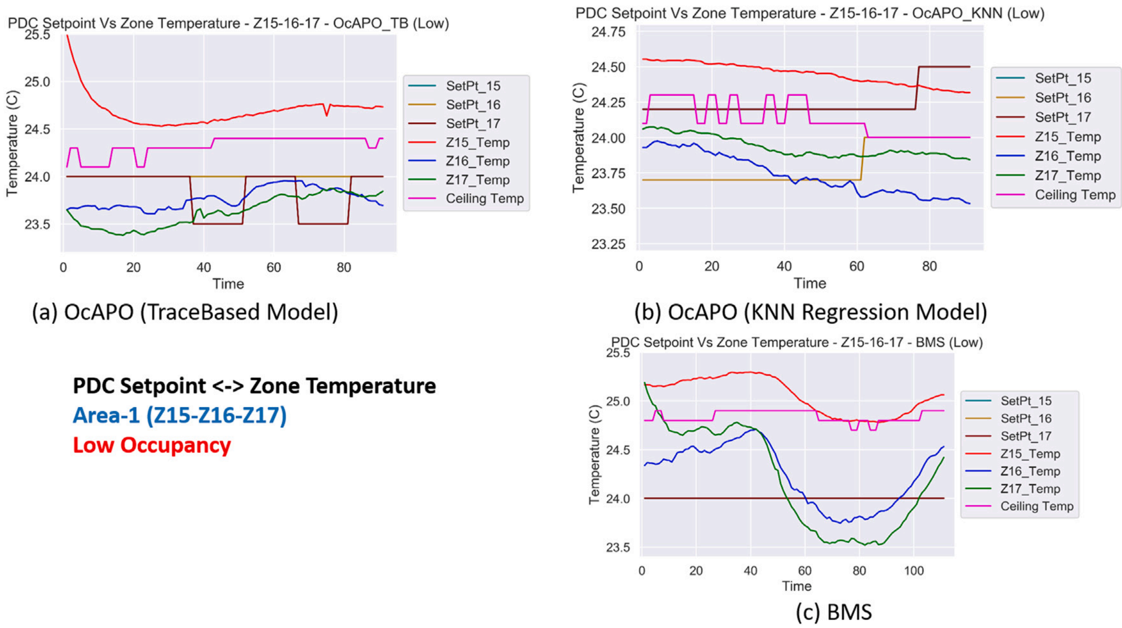

Fig. 18. PDC setPoint by BMS, OcAPO (TraceBased and KNN) Vs ceiling and zone temperature — high occupancy (Area-1).

study the opening percentage (%) of the valves for each methodology (BMS and the two *OcAPO* variants). Figs. 19–21 reports the PDC valve opening (%) for Area-3, 2 and 1 respectively for low-occupancy. As observed, the BMS module opens the valves up to 10% in Area-3 and 20% in Area-2; however, we have seen that the ambient temperature readings often significantly exceed the desired value. *OcAPO's* Trace-based version performs more aggressive PDC setpoint control, resulting in average valve opening percentage being above 90% under low and high occupancy and 0% under no occupancy across all areas. The KNN-regressor version of *OcAPO* performs more gradual control of PDC setpoints, resulting in average valve opening percentages of (a) 66.6% under high occupancy and 36.2% at no to low occupancy in Area-2; (b) 50% under high and low occupancy and less than 20% under no occupancy in Area-3; and (c) 100% under high occupancy and 36.9% at no to low occupancy in Area-1. Qualitatively speaking, a less aggressive

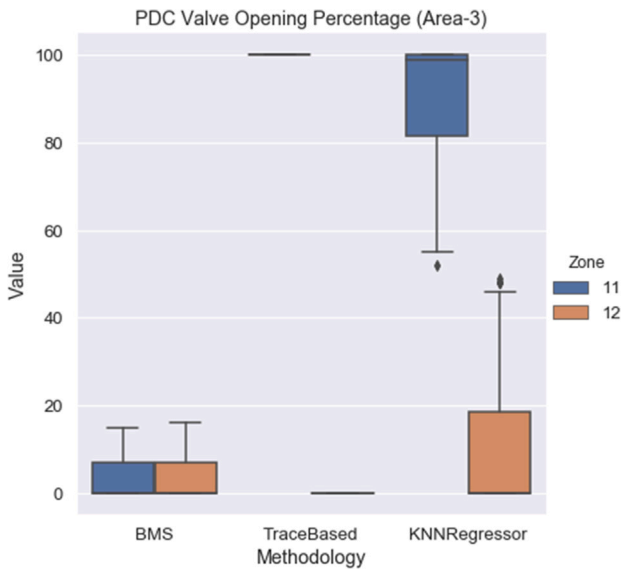

Fig. 19. PDC valve opening percentage for BMS Vs OcAPO (TraceBased and KNN regressor) across all 3 zones (Area-3).

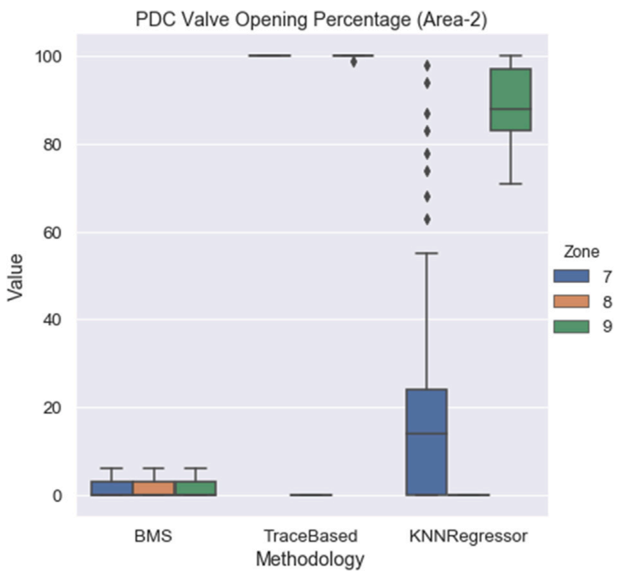

Fig. 20. PDC valve opening percentage for BMS Vs OcAPO (TraceBased and KNN regressor) across all 3 zones (Area-2).

valve opening percentage translates into lower heat exchange for the chilled water pipes, thereby translating into reduced cooling energy overheads.

6.5. Occupancy estimation

State-of-the-art object detectors experience errors (a) in partially-occluded environments, (b) for more distant objects that are smaller in size; such errors can cause *OcAPO* to perform incorrect PDC setpoint control. While the camera-based occupancy estimator had an accuracy of 93.7% (in classifying occupancy across the 4 bins), the use of temporal smoothing (over multiple readings within a 15-minute window) helped reduce the estimation error by an additional 23.8%. Overall, over a 6-hour operational period, occupancy estimation errors caused *OcAPO* to specify incorrect PDC setpoints thrice (i.e., for a total of 45 min).

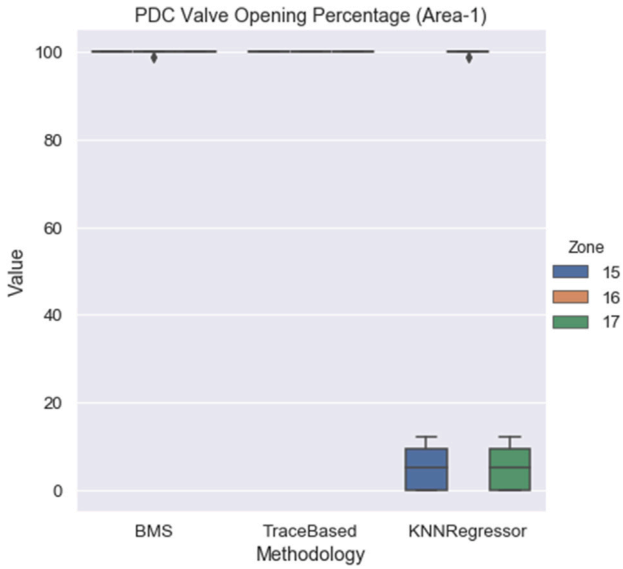

Fig. 21. PDC valve opening percentage for BMS Vs OcAPO (TraceBased and KNN regressor) across all 3 zones (Area-1).

7. Challenges encountered in practice and discussion

We now discuss the various practical challenges we faced in the model-building process and highlight the opportunities for building AI model-based PDC control in open-plan workspaces.

7.1. Challenges

Restricted Access to Control PDC: Our efforts were restricted to controlling the PDC setpoints and not the PDC valves. However, the energy consumption is driven by the PDC valves. It was challenging to understand the relationship between the status of the PDC valve and the setpoints as it is driven by the ceiling room temperature (which is affected by many factors).

Understanding occupancy patterns for predictive occupancy: As discussed in Section 4.2, occupancy plays a major role in providing occupant comfort across zones. Moreover, the valves must be opened sufficiently prior to the arrival of occupants to provide immediate thermal comfort. Thus, modeling the predictive occupancy is the key to successfully adjusting setpoints across zones for a given area. Our testbed is an experiential, collaborative space, where the set of students visiting the collaborative zone varies significantly, and there is no dedicated place for each student. The occupancy is thus bound to change at the zone level. Moreover, due to changing COVID restrictions, *the occupancy pattern was severely affected and non-stationary throughout 2020, 2021 and significant portions of 2022.* The Collaboration Space selected for the study was, initially (i.e., pre-COVID), a 24*7 space (including sleeping bunks) accessible to all the students. While we initially envisaged using the place for 24*7 observations and experimentation, COVID caused the operational hours of the space to be reduced to only 8.00 am–11.00 pm, Monday–Saturday. We thus did not have sufficient operational data under "normal" conditions to build a model that can accurately predict the changing occupancy over time. Hence, building a model that provides predictive estimates of zone-wise occupancy proved to be very challenging and remains a work in progress. This prediction problem will be significantly simpler in spaces such as classrooms that involve regularly-scheduled activities. Nonetheless, we believe that reasonably accurate *aggregate level occupancy prediction* may be possible for even open-floor, collaboration spaces (which are increasingly favored as part of the move towards 'open offices').

Gathering data to build model: As observed in Section 4.1, the outside weather condition also significantly impacts the indoor temperature. Thus, for the model to work under varying weather and occupancy conditions, it is essential to understand the effect of occupancy and its impact on the indoor temperature in tandem with the different weather conditions. Getting such data, driven by natural occupancy, under different weather and setpoint conditions is a considerable challenge. As an example of such practical challenges, we observed that the collaborative zone was highly occupied only during the late afternoon (2–6). This effectively prevented us from capturing and studying the impact of varying occupancy levels (e.g., High) during the Morning and Evening periods.

Scalability and Sensitivity of *OcAPO* : It is useful to investigate the overhead of expanding *OcAPO* to activate occupancy-aware PDC control over a larger area, such as our entire campus. The SMU campus comprises 10 buildings and approximately 20,000 square meters of shared-occupancy spaces. This necessitates the concurrent processing of approximately 120 camera feeds, assuming a camera deployment density comparable to our collaborative workspace. This translates to an average total network bandwidth requirement of 500 Mbps. Drawing from our implementation experience, we anticipate that achieving campus-wide execution of *OcAPO* can be accomplished using five servers, each equipped with 246 GB of memory and four GPUs. We expect that areas equipped with surveillance cameras (a prerequisite for the effective deployment of *OcAPO*) will already have servers in place to collect the video feeds; deploying *OcAPO* will thus not require any significant additional computing infrastructure. *OcAPO* can be implemented in any setting with a PLC backend governing PDC/HVAC valves, adjusting them based on room temperature and PDC setpoint. Expanding the *OcAPO* system will also entail gathering data to assess the impact of occupancy and external temperature across different areas and rooms. The data collection strategy may vary depending on the deployment area, as each location will exhibit unique occupancy patterns and sensitivity to occupancy, influenced by factors such as the presence of windows/doors or the height of ceilings.

Energy Cost Our testbed did not have separate energy meters to monitor zone-level energy consumption; instead, the building was instrumented to capture only the *building-level* chiller energy consumption. In fact, by its very design, PDC does not employ any separate fans or blowers, making it difficult to disaggregate the energy demand for individual zones. However, as highlighted in Section 6.3, we observed that the PDC valve opening percentage (an *indirect* indicator of energy consumption) decreased by 60% for Area-1 under low-occupancy conditions with the KNN regressor model. Additionally, in Area-2 and Area-3, PDC valves opened 10% less in unoccupied zones compared to the rule-based PDC control strategy.

7.2. Opportunities for PDC setpoint control

Accommodating Transient Artifacts: *OcAPO* relies on the data collected over a designated period to adjust the PDC setpoints. However, the relationship between PDC setpoint and table-level ambient temperature can experience short-term deviations, due to usage artifacts (e.g., users operating a heat-generating GPU server), which *OcAPO* currently is incapable of tackling. Adapting to such localized "heat island" effects requires additional and extensive sensor instrumentation and is also constrained by the longer response transient of PDC.

Modeling Historical Time-Varying Changes: *OcAPO* currently performs adaptation based on a model of ambient temperature evolution over a 20 min window, and is thus oblivious to the longer term correlation between ambient temperature readings and weather conditions. In reality, it is likely that adaptation polices that incorporate longer-term memory (e.g., 12 h history) might offer more optimal energy-vs.-comfort tradeoffs. However, developing such "trajectory-based models" is hard and requires significantly larger volumes of longitudinal training data.

8. Conclusion

Motivated by empirical observations of a material, occupancy-dependent discrepancy between the temperature recorded by ceiling-mounted sensors (which drives the PDC control loop) and occupant-level ambient temperature in an open plan collaborative ~1000 m^2 workspace, we have introduced two versions of *OcAPO*, an occupancy-aware, adaptive PDC setpoint control system. *OcAPO* takes real-time and diurnal occupancy estimates (obtained automatically by analyzing camera feeds), together with current weather conditions, as input to find an optimal PDC setpoint for all zones combined in a given area. In the data-driven trace based version, *OcAPO* utilizes a trace-based lookup table to dynamically specify an optimum tuple of PDC setpoint values for a given area. The trace based model only utilized coarse-grained observational data for PDC control—this resulted in setpoint settings that are less granular than ideal. Through a second phase of data collection, we collected finer-grained data (at increments of 0.2 °C) over a smaller region of the decision space, and then built a KNN-regressor based model for fine-grained PDC control. A real-world deployment of *OcAPO* established that it can maintain an optimum thermal comfort (range: $23.5-24.0$ °C) with a tolerance of 0.2 °C, in contrast to a conventional BMS system (where the setpoint is always maintained at 24.0 °C) which exceeds the optimal thermal comfort by ≥ 1.1 °C. Overall, *OcAPO* is able to maintain an average ambient temperature of 0.1 °C under low, and 0.2 °C under high occupancy conditions.

CRediT authorship contribution statement

Anuradha Ravi: Writing – original draft, Visualization, Supervision, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Dulaj Sanjaya Weerakoon:** Visualization, Data curation. **Archan Misra:** Writing – review & editing, Validation, Supervision, Resources, Project administration, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The data that has been used is confidential.

Acknowledgment

This work was supported partially by the National Research Foundation, Singapore under its NRF Investigatorship grant (NRF-NRFI05-2019-0007),by the BCA GBIC grant, Singapore and US ONR Grant #N00014-23-1-2119. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

References

[1] M. González-Torres, L. Pérez-Lombard, J.F. Coronel, I.R. Maestre, D. Yan, A review on buildings energy information: Trends, end-uses, fuels and drivers, Energy Rep. 8 (2022) 626–637.
[2] F. Oldewurtel, D. Sturzenegger, M. Morari, Importance of occupancy information for building climate control, Appl. Energy 101 (2013) 521–532.
[3] S. Azizi, R. Rabiee, G. Nair, T. Olofsson, Effects of positioning of multi-sensor devices on occupancy and indoor environmental monitoring in single-occupant offices, Energies 14 (19) (2021).
[4] M. Gruber, A. Truschel, J.-O. Dalenback, CO2 sensors for occupancy estimations: Potential in building automation applications, Energy Build. 84 (2014) 548–556.
[5] N. Li, Z. Yang, B. Becerik-Gerber, M. Orosz, Towards Energy Savings From a Bimodal Occupancy Driven Hvac Controller in Practice, China, 9-12 October, 2013, pp. 680–689.
[6] S.R. West, J.K. Ward, J. Wall, Trial results from a model predictive control and optimisation system for commercial building HVAC, Energy Build. 72 (2014) 271–279.
[7] T.C. Seng, T. Soe, Getting nature to help in energy efficiency of the air conditioning system, Energy Procedia 143 (2017) 230–236.
[8] C. Nesler, Adaptive control of thermal processes in buildings, IEEE Control Syst. Mag. 6 (4) (1986) 9–13.
[9] B. Dong, K.P. Lam, C.P. Neuman, Integrated building control based on occupant behavior pattern detection and local weather forecasting, in: Proceedings of Building Simulation 2011: 12th Conference of International Building Performance Simulation Association, Vol. 3, 2011, pp. 193–200.
[10] T. Kitzberger, J. Kotik, T. Pröll, Energy savings potential of occupancy-based HVAC control in laboratory buildings, Energy Build. 263 (2022) 112031.
[11] H. Rajabi, Z. Hu, X. Ding, S. Pan, W. Du, A. Cerpa, MODES: Multi-sensor occupancy data-driven estimation system for smart buildings, in: Proceedings of the Thirteenth ACM International Conference on Future Energy Systems, in: e-Energy '22, Association for Computing Machinery, New York, NY, USA, 2022, pp. 228–239, http://dx.doi.org/10.1145/3538637.3538852.
[12] N. Cao, J. Ting, S. Sen, A. Raychowdhury, Smart sensing for HVAC control: Collaborative intelligence in optical and IR cameras, IEEE Trans. Ind. Electron. 65 (12) (2018) 9785–9794, http://dx.doi.org/10.1109/TIE.2018.2818665.
[13] B. Balaji, J. Xu, A. Nwokafor, R. Gupta, Y. Agarwal, Sentinel: Occupancy based HVAC actuation using existing wifi infrastructure within commercial buildings, in: SenSys 2013 - Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems, 2013.
[14] S. Nagarathinam, H. Doddi, A. Vasan, V. Sarangan, P. Venkata Ramakrishna, A. Sivasubramaniam, Energy efficient thermal comfort in open-plan office buildings, Energy Build. 139 (2017) 476–486.
[15] H. Burak Gunay, W. O'Brien, I. Beausoleil-Morrison, Development of an occupancy learning algorithm for terminal heating and cooling units, Build. Environ. 93 (2015) 71–85, http://dx.doi.org/10.1016/j.buildenv.2015.06.009, URL https://www.sciencedirect.com/science/article/pii/S0360132315300238.
[16] J.R. Dobbs, B.M. Hencey, Predictive HVAC Control Using a Markov Occupancy Model, in: Proceedings of the American Control Conference, No. June, 2014, pp. 1057–1062.
[17] J. Shi, N. Yu, W. Yao, Energy efficient building HVAC control algorithm with real-time occupancy prediction, Energy Procedia 111 (September 2016) (2017) 267–276.
[18] J. Hou, H. Li, N. Nord, G. Huang, Model predictive control under weather forecast uncertainty for HVAC systems in university buildings, Energy Build. 257 (2022) 111793.
[19] Data predictive control using regression trees and ensemble learning, in: 2017 IEEE 56th Annual Conference on Decision and Control, Vol. 2018-January, No. September, CDC 2017, 2017, pp. 4446–4451.
[20] F. Smarra, A. Jain, T. de Rubeis, D. Ambrosini, A. D'Innocenzo, R. Mangharam, Data-driven model predictive control using random forests for building energy optimization and climate control, Appl. Energy 226 (February) (2018) 1252–1272.
[21] V.L. Erickson, M.Á. Carreira-Perpiñán, A.E. Cerpa, OBSERVE: Occupancy-based system for efficient reduction of HVAC energy, in: Proceedings of the 10th ACM/IEEE International Conference on Information Processing in Sensor Networks, 2011, pp. 258–269.
[22] T. Wei, Y. Wang, Q. Zhu, Deep Reinforcement Learning for Building HVAC Control, in: Proceedings - Design Automation Conference, Vol. Part 12828, No. 2, 2017.
[23] G. Gao, J. Li, Y. Wen, DeepComfort: Energy-efficient thermal comfort control in buildings via reinforcement learning, IEEE Internet Things J. 7 (9) (2020) 8472–8484.
[24] M. Biemann, F. Scheller, X. Liu, L. Huang, Experimental evaluation of model-free reinforcement learning algorithms for continuous HVAC control, Appl. Energy 298 (March) (2021) 117164.
[25] MARCO - Multi-Agent Reinforcement learning based COntrol of building HVAC systems, in: e-Energy 2020 - Proceedings of the 11th ACM International Conference on Future Energy Systems, No. Ml, 2020, pp. 57–67.
[26] M. Esrafilian-najafabadi, F. Haghighat, Energy & Buildings Occupancy-based HVAC control using deep learning algorithms for estimating online preconditioning time in residential buildings, Energy Build. 252 (2021) 111377.
[27] D. Azuatalam, W.L. Lee, F. de Nijs, A. Liebman, Reinforcement learning for whole-building HVAC control and demand response, Energy AI 2 (2020) 100020.
[28] B. Chen, Z. Cai, M. Bergés, Gnu-RL: A precocial reinforcement learning solution for building HVAC control using a differentiable MPC policy, in: BuildSys 2019 - Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, 2019, pp. 316–325.
[29] P. Fanger, Thermal Comfort: Analysis and Applications in Environmental Engineering, Danish Technical Press, 1970, URL https://books.google.com.sg/books?id=S0FSAAAAMAAJ.
[30] S. Lee, I. Bilionis, P. Karava, A. Tzempelikos, A Bayesian approach for probabilistic classification and inference of occupant thermal preferences in office buildings, Build. Environ. 118 (2017) 323–343, http://dx.doi.org/10.1016/j.buildenv.2017.03.009, URL https://www.sciencedirect.com/science/article/pii/S0360132317300951.
[31] D. Li, C.C. Menassa, V.R. Kamat, Personalized human comfort in indoor building environments under diverse conditioning modes, Build. Environ. 126 (2017) 304–317, http://dx.doi.org/10.1016/j.buildenv.2017.10.004, URL https://www.sciencedirect.com/science/article/pii/S0360132317304535.
[32] J. Kim, S. Schiavon, G. Brager, Personal comfort models – A new paradigm in thermal comfort for occupant-centric environmental control, Build. Environ. 132 (2018) 114–124.
[33] M. Luo, Z. Wang, K. Ke, B. Cao, Y. Zhai, X. Zhou, Human metabolic rate and thermal comfort in buildings: The problem and challenge, Build. Environ. 131 (2018) 44–52, http://dx.doi.org/10.1016/j.buildenv.2018.01.005, URL https://www.sciencedirect.com/science/article/pii/S0360132318300052.

[34] N. Morresi, S. Casaccia, M. Sorcinelli, M. Arnesano, A. Uriarte, J.I. Torrens-Galdiz, G.M. Revel, Sensing physiological and environmental quantities to measure human thermal comfort through machine learning techniques, IEEE Sens. J. 21 (10) (2021) 12322–12337.

[35] W. Hu, Y. Luo, Z. Lu, Y. Wen, Heterogeneous transfer learning for thermal comfort modeling, in: Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, BuildSys '19, Association for Computing Machinery, New York, NY, USA, 2019, pp. 61–70, http://dx.doi.org/10.1145/3360322.3360843.

[36] N. Somu, A. Sriram, A. Kowli, K. Ramamritham, A hybrid deep transfer learning strategy for thermal comfort prediction in buildings, Build. Environ. 204 (2021) 108133, http://dx.doi.org/10.1016/j.buildenv.2021.108133, URL https://www.sciencedirect.com/science/article/pii/S0360132321005345.

[37] M. Aftab, C. Chen, C.-K. Chau, T. Rahwan, Automatic HVAC control with real-time occupancy recognition and simulation-guided model predictive control in low-cost embedded system, Energy Build. 154 (2017) 141–156, http://dx.doi.org/10.1016/j.enbuild.2017.07.077, URL https://www.sciencedirect.com/science/article/pii/S0378778817305091.

[38] V.L. Erickson, S. Achleitner, A.E. Cerpa, POEM: Power-efficient occupancy-based energy management system, in: 2013 ACM/IEEE International Conference on Information Processing in Sensor Networks, IPSN, 2013, pp. 203–216, http://dx.doi.org/10.1145/2461381.2461407.

[39] J. Redmon, A. Farhadi, YOLOv3: An incremental improvement, 2018, CoRR abs/1804.02767.