Sustainable Computing: Informatics and Systems 31 (2021) 100585

Contents lists available at ScienceDirect

Sustainable Computing: Informatics and Systems

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

The energy saving potential of retrofitting a smart heating system: A residence hall pilot study

Yannick De Bock [a,*], Andres Auquilla [a,c], Ellen Bracquené [a], Ann Nowé [b], Joost R. Duflou [a]

[a] Department of Mechanical Engineering, KU Leuven - Member Flanders Make, Celestijnenlaan 300, 3000 Leuven, Belgium
[b] Computational Modelling Lab, Vrije Universiteit Brussel, Pleinlaan 2, 1050 Brussel, Belgium
[c] Department of Computer Science, University of Cuenca, Cuenca, Ecuador

ARTICLE INFO

Keywords:
Smart thermostat
Occupancy prediction
Energy conservation
Energy management
Smart home
Building automation

ABSTRACT

Energy conservation is of increasing importance in contemporary society. A large fraction of energy end-use can be attributed to space conditioning. Therefore, intelligent control systems were devised and commercialised in the form of smart thermostats. Hereto, the availability of occupancy information is essential such that heating and/or cooling schedules can be tailored to user needs. This way energy savings can be obtained without jeopardising user satisfaction. However, preceding studies generally rely on simulations to estimate the potential reduction in energy consumption. This work aims at quantifying the potential based on a real life experiment. The development of a smart heating system is presented along with the results of an actual field test of retrofitting this system in 14 single-user student rooms of a university residence hall. An experiment was conducted in which the heating was automatically steered for 1 week (26 March 2018–01 April 2018). Total energy savings range between 26.9% and 59.5% and calculated thermal comfort was not significantly affected by the autonomous control. Furthermore, an environmental impact reduction of 3.2 to 12.9 EcoPoints is estimated for the controlled week, resulting in a reduction of 37.5 to 150.2 $kgCO_2eq$.

1. Introduction and related work

Encouraging environmentally conscious behaviour (ECB) of consumers has been the subject of a substantial number of research efforts such as [1–3]. Advances in electronics and information technology have paved the way to automate some of these resource minimisation strategies [1]. The most well-known examples are probably smart thermostats, which have received substantial attention as result of a growing environmental awareness and the large proportion of space heating in total building energy consumption (over 40% for both residential and commercial buildings in the US [4,5], and 68% for residential space heating in the EU [6]). These thermostats try to overcome the limitations of simple programmable thermostats, which do not necessarily save energy compared to traditional, manually controlled thermostats [7] as a result of improper use due to complex user interfaces, lengthy user manuals, etc. [8]. To realise this, the availability of occupancy information is essential such that heating schedules can be tailored to user needs, which comprises reaching the desired temperature in anticipation of user arrival and automatically switching to a setback temperature in case of vacancy. This way energy savings can be obtained without jeopardising user satisfaction.

Such intelligent heating systems have been extensively researched and a number of smart thermostats are commercially available, such as the Nest,[1] Heat Genius[2] and Anna[3] solutions, but their intelligent features are not always well-received [9] These systems typically rely on motion sensors to infer occupancy schedules or on geofencing. Kleiminger respectively categorises these as schedule-based and context-aware approaches [10], of which an overview is discussed next. Starting with the former, in [11] a sensor network and machine learning techniques are used to model and predict the number of occupants in order to optimise heating control. Erickson et al. employed a combined network of cameras and motion sensors and a blended Markov chain for prediction in [12]. The Smart Thermostat [13] relies on basic sensing technology and a hidden Markov model to exploit occupancy and sleep patterns. Neural networks in combination with a look-up table to exploit the periodicity of human conduct, registered by a network of sensors, are used for occupancy-based conditioning in Mozer's Neurothermostat [14]. Preheat [15], finally, relies on motion

* Corresponding author.
E-mail address: yannick.debock@kuleuven.be (Y. De Bock).
[1] https://nest.com/
[2] https://www.geniushub.co.uk/
[3] http://www.whoisanna.com/

https://doi.org/10.1016/j.suscom.2021.100585
Received 16 February 2020; Received in revised form 22 February 2021; Accepted 10 June 2021
Available online 26 June 2021

sensors to detect occupancy and uses k-nearest neighbours to identify similar days, which are then used for prediction. Gupta's GPS thermostat [16], which evidently uses GPS information to predict occupancy, is an example of a context-aware technique. Contextual information is also used in TherML [17], which also uses GPS data, and [18] which uses cellular and WiFi connection information to predict occupancy. Krumm and Bush [19], finally, make use of both contextual information and occupancy schedules as they represent the task of learning presence schedules as a linear matrix problem, solve it using least squares and combine it with the driving time heuristic of Gupta's GPS thermostat.

The question this paper aims to answer is whether such smart heating systems effectively save energy, and whether these savings are obtained at the expense of user comfort or not. Where other studies generally rely on simulations and report savings varying from 10% till 42% [20], this research quantifies the savings potential based on a real-life experiment in a residence hall. To this end, an experiment was conducted in 17 student rooms in a residence hall in Leuven, Belgium. In a first stage, i.e. the learning phase, user occupancy and preferences of these rooms are gathered and used to train prediction models. Then, in the control stage, heating is steered automatically based on predicted occupancy and a comparison of the energy consumption is made between the manual and automatic heating control.

Savings from actual experiments are also reported in [21] where AC energy consumption is reduced by 13% in a conference room, and [7] where occupancy-responsive thermostats were retrofitted in three dormitories resulting in 5%–8% of savings in heating energy consumption. However, the implemented thermostats in [7] were reactive, not predictive, meaning that rooms cannot be preheated in anticipation of user arrival. Also, the underlying algorithms nor the occupancy detection accuracy have been discussed. Furthermore, the effect on thermal comfort and the environmental impact were not investigated. These limitations are tackled in this work, which is structured as follows. Section 2 outlines the design of the experiment, followed by a description of the underlying algorithms in Sections 3 and 4 , which respectively discuss the learning and control phase. The performance of the smart heating system is evaluated in Section 5 and finally, conclusions are formulated in Section 7.

2. Methodology

2.1. Experiment design

For this study 17 single-user student rooms of a residence hall in Leuven, Belgium, were equipped with two multi-purpose sensors (Aeotec Multisensor 6[4] and Zipato multisensor quad[5]) which together record motion, temperature, light intensity, humidity and the position of the door, i.e. open or closed. All residents were between 18 and 23 years of age, 11 of them were female and 6 were male. Two types of rooms can be distinguished, those consisting of only one space (10) and those with a separate bedroom (7). All sensor data are accumulated by two home controllers (Vera Edge[6]), which are strategically placed to maximise signal strength. This sensor network is extended by a bi-directional, wireless Danfoss thermostatic radiator valve[7] (TRV). On the one hand, the valve transmits the user-defined set point to the controller such that it can be stored in the database. On the other hand, the controller can set a new desired temperature, which, for example, allows to switch to a setback temperature in case of vacancy. Furthermore, at any point in time, the user can override the system by adjusting the TRV to the desired set point. The sensors communicate with the home controllers using the Z-Wave protocol. Energy consumption, finally, is measured by an ultrasonic thermal energy metre (Kamstrup Multical 302[8]), which measures the flow rate and change in temperature between outflow and return of the heat transfer liquid (i.e. radiated heat), for each room individually at a resolution of 15-minutes. Figs. 1 and 2 present an overview of the sensing set-up. The existing reference heating regime is a continuous heating strategy with a maximum temperature of 21 °C.

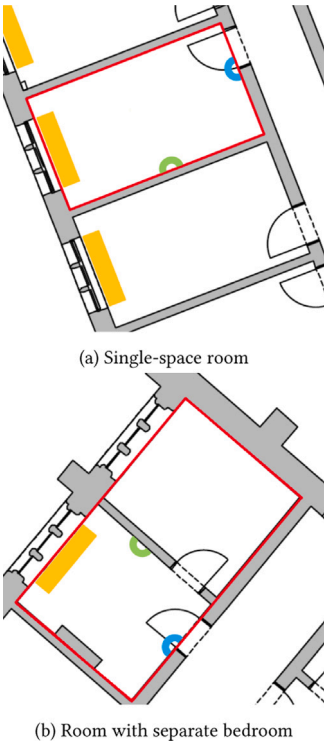

(a) Single-space room

(b) Room with separate bedroom

Fig. 1. Overview of the sensing environment, where the red box marks the room, the green and blue semicircles signify the multi-sensors, and the yellow box represents the radiator with TRV and energy metre.

Occupancy patterns depend on the academic calendar, as lecture schedules, assignments, extracurricular activities of the first and second semester can be significantly different. Moreover, occupancy during normal academic periods, holidays and exam periods is also expected to differ substantially. These phenomena, in combination with climatological aspects, limit the timing of the experiment. Training data were collected for four weeks, from 26 February 2018 until 25 March 2018 (offline phase), followed by one week, 26 March 2018 till 1 April 2018, for automated heating control (online phase). Table 1 presents an overview of the dataset.

2.2. Smart heating system

The operation of the implemented smart heating system can be split into two stages, an offline and an online phase. Gathering occupancy data and identifying patterns for each room is the focus of the first stage. This entails transforming the raw sensor data into properly formatted occupancy and set point information, and exploiting the historical data to model user behaviour. The occupancy detection system distinguishes three states: absent, present (active) and sleeping. Furthermore, the thermal characteristics of the room need to be learned. Hereto, the heating controller is simulated, such that the required time

[4] https://aeotec.com/z-wave-sensor
[5] https://www.zipato.com/product/multisensor-quad
[6] http://getvera.com/controllers/veraedge/
[7] https://www.smartheating.danfoss.com/en/a-solution-for-every-home/danfoss-link/connect-thermostat/
[8] https://www.kamstrup.com/en-en/products-solutions/thermal-energy-meters/multical-302

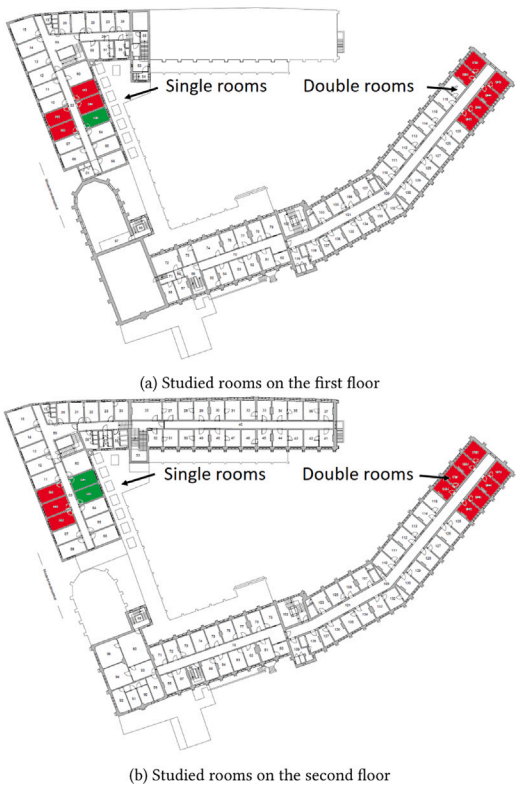

(a) Studied rooms on the first floor

(b) Studied rooms on the second floor

Fig. 2. Location of the monitored rooms (coloured) in the entire building. The 3 green rooms are not included in the control stage as either the resident did not want to participate or some sensor(s)/actuator malfunctioned.

Table 1
Dataset summary.

Equipment	*Sensors*	Aeotec Multisensor 6, Zipato Multisensor quad, Kamstrup Multical 302
	Actuator	Danfoss Link Connect TRV
	Controller	2LVera Edge
Data	*Period*	Learning: 26/02–25/03/2018 Control: 26/03–01/04/2018
	Rooms	17 (learning), 14 (control)
	Variables	Temperature, Humidity, Light, Motion, Door, Set points, Energy consumption
	Residents	18-23 years old 11 female, 6 male

to reach a certain temperature can be estimated. This enables heating the room just-in-time for expected user arrival. Moreover, this allows to significantly lower the temperature in case of vacancy. The online control stage computes which pattern is applicable for the current day and dynamically updates this prediction throughout the day. The set point of the expected state, i.e. present, absent or sleeping, is then sent to the corresponding TRV. An overview of the entire system is presented in Fig. 3.

The system is evaluated in terms of energy and environmental impact savings, and thermal comfort. Two methods were used to compute the realised energy savings. First, the degree day method is used, which normalises the energy consumption for differing weather conditions based on the number of heating degree days. The second method is based on ASHRAE Guideline 14 [22] and [7] in which a regression model of the building's energy consumption is built and used to project the consumption for the test week, allowing a comparison with the actual, measured consumption for that week. The environmental impact reduction is computed by a life cycle impact assessment (LCIA) study. Fanger's Predicted Mean Vote (PMV) and Predicted Percentage Dissatisfied (PPD) [23] are used to evaluate thermal comfort. These metrics estimate the thermal sensation of a large group of subjects for a particular combination of air temperature, mean radiant temperature, relative humidity, air speed, metabolic rate, and clothing insulation. However, as some of these variables are difficult to measure, assumptions were made as discussed in Section 5.3. In addition, the number of system overrides is also investigated.

3. Offline learning

Prior to controlling the heating system, its characteristics and user preferences must be inferred. This section discusses the three components that constitute this learning stage, i.e. occupancy detection, user modelling and heater simulation.

3.1. Occupancy detection

A wide variety of technologies is available for occupancy detection, ranging from motion sensors, over ultrasonic sensors to cameras. An overview of approaches is presented in [24]. As many of the presented technologies are either expensive, intrusive, complex to install or a hassle for the user (e.g. wearables based on for example RFID or GPS), wireless networks of simple sensors are by far the most popular approach in literature [20]. Therefore, two multi-purpose sensors were installed in each room, one on a wall facing a desk and the other one on the door frame (see Fig. 1). These sensors register motion, temperature, light intensity, humidity and the position of the door, allowing to distinguish between a "Present", "Absent" and "Sleeping" state by means of a simple post-processing logic. The raw sensor data are first formatted into time-series of 15-minute intervals and cleaned by noise removal techniques. Then, detection is performed by logical inference from the preprocessed data. To this end, the following series of user actions is defined:

Enter $D \wedge M \Longrightarrow \}\}Present\varepsilon$

Go to sleep $\neg P_{Sleeping} \wedge (P_{Present} \vee M) \wedge (H \geq 22 \vee H \leq 5) \wedge (L - L_{t-1} < -5) \Longrightarrow \}\}Sleeping\varepsilon$

Wake up $P_{Sleeping} \wedge (H > 5) \wedge (L - L_{t-1} > 5) \Longrightarrow \}\}Present\varepsilon$

Leave $\neg P_{Absent} \wedge D \wedge \neg M \Longrightarrow \}\}Absent\varepsilon$

Sleep $\neg P_{Absent} \wedge (H \geq 22 \vee H \leq 8) \Longrightarrow \}\}Sleeping\varepsilon$

Stay home $P_{Present} \wedge M \Longrightarrow \}\}Present\varepsilon$

Move $M \Longrightarrow \}\}Present\varepsilon$

Stay outside $P_{Absent} \vee (\neg M_{t-1} \wedge \neg M_{t-2}) \Longrightarrow \}\}Absent\varepsilon$

where M denotes that one of the sensors detected movement, D that the door is open, P_{state} the previous state, H the hour of the day [0-23], L the lighting level in the room [lux, $\in \mathbb{Z}_{>0}$] and t the current timeslot [0-95]. The actions are evaluated in this specific order. If none of the actions are valid, the state of the previous time slot is assumed. Some safety measures were built into the detection logic. For example, if the user enters the room (P_{Absent}) but somehow the sensor missed the event of opening the door, the action "Move" catches this in case one of the motion detectors is triggered.

To evaluate the detection algorithm, three residents were asked to note their presence on a timetable attached to the door for a period of two weeks. On average a detection accuracy of 92% was achieved, validating the implemented approach for further use. Mismatches can be attributed to imprecise registration of the user, sensor failures (missing data, false positives and false negatives), the aforementioned safety measures, and to the logic that might not cover all sorts of events.

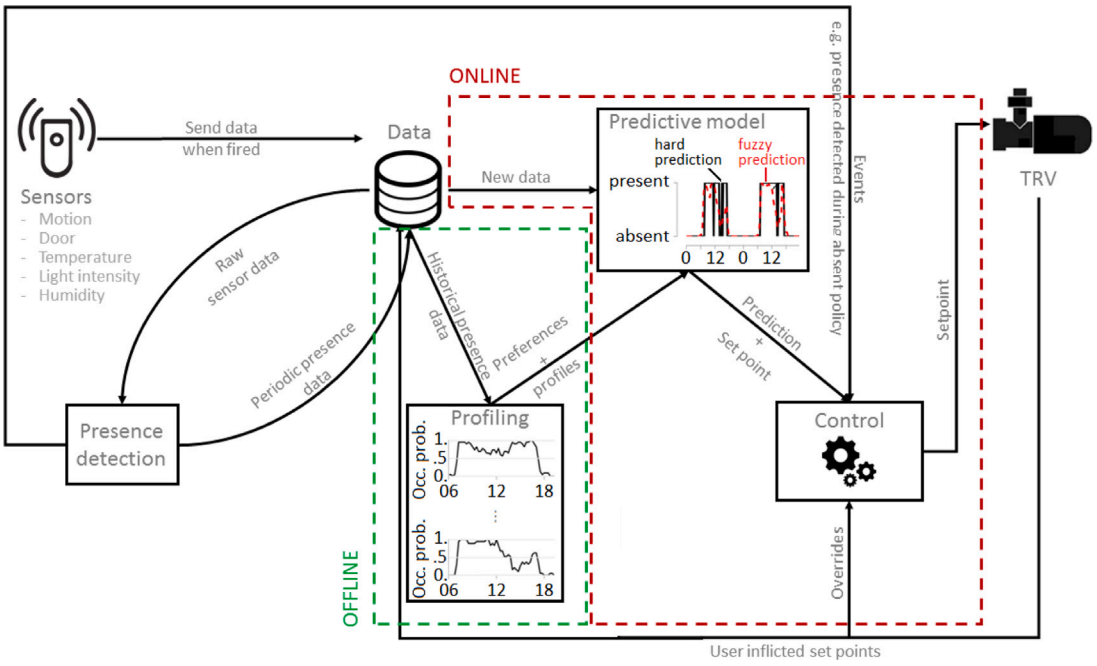

Fig. 3. Overview of the smart heating system. The offline and online stages are respectively indicated by green and red dashed lines. The remaining components are part of both stages.

3.2. User modelling

For each room, the gathered occupancy data are segmented into days and organised into 15-minute resolution arrays. Typical days are then identified by the user modelling technique presented in [25], which essentially groups similar days into clusters. This method uses a Dirichlet process mixture (DPM) model [26,27], implemented by Gibbs sampling [28], to infer the underlying patterns from the data. The grouping is performed by the stick-breaking process [29], which essentially generates weights representing the probability of a day belonging to a certain cluster. As opposed to many standard unsupervised learning techniques, the number of patterns does not have to be specified in advance. Instead an upper bound K, here set to 20, has to be defined and the different clusters C ($C \ll K$) will automatically emerge as a result of the sampling process. Another advantage over many of the discussed occupancy prediction algorithms is that this approach can deal with multiple activities, in this case "Present", "Absent" and "Sleeping". In addition, this strategy returns intuitive user profiles providing valuable insights into the users' behaviour and outliers are automatically detected and grouped in an outlier cluster.

A graphical representation of the model is depicted in Fig. 4. Here L represents the number of users/rooms (as the modelling is performed for each room independently, L is always equal to 1), N the number of days in the data set, and K a predefined upper bound on the number of clusters. The day of the week q_n is parametrised by multinomial distribution σ_k with prior Dirichlet distribution γ, class assignment $\mathbf{z}_n$ is generated from the multinomial distribution π_k with prior Beta distribution α (stick-breaking process), and the observation $\mathbf{x}_{n,l}$ is parametrised by the multinomial distribution of the classes μ_k with prior Dirichlet distribution β.

As described in [25], the stick-breaking weights π are initially uniformly assigned, i.e. $\pi_k = \frac{1}{K} \forall k \in K$. Then the observed days x are randomly allocated to the K classes by sampling the multinomial distribution $\mathcal{M}(\mathbf{z}_n|\pi)$, where $\pi = [\pi_1, \ldots, \pi_K]$. As a result, each cluster k will have approximately the same number of data points. After this initialisation, the weights are updated based on the number of days in each cluster. Then, each day is reallocated, while fixing the assignments of the other data points, based on the new weights and the likelihood

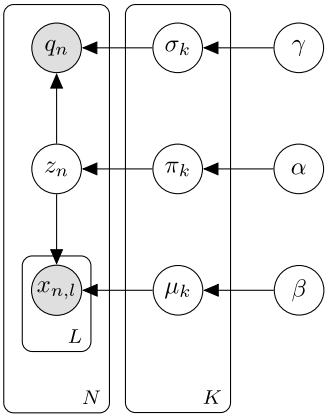

Fig. 4. Graphical model for identifying typical days in occupancy data [25,30].

of the data points' weekday and occupancy pattern given the cluster's weekday distribution and average occupancy profile respectively. The process is finally repeated until convergence. As such the K clusters will dynamically evolve into C typical days. An overview of the number of identified day types, for the different rooms, is presented in Table 2, and an example of those day types for one of the rooms is shown in Fig. 5. The first day type clearly represents holidays, weekends or days without lectures (when most students typically return home[9]). The second and third cluster depict regular weekdays with lectures in the morning, and for the second cluster also in the afternoon. These assumptions are also evident from the weekday distribution of the cluster assignments as illustrated in Table 3.

In addition to identifying occupancy patterns, the residents' temperature preferences have to be inferred. To this end, the most frequently user inflicted set point during training was selected as the desired temperature in case of presence. For the absent state a setback temperature

[9] In Belgium most students typically return home during weekends and holidays.

Table 2
Overview of the number of occupancy patterns for each room.

Room	# Clusters	Room	# Clusters
1	2	**10**	2
2	3	**11**	2
3	3	**12**	2
4	3	**13**	2
5	3	**14**	2
6	2	**15**	2
7	3	**16**	3
8	2	**17**	3
9	2		

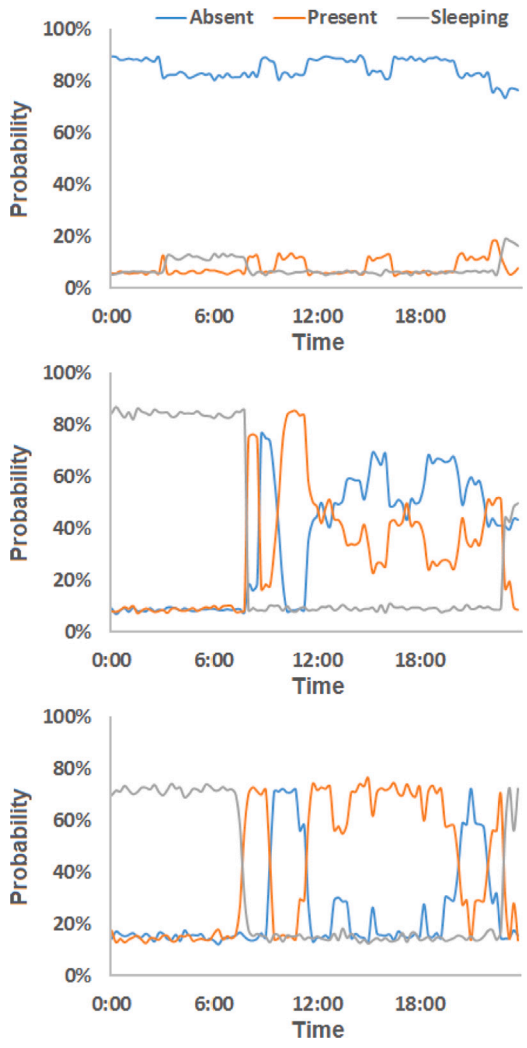

Fig. 5. Typical occupancy patterns as obtained for room 5. Some noise, preventing the probabilities to reach 0 or 100, occurs as a result of the sampling process.

of 16 °C was defined; and for the sleeping state a set point of 19 °C was specified.

3.3. Heater simulation

Besides modelling user behaviour, enabling anticipation of future heating requirements, the system also has to learn the thermal characteristics of the room and more specifically, the time required to reach a certain temperature. This enables preheating the room in such a way that the comfort temperature is reached at precisely the time of expected resident arrival. To this end a simplified model of the room was built to simulate the heater control and its effect on the room temperature. The model is dependent on heat gain from the radiator and heated adjacent rooms, heat losses through (exterior) walls and windows, and room characteristics such as dimensions and thermal resistance.

Table 3
Weekday distribution of cluster assignments for room 5.

Weekday	Cluster 1	Cluster 2	Cluster 3
Mon.	10.3%	11.8%	8.1%
Tue.	8.4%	31.3%	10.6%
Wed.	4.9%	7.0%	44.9%
Thu.	4.3%	31.2%	8.2%
Fri.	23.8%	6.6%	9.9%
Sat.	25.0%	7.0%	9.3%
Sun.	23.3%	5.1%	9.1%

3.3.1. Room characteristics

The single rooms measure 5 m by 2.7 m and are 3 m high, making up a volume of 40.5 m^3. Double-spaced rooms are, since they only have one radiator, treated as one large room with dimensions of 3.5 m by 6 m by 3 m, thus comprising a volume of 63 m^3. Half of the single rooms are oriented north, while the other half face south with the external window. Of the 7 double-spaced rooms, 4 are south-west oriented and 3 north-east. As the rooms were constructed in the 1950s or earlier, they are rather poorly insulated and have single pane windows. However, since the specific details are unknown, the thermal resistance of the room is estimated by simulation.

3.3.2. Dynamic temperature calculation

The rate of change of the room temperature can be computed from the exchange of thermal energy between the heat sources and the room on the one hand and between the room and the outdoor environment on the other hand, respectively known as heat gain and heat loss. The former is given by the heater output, which will be simulated, and the latter is defined in Eq. (1) where T_{room}, $T_{outside}$ and T_i respectively represent the room, outside, and adjacent room temperature, R represents the thermal resistance and N the number of adjacent rooms. Finally, the rate of temperature change is calculated by the difference between heat gain and heat loss, divided by the heat capacity of the air in the room, as specified in Eq. (2) where $m_{room\ air}$ and c_{air} respectively represent the mass of the air in the room and the specific heat capacity of air.

$$\frac{dQ_{loss}}{dt} = \frac{(T_{room} - T_{outside})}{R_1} + \sum_{i=0}^{N} \frac{(T_{room} - T_i)}{R_2} \tag{1}$$

$$\frac{dT_{room}}{dt} = \frac{1}{m_{room\ air} c_{air}} \left(\frac{dQ_{gain}}{dt} - \frac{dQ_{loss}}{dt}\right) \tag{2}$$

$$R = \frac{D}{kA} \tag{3}$$

Before simulating the heater output, the thermal resistance of the room is estimated. Hereto, Eq. (2) is used to approximate the measured indoor temperature and a grid search is performed on the R parameters. The heat gain is in this case directly given by the measurements of the heat metre on the radiator. Thermal resistance (Eq. (3)) is defined as the depth, i.e. thickness, of the surface D over the thermal conductivity k of the material, times the surface area A. For example, the thermal resistance of the inner wall R_2 of room 1 was found to be $3e^{-6} hK/J$. The surface area A equals $13.5m^2$ and the thickness D of the walls is around 10 to 15 cm. Therefore, the walls have a thermal conductivity k_2 of 0.686–1.029 J/hmK which corresponds to the conductivity of building bricks as specified in [31].

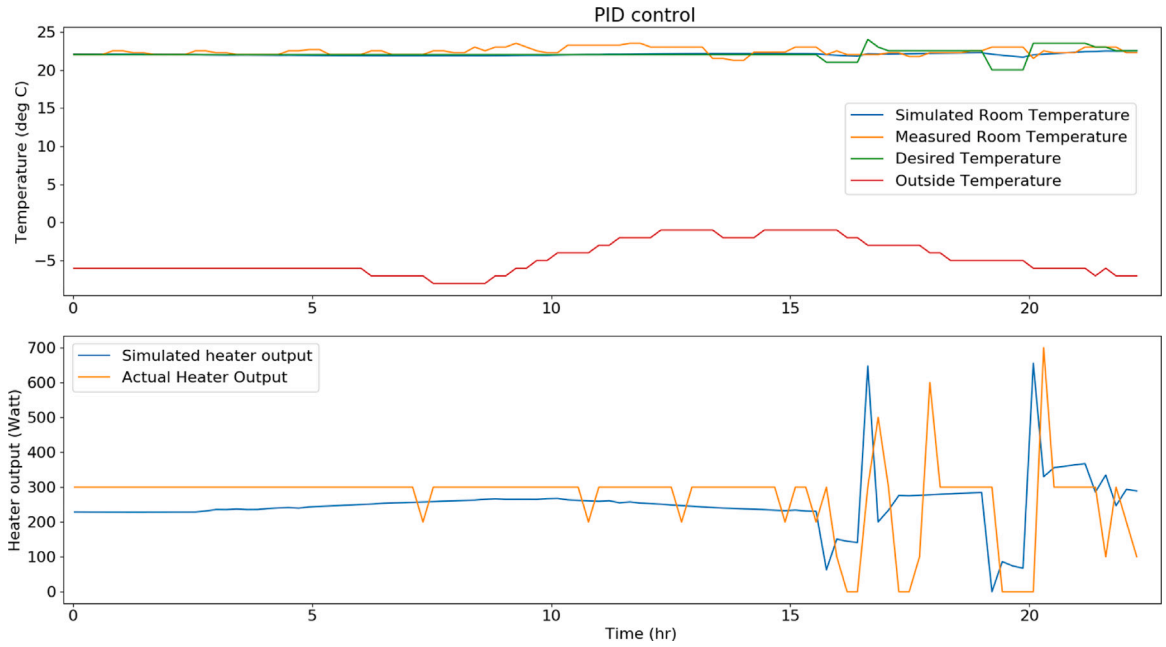

Fig. 6. Example of the heater simulation, for room 1 for 1 day.

3.3.3. PID control

Lastly, the behaviour of the heater is simulated to estimate the rate of heat gain from the radiator. The implemented radiator valves are proportional–integral–derivative (PID) controllers, which continuously compute the difference between the set point and current temperature, i.e. the error $e(t)$, and correct for this with proportional (P), integral (I) and derivative (D) gains, hence their name. As these parameters, set by Danfoss, are unknown, they have to be estimated. The proportional, integral and derivative terms are respectively defined in Eqs. (4)–(6), where K_p, K_i and K_d are tuning parameters, e is the error value, t represents the current time and $\tau \in [0, t]$. The heater output Q_{gain} is then given by $P+I+D \in [0, Q_{max}]$, where Q_{max} represents the maximum power of the heater. Q_{max} was extracted from the historical energy metre measurements.

$$P = K_p e(t) \tag{4}$$

$$I = K_i \int_0^t e(\tau) d\tau \tag{5}$$

$$D = K_d \frac{de(t)}{dt} \tag{6}$$

The values for the tuning parameters K_p, K_i and K_d are then found by performing a grid search while minimising the difference between the measured and simulated heater output and room temperature. Fig. 6 presents an example of the simulation, using these optimised PID parameters for room 1 for 1 day. The variables of the model are given in Table 4. Although the model is a simplified representation of the real environment, it is sufficiently rich to capture the main dynamic thermal behaviour of the room. This model enables the calculation of the required time to reach a certain temperature. As such, the heating can be controlled in anticipation of the expected occupancy and weather conditions, making up for the lack of predictive abilities of more advanced controllers such as model predictive control (MPC).

4. Online control

In this section the online control stage, comprising occupancy prediction and heating control is discussed. The number of rooms included in the experiment is in this stage reduced to 14 as either some sensors were malfunctioning or the residents preferred not to participate.

Table 4
Model parameters for room 1.

Variable	Description	Value	Unit
ρ_{air}	Density of air	1.205	kg/m^3
$m_{room\ air}$	Mass of air in the room (ρ_{air} * volume)	48.8	kg
c_{air}	Specific heat capacity of air	10005.4	J/kgK
R_1	Thermal resistance with outdoor environment	$12e^{-6}$	hK/J
R_2	Thermal resistance with adjacent room(s)	$3e^{-6}$	hK/J
Q_{max}	Maximum heater output	1500	W
K_p	Proportional gain	85	–
K_i	Integral gain	0.01	–
K_d	Derivative gain	80	–

4.1. Occupancy prediction

Once the residents' occupancy behaviour and heating control are modelled, the system can take over and intelligently steer the heating system by anticipating future events. The user's occupancy state is evaluated in real-time and serves as an input for the prediction algorithm. More specifically, the occupancy vector of the current day (until the current time slot) is matched with the patterns identified in Section 3.2, resulting in a likelihood of this day belonging to any of those profiles. As human conduct typically exhibits periodic activities on a weekly basis, e.g. academic lectures, weekday information is also taken into account by selecting the likelihood of each pattern occurring on the current weekday. Furthermore, the general likelihood of each pattern, i.e. the ratio of the number of times a pattern occurs to the total number of training days, also conveys valuable information about the future occupancy probability. These likelihoods are then multiplied to find the most likely pattern for the current day. Finally, the prediction is given by the occupancy information of the most likely pattern from the current time slot till the end of the day.

However, since the prediction relies on probabilities for each of the states, a threshold has to be applied. Therefore the pattern's probabilities are first binarised by a user tolerance (UT) parameter as defined in [32]. For each state the signal is transformed by discarding the lowest UT% of the state's probabilities and bringing the remaining probabilities to one, as illustrated in Fig. 7. As a result multiple states could have a value of one simultaneously. In this case, to limit user inconvenience, the state with the highest set point is selected (except

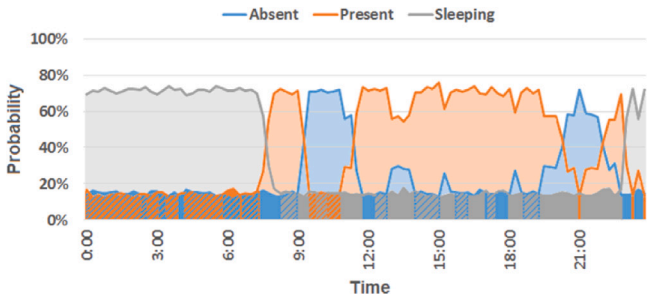

(a) Original occupancy pattern in which the dark shaded areas will be discarded, and the light shaded areas will become 100%.

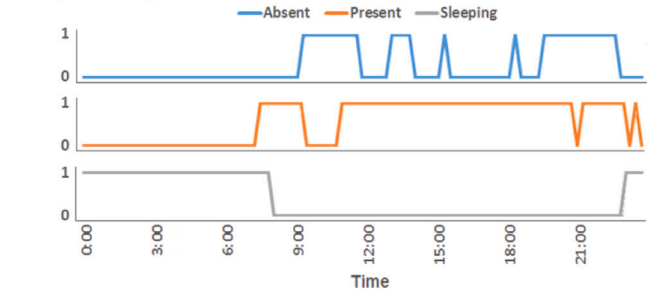

(b) Result of the binarisation process by applying the UT threshold.

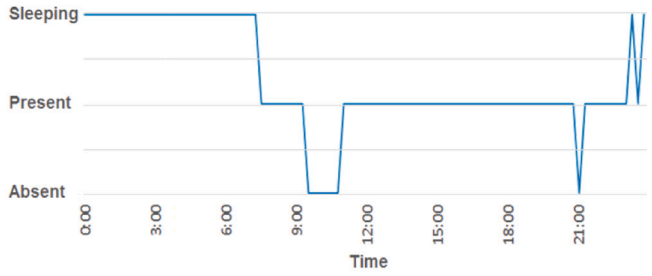

(c) Final discretised occupancy pattern.

Fig. 7. Example of the discretisation process of the third pattern of Fig. 5.

between midnight and 6 am, where the sleeping state is preferred over the present state). For this experiment a UT of 10% was used. In further research, the tolerance level of a user will be inferred from his feedback to the system, such as observed overrides.

4.2. Heating control

Control of the heating system is subject to both current and predicted occupancy, and desired temperature. Furthermore, during automatic control, the user can still perform overrides. To differentiate between system changes and overrides, each system initiated set point change is stored in an array. If during the next time slot the most recent set point in this array differs from the set point of the radiator, an override occurred.

Initially, the temperature set point corresponding to the current occupancy state is communicated to, and implemented by the radiator valve. Then, the system identifies the next change in the expected occupancy pattern, i.e. where the predicted state is different from the current state, and acts as follows:

- if the set point resulting from the expected state change is higher than the current set point, the system computes the time required to reach that set point, using the simulator of Section 3.3.3 and acts accordingly;
- if on the other hand, that set point is lower than the current desired temperature, the temperature is lowered half an hour before the anticipated event. This way, some extra energy can be saved and the user will not suffer any or very limited inconvenience.

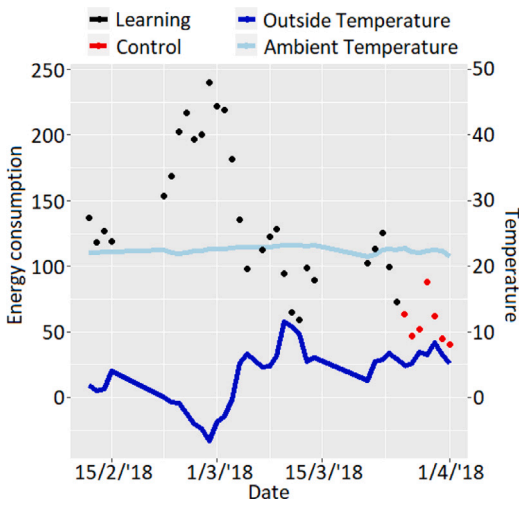

Fig. 8. Total energy consumption, outside and ambient temperature during the learning and control period.

Finally, the system checks whether the current set point corresponds with the set point of the current occupancy state and that no override occurred. If necessary, faulty predictions (e.g. presence is detected while absence was expected) are corrected. In the event of an override, the temperature of the override is maintained until the next change, as is the case in today's thermostats.

5. Performance analysis

It was envisaged to demonstrate energy savings per room by comparing the weather (and occupancy) normalised energy consumption of the learning and testing periods. However, as was concluded in [7], individual rooms cannot be considered independently. The building must be evaluated as a whole, as the temporal and spatial distribution of vacancy events affect the way thermal energy flows within a building and each room can react in a unique way. Attempts of assessing the smart heating system per room corroborate these findings.

Therefore, data of all monitored rooms were aggregated and treated as a complete building. Since the monitored rooms are adjacent (see Fig. 2), heat losses and gains to and from neighbouring rooms (under study) are incorporated. The data is then used to evaluate the smart heating system in terms of energy savings and thermal comfort. Moreover, the performance of the underlying algorithms is assessed.

5.1. Energy savings

From Fig. 8 it is clear that consumption during the control period was generally lower, although in line with the downward trend. Furthermore, the lower values might be attributed to the mentioned climatological and behavioural differences between the two periods. The analytical assessment of energy savings is therefore performed by two industry standard approaches, the degree day method [33] and an approach based on ASHRAE guideline 14 [7,22]. Both compare the energy consumption during the learning and control period while correcting for differences in weather conditions. In case of the ASHRAE based approach, changes in occupancy are also taken into account. However, it must be noted that some other exogenous factors, that these methods do not take into consideration, could still have changed between the two periods.

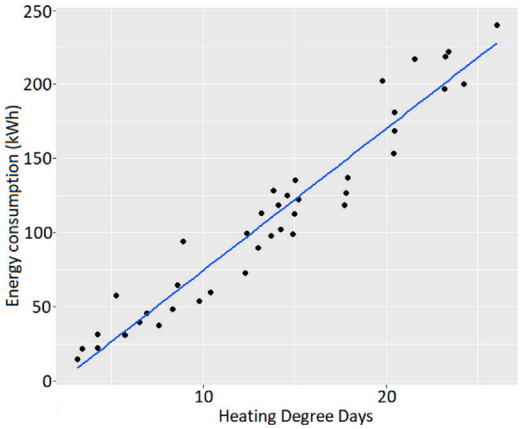

Fig. 9. Scatter plot with a regression line ($R^2 = 0.93$), illustrating the linear relationship between energy consumption and the number of HDDs during the learning period. As the accuracy of linear regression improves with more data, some extra days prior to the learning period and after the control period were included.

5.1.1. Degree day method

The degree day method is a well-known approach for weather normalisation of energy consumption, which factors out the differences in outside air temperature between two periods, allowing to fairly compare the normalised consumption figures. Heating degree days (HDDs) express how much (in degrees), and for how long (in days), the outside temperature was below a specified base temperature, which in turn is defined as the outside temperature above which heating is not required. This base temperature depends on the set point temperature, insulation level of the building, and sources of internal heat gain such as the heat dissipation of the residents and equipment. Since the experiment is performed in individual rooms, without cooking facilities the internal heat gain is assumed to be very limited and thus ignored. The set point temperature, however, does change over time and affects the energy consumption in a direct way. Therefore, the number of daily HDDs is computed (using BizEE's Degree Days calculator[10]) for the weighted average of the set point temperatures of the respective day, and its relation to energy consumption during the learning period is illustrated in Fig. 9. The regression model has an adjusted R-squared of 0.9318 and its formula is given by:

$$y = 9.5856x - 21.5947, \tag{7}$$

where x and y respectively represent the number of HDDs and energy consumption (in kWh).

Now the linear relationship has been clearly established, energy savings can be computed. However, since the set points during the control period are determined by the smart heating system, the HDDs for manual operation during that period cannot be computed. Hence, the regression model cannot be used to project the energy consumption of manual control into the future. Nevertheless, the HDDs for automatic control can be derived. Energy savings can then be estimated by dividing the measured consumption by the HDDs for both the learning and control period. Finally, the resulting numbers are multiplied with an average HDD value (e.g. the average HDD for 21 °C of 2017) to get normalised values, allowing to fairly compare the pre- and post-retrofit energy consumption (see Eq. (8), where $\tilde{Q}$ represents the normalised energy consumption). The normalised consumption of each day of the control period is compared with the average normalised consumption of the corresponding weekday during the learning period. Table 5 presents the obtained savings. Average savings amount to 39.0% (36.0 kWh). Substantial savings are achieved from Monday till

Table 5
Obtained absolute and relative savings by the degree day method.

	Savings	
	kWh	%
Mon.	30.1	36.8%
Tue.	40.0	49.2%
Wed.	32.0	37.8%
Thu.	13.4	14.1%
Fri.	27.1	27.0%
Sat.	59.4	55.8%
Sun.	46.8	51.0%
Overall	36.0	39.0%

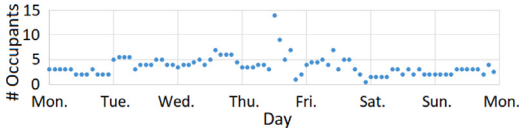

Fig. 10. Total occupancy rate of the residence hall during the control period.

Wednesday, on Friday and over the weekend. The latter should come as no surprise as many Belgian students typically return home for the weekend. On Thursday rather limited savings are accomplished which might be explained by an elevated occupancy rate, as shown in Fig. 10. On Friday, some students already return home. However, this behaviour is rather stochastic and/or not sufficiently captured by the profiling process as a result of limited training data. This, in combination with the rather conservative prediction method, results in slightly lower energy savings. The savings computed with the method based on ASHRAE guideline 14 should provide a better estimate as the occupancy rate is incorporated in the model.

$$\begin{aligned} \tilde{Q} &= \frac{Q}{HDDs} \cdot \overline{HDDs}_{21\ °C,\ 2017} \\ \text{Savings} &= \tilde{Q}_{Learning} - \tilde{Q}_{Control} \end{aligned} \tag{8}$$

For a more elaborate assessment of the system, it is compared with several baselines. The first baseline consists of a separate weekday and weekend schedule, which are applicable to all rooms. During weekdays the sleeping set point (19 °C) is applied from 23h until 06h30, the present set point from 06h30 until 8h30 and again from 16h00 until 23h, and finally the absent set point (16 °C) from 08h30 until 16h00. This schedule aims at simulating school days. For weekends, a distinct schedule is defined for residents who return home and for those who spend the weekend at the residence hall. The former is defined as follows: 16 °C the entire day on Saturday; and 16 °C from midnight until 17h00,[11] 22 °C from 17h00 until 23h00 and 19 °C from 23h00 until midnight on Sunday. In the latter, 19 °C is assigned from 23h00 until 09h00 and 22 °C from 09h00 until 23h00. The ratio of people returning home was estimated at 75%. The second baseline, is more oriented towards comfort. It implements 22 °C when the likelihood of being present is larger than 25%, and the set point corresponding to the most probable occupancy state otherwise (i.e. 16 °C in case of absence and 19 °C in case of sleeping). The third and final baseline is

[10] www.degreedays.net#generate

[11] Many students return to the residence Sunday in the evening.

Table 6
Obtained relative savings with the baselines.

	Baseline 1	Baseline 2	Baseline 3
Mon.	−3.7%	−7.4%	4.8%
Tue.	−10.2%	−10.2%	−2.6%
Wed.	−6.4%	−8.0%	0.8%
Thu.	4.5%	1.0%	10.8%
Fri.	12.3%	18.0%	23.0%
Sat.	36.2%	40.7%	42.2%
Sun.	17.0%	26.8%	32.7%
Overall	7.1%	8.7%	16.0%

Table 7
Discomfort and energy wastage of the different approaches.

	Discomfort	Energy waste
Baseline 1	39.2%	35.8%
Baseline 2	15.1%	27.6%
Baseline 3	45.0%	11.3%
Proposed approach	33.6%	9.5%

the most similar to the presented approach as its schedules are defined by the average occupancy patterns of each weekday. As such a separate schedule for each day of the week is generated. The heating degree days can be computed according to the schedules of the respective baselines, and can then be fed into Eq. (7) to estimate the corresponding energy consumption. The savings achieved using the baselines are reported in Table 6. These results are substantially lower, and sometimes even yield higher consumption than with manual control, than those obtained with the presented approach (Table 5). In addition, the efficiency of the presented approach and baselines is analysed in terms of their associated discomfort and energy waste in Table 7. Discomfort is computed by the fraction of time slots where the absent set point was implemented while the resident was present, by the total number of presence time slots. And energy waste is calculated as the fraction of time slots when the present set point was activated while the occupant was absent, by the total number of absence time slots. As expected, lowest discomfort is obtained with the second baseline, followed by the proposed approach. In terms of avoiding energy wastage, the proposed approach scores best.

5.1.2. ASHRAE guideline 14

ASHRAE guideline 14 consists of three approaches to compute energy savings. In this case, the whole building approach was applied. With this approach savings are estimated by developing a change-point regression model[12] on the learning data which describes consumption as a function of the outside temperature, the average outside temperature over the past 24h, the average occupancy state of the building, the day of the week and the time of the day. For this method, the data were formatted in arrays of 2h intervals. A variety of different models, with different combinations of variables and change-points, was built and 10-fold cross validation was used to select the best model, i.e. the model with the highest adjusted coefficient of determination (R^2). In addition, this model had the lowest coefficient of variation of the root mean squared error (CV-RMSE), which expresses the uncertainty of the model. The adjusted R^2 equals 0.88 and the CV-RMSE is 13.7% which is lower than the maximum 20% according to the ASHRAE guideline

[12] All computations were made in R using Muggeo's segmented package [34].

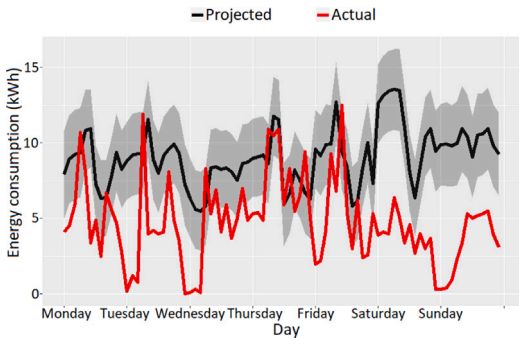

Fig. 11. Projected and actual energy consumption during the control period.

14. The formula of the segmented regression model is defined as:

$$\begin{aligned} Q &= \beta_0 + \beta_1 \times T_{Out} + \beta_2 \times (T_{Out} - C_2)^+ + \beta_3 \\ &\times T_{24h} + \beta_4 \times (T_{24h} - C_4)^+ + \beta_5 \times (T_{24h} \\ &- C_5)^+ + \beta_6 \times (T_{24h} - C_6)^+ + \beta_7 \times Occ \\ &+ \beta_8 \times (Occ - C_8)^+ + \beta_9 \times Mon + \beta_1 0 \\ &\times Tue + \beta_1 1 \times Wed + \beta_1 2 \times Thu + \beta_1 3 \\ &\times Fri + \beta_1 4 \times Sat + \beta_1 5 \times AM + \beta_1 6 \times \\ &PM + \beta_1 7 \times EV \\ &= 5.601 - 0.179 \times T_{Out} - 0.634 \times (T_{Out} - \\ &5.807)^+ - 6.605 \times T_{24h} + 15.042 \times (T_{24h} + \\ &1.648)^+ - 8.641 \times (T_{24h} + 0.915)^+ + 4.037 \\ &\times (T_{24h} - 11.468)^+ - 4.738 \times Occ + 6.669 \times \\ &(Occ - 2.133)^+ - 2.413 \times Mon - 2.375 \times \\ &Tue - 2.703 \times Wed - 2.789 \times Thu - 1.793 \\ &\times Fri - 0.402 \times Sat - 0.740 \times AM - \\ &0.669 \times PM - 1.931 \times EV \end{aligned} \quad (9)$$

where:

Q represents the energy consumption [kW],
T_{Out} the average outside temperature per 2 h interval [° C],
T_{24} the average outside temperature of the previous 24 h [° C],
Occ the average occupancy state of all rooms,
$(\ldots)^+$ terms that are evaluated when they are positive,
β_i model coefficients,
C_i change point beyond which β_i is applicable.

Finally, Mon, Tue, Wed, Thu, Fri, Sat, AM, PM, EV are dummy variables which indicate the interval's weekday (Monday till Sunday) and time of day (AM = 7 h–12 h, PM = 13 h–18 h, EV=19 h–0 h). The dummies for Sunday and night time do not have to be included in the model as they can be derived from the other variables.

The expected energy consumption during the testing period, with manual control, is estimated by inserting the observed environmental and behavioural conditions of the testing period into the regression model. The difference between this projected and measured consumption returns the energy savings, which amount to 366.5 kWh or 47.9%. A 90% prediction interval was used to provide lower and upper bounds on the savings, leading to savings between 26.9% (146.3 kWh) and 59.5% (586.7 kWh). Fig. 11 depicts the predicted and actual consumption and Table 8 presents a breakdown of the energy savings per day. From Monday till Thursday these savings, although slightly higher, match relatively well with the results obtained with the degree day method. However, from Friday until Sunday, the degree day method underestimates the savings to a greater extent since the impact of the occupancy rate is more significant as many students return home.

Table 8
Obtained absolute and relative savings according to ASHRAE guideline 14.

	Lower		Fit		Upper	
	kWh	%	kWh	%	kWh	%
Mon.	6.4	9.2%	38.7	37.7%	70.9	52.6%
Tue.	31.9	40.5%	63.5	57.5%	95.1	67.0%
Wed.	7.3	12.4%	38.1	42.5%	68.8	57.2%
Thu.	−19.3	−27.9%	12.1	12.0%	43.5	33.0%
Fri.	13.7	18.1%	45.0	42.1%	76.4	55.3%
Sat.	54.1	54.5%	85.6	65.5%	117.1	72.2%
Sun.	47.7	53.9%	80.6	66.5%	113.5	73.6%

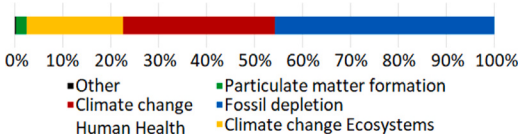

Fig. 12. Contribution of different impact categories to the overall avoided environmental impact using Europe ReCiPe H/A method [36].

5.2. Environmental impact

Additionally, the potential environmental impact reduction of the smart heating system is assessed using the lifecycle assessment (LCA) methodology. Background data from the Ecoinvent 3 LCI database[13] [35] are used and the overall environmental impact is characterised as a single score with the Europe ReCiPe H/A method [36]. The environmental benefit is calculated based on the measured energy savings and the environmental burden is estimated based on the required IT equipment, i.e. sensors, home controller, (single-board) computer, and data storage.

In total, an environmental impact reduction of approximately 8 EcoPoints is calculated for the control week, translating to 280 EcoPoints per year.[14,15] From Fig. 12 it is clear that the smart heating retrofit has the largest effect on the categories climate change (both Human Health and Ecosystems) and fossil depletion.

The climate change impact category provides the largest contribution to the single point score (>50%) and is given in terms of saved kilogrammes of equivalent carbon dioxide ($kgCO_2eq$), which is 0.256 $kgCO_2eq/kWh$ for the selected heating type (gas boiler). For the control week this means a reduction of 93.8 $kgCO_2eq$. and this can be extrapolated to an annual[14] impact reduction of 3257.0 $kgCO_2eq$. As was illustrated in previous work, for other heating types, such as fuel oil, the impact reduction would be even larger [37].

Furthermore, the environmental impact assessment is completed with an estimation of the environmental burden that is associated with the use of the IT equipment required for the smart heating system. The operation of the IT equipment is estimated to result in 0.065 EcoPoints/week and 0.58 $kgCO_2eq/week$. This environmental burden amounts to only 0.6% of the environmental benefit described above and is therefore considered to be negligible.

Finally, the potential cost savings are estimated. According to the Ecoinvent 3 LCI database, the gas and electricity consumption of the selected heating type are respectively 0.0972 m^3 and 0.0108 kWh per kWh of heat produced. Cubic metres of gas are converted to kilowatt-hours by a conversion coefficient of 10.224 kWh/m^3 for low calorific gas,[16] that is available in the region of Leuven, resulting in a gas consumption of 0.993 kWh per kWh of heat. Multiplying these numbers with the corresponding energy tariffs (0.035 €//kWh for gas and 0.3 €//kWh for electricity) and summing them returns a cost saving of 0.038 €//kWh or 13.81 €/ for the control week and 479.45 €//year[14]. The investment in IT equipment for the 14 student rooms amounts to approximately 2450 €/ but could be reduced to around 1600 €/ by discarding one of the multi-sensors, which means the payback period equals 3.3 to 5.1 years.[17]

5.3. Thermal comfort

However promising the realised savings, they should not be obtained at the expense of thermal comfort. Therefore, a pre- and post-retrofit assessment of thermal comfort is performed. Hereto, the widely popular PMV and PPD metrics [23] are used. PMV returns a score for the thermal sensation of a population ranging from −3 (cold) to +3 (hot), with ideal values between −0.5 and 0.5. PPD, on the other hand, indicates the percentage of people that will be dissatisfied under the given conditions. ASHRAE standard 55 requires that at least 80% of the occupants will be satisfied (or PPD ≤ 20%) [38]. Both PMV and PPD are computed using the CBE Thermal Comfort Tool [39] , which takes the following six comfort parameters as input:

1. Air temperature [° C]
2. Mean radiant temperature [° C]
3. Air speed [m/s]
4. Relative humidity [%]
5. Metabolic rate [met]
6. Clothing level [clo].

As the available information is limited and some of these parameters are hard to measure, assumptions have to be made. The air speed is set to 0.05 m/s which is characteristic of most indoor environments; the metabolic rate[18] is set to 1 met (seated relaxed) when residents are present or absent and to 0.8 met (reclining/sleeping) when residents are sleeping; and the clothing level is set to 1 clo (typical winter indoor clothing) for both the present and absent state and to 2.34 clo for the sleeping state.[19] The mean radiant temperature can be assumed to be equal to the air temperature [41], as is the case in many indoor climate studies. Tables 9 and 10 show the PMV and PPD results for the different occupancy states per room. From these tables it is clear that calculated thermal comfort was not negatively affected to realise the attained savings. On average thermal comfort even slightly improved. For the absent state, thermal comfort worsened for eight out of fourteen rooms. Which means that room temperature dropped to a marginally lower level during the control period, further illustrating the effectiveness of the autonomous system. A dependent t-test is performed to assess if the results during the learning and control period were statistically different. The respective p-values are 0.67, 0.06, 0.51, 0.38, 0.36 and 0.46, indicating that thermal comfort scores with and without automatic control are not significantly different (with a significance level of 0.05).

However, it must be noted that the PMV and PPD metrics represent the average score of a population. Therefore, it does not necessarily

[13] The LCI dataset that was used for the LCA study was: Heat, district or industrial, natural gas Europe without Switzerland| heat production, natural gas, at boiler modulating > 100 kW.

[14] In [37] it was estimated that there are 243 number of heating days (days where the average temperature between 7 and 18 h is lower than 20.5 °C) in Leuven.

[15] EcoPoints for the lower and upper bound of the energy savings respectively amount to 3.2 and 12.9 for the control week, resulting in 111 and 448 EcoPoints per year.

[16] https://www.vreg.be/nl/aardgas-omrekening-van-m3-naar-kwh

[17] Monetary savings for the lower and upper bound of the energy savings respectively amount to 5.51 €/ and 22.11 €/ for the control week, resulting in 191.29 and 767.61 €/ per year and a payback period of 12.81–3.19 years (or 8.36–2.08 years in case of the one multi-sensor scenario).

[18] https://www.engineeringtoolbox.com/met-metabolic-rate-d_733.html

[19] 2.34 clo corresponds to a person wearing sleepwear, sleeping on a mattress covered by a quilt as specified in [40].

Table 9
Comparison of average thermal comfort for the present and sleeping state during learning and control period. L and C respectively represent the learning and control phase.

Room	$PMV_{present}$		$PPD_{present}$		$PMV_{sleeping}$		$PPD_{sleeping}$	
	L	C	L	C	L	C	L	C
1	−0.25	0.00	6.9%	7.2%	0.28	0.38	7.1%	9.7%
2	−0.70	−0.25	16.3%	7.3%	−0.05	0.11	5.9%	5.8%
3	−0.83	−0.72	22.0%	17.0%	−0.40	−0.33	11.3%	7.9%
4	−0.49	−0.41	12.0%	8.8%	0.06	0.05	6.5%	5.4%
5	0.04	−0.31	9.3%	7.9%	0.12	0.37	6.3%	7.9%
6	−0.18	−0.35	7.8%	8.5%	0.22	−0.02	6.6%	5.5%
7	−0.02	0.06	6.1%	5.6%	0.30	0.35	7.3%	7.8%
8	−0.28	−0.35	7.4%	8.2%	0.22	0.16	6.4%	5.9%
9	−0.95	−0.68	24.5%	14.9%	−0.28	−0.13	6.7%	5.4%
10	−0.29	−0.20	7.0%	5.8%	0.30	0.35	6.9%	7.5%
11	−0.23	−0.06	7.1%	5.2%	0.17	0.34	6.1%	7.5%
12	0.10	−0.02	5.6%	5.1%	0.44	0.47	9.1%	9.7%
13	−0.30	−0.60	7.3%	12.7%	0.19	−0.03	6.0%	5.0%
14	0.85	0.56	21.6%	15.5%	1.10	0.89	31.2%	24.8%
Overall	−0.25	−0.24	11.5%	9.3%	0.19	0.21	8.8%	8.3%

Table 10
Comparison of average thermal comfort for the absent state during learning and control period. L and C respectively represent the learning and control phase.

Room	PMV_{absent}		PPD_{absent}	
	L	C	L	C
1	−0.33	−0.21	7.8%	8.1%
2	−0.72	−0.43	17.0%	10.0%
3	−0.79	−0.84	20.6%	21.1%
4	−0.42	−0.53	10.9%	11.2%
5	−0.32	−0.53	9.6%	11.1%
6	−0.34	−0.57	9.3%	13.7%
7	−0.04	0.06	5.8%	6.3%
8	−0.27	−0.44	6.8%	9.6%
9	−0.99	−0.73	26.0%	16.2%
10	−0.34	−0.26	7.7%	6.4%
11	−0.38	−0.16	9.1%	6.2%
12	0.09	−0.21	5.4%	6.7%
13	−0.32	−0.64	7.5%	13.8%
14	0.87	0.26	21.5%	10.0%
Overall	−0.31	−0.38	11.8%	10.8%

mean that the particular resident is more comfortable when PMV and PPD scores improve. For example, room 14 has substantially higher PMV and PPD values, which is due to the resident regularly specifying an unusually high ($\geq 25°$ C) desired temperature. Since the intelligent control system assumed a 22 °C comfort temperature (present state) for this resident, PMV and PPD values improved to some extent. However, as apparent from the number of system overrides given in Table 11, the resident often adjusted the heating to this unusually high set point. On average, however, only 5.79 overrides or less than one per day were performed per room during the control period. The number of user interactions, i.e. set point changes, during the learning period are reported as well. This allows to analyse the impact on user convenience. The fewer the user has to intervene, the better. In general, the number of interactions is somewhat lower during the control period. It must be noted however, that during manual operation, the resident might occasionally forget to lower the heating to a setback temperature or deliberately maintain the comfort temperature (e.g. to avoid a cold room upon arrival) during absence which might lead to an underestimation of the control actions during this period.

5.4. Predictive performance

Finally, the performance of the smart heating algorithm itself is evaluated. To this end, 5 standard classification metrics, accuracy, Matthews Correlation Coefficient (MCC), precision, recall and F1-score,

Table 11
Number of system overrides during the learning and control period for each room.

Room	Number of overrides		Average daily number of overrides	
	L	C	L	C
1	45	15	1.61	2.14
2	15	3	0.54	0.43
3	33	7	1.18	1.00
4	15	3	0.54	0.43
5	13	0	0.46	0.00
6	17	3	0.61	0.43
7	22	2	0.79	0.29
8	15	2	0.54	0.29
9	33	8	1.18	1.14
10	31	3	1.11	0.43
11	79	7	2.82	1.00
12	21	9	0.75	1.29
13	18	1	0.64	0.14
14	57	18	2.04	2.57
Overall	29.57	5.79	1.06	0.83

are computed. Accuracy is defined as the proportion of data that has been correctly classified, and is given by the formula:

$$Accuracy = \frac{T_P + T_N}{n} \tag{10}$$

where T_P and T_N respectively represent the number of true positives and true negatives, and n is the total number of samples. Since the frequency of the different activities, i.e. present, absent and sleeping, is typically uneven, the results might be affected by the accuracy paradox. This paradox entails that if one class dominates the dataset, constantly predicting this class will still yield a high accuracy score. Therefore, in order to account for this imbalance, MCC scores [42] were computed as well. MCC scores range from −1 (total disagreement) to +1 (perfect prediction), where 0 represents random prediction. For a more detailed evaluation, per activity class, precision, recall and F1-scores are reported. Precision represents the ratio of true positives to the sum of true positives and false positives, $\frac{T_P}{T_P+F_P}$. Recall, or sensitivity is defined as the true positive rate (TPR), i.e. $\frac{T_P}{T_P+F_N}$. The F_1 score combines precision and recall as in Eq. (11)

$$F_1 = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall} \tag{11}$$

The predictive capabilities are reviewed in Table 12. Generally speaking, accuracy and MCC scores are rather modest. For the absent state, precision is fairly high and recall is quite low. The opposite is true for the present state. Lastly, both precision and recall are reasonably high for the sleeping state. These results can be justified by the user tolerance parameter (UT). As a result of applying UT to the identified occupancy patterns, absence will only be predicted when the algorithm is really sure the room will not be occupied, leading to high precision and low recall scores. Since the present state is favoured by the prediction algorithm, recall is rather high while precision is fairly low. Sleeping is more predictable, as bedtime is relatively fixed, resulting in high scores for both precision and recall. For a few rooms some scores are not available, which means that the corresponding state has not occurred. Rooms 4 and 5 have no scores for the sleeping state, which might be the result of a malfunctioning light sensor. According to Table 12 both room 9 and 13 are never vacant. In this case, the occupancy detection system has failed to detect "Leave" events. If it were not for this flaw in the detection system, even more savings might have been achieved.

6. Discussion

The encouraging results suggest that the presented system can generate substantial energy savings, while maintaining thermal comfort and user convenience. The latter, measured by the number of user

Table 12
Assessment of the underlying algorithm's predictive abilities. A, P and S respectively represent the absent, present and sleeping state.

Room	Accuracy	MCC	Precision			Recall			F1-score		
			A	P	S	A	P	S	A	P	S
1	.75	.65	.82	.61	.94	.18	.93	.97	.30	.74	.95
2	.66	.52	.94	.19	.84	.56	.84	.93	.70	.31	.88
3	.69	.36	.81	.32	.74	.74	.48	.65	.78	.38	.69
4	.88	.44	.89	.96		.99	.28		.94	.44	
5	.65	.38	.87	.50		.55	.85		.67	.63	
6	.66	.35	.98	.09	.55	.64	.67	1.0	.77	.16	.71
7	.56	.41	.87	.19	.87	.44	.68	.94	.59	.30	.90
8	.72	.53	.92	.20	.78	.68	.66	.90	.78	.30	.84
9	.88	.80		1.0	.98		.84	.95		.91	.97
10	.74	.64	.92	.44	.91	.62	.81	1.0	.74	.57	.96
11	.56	.39	.94	.14	.73	.49	.67	.97	.64	.23	.83
12	.74	.34	.85	.35	.52	.80	.33	1.0	.83	.34	.68
13	.97	.95		1.0	.95		.97	.98		.98	.97
14	.66	.53	.70	.43	.94	.42	.72	.95	.53	.54	.95
Avg.	.72	.52	.87	.46	.81	.59	.70	.94	.69	.49	.86

interactions, was, however, expected to be more significantly improved. This might, in part, be due to residents not switching to a setback temperature, either intentionally to avoid discomfort upon arrival or involuntary, when leaving the room. As such the number of interactions during the learning period could be underestimated. The profiling algorithm was able to extract meaningful patterns from limited training data comprising multiple activities (i.e. absent, present, sleeping). The required number of days to construct reliable profiles naturally depends on the rate with which recurrent behaviour occurs. Typically humans exert daily, weekly or bi-weekly patterns. For example, universities in Belgium often employ separate course schedules for even and uneven weeks. As such, one month of data should suffice to capture most patterns. Since throughout the (academic) year the residents' occupancy schedules may shift, the need to retrain the profiles might arise. However, as discussed in [25] and [43], the algorithm can be extended with a mechanism that deals with evolving data. In [25], the stability and predictive power of the identified profiles is also demonstrated for larger data sets.

Although the presented system seems promising, some important limitations need to be addressed. First, occupancy detection is purely based on motion and door sensors. Therefore, it is susceptible to misclassifications when residents perform stationary activities such as reading or studying. Extending the system with for example smart metres to measure electricity consumption or network connectivity loggers can help overcome this issue. Second, data is handled at a 15-minute granularity, meaning that occupancy can only be derived once every 15 min. This granularity was initially chosen as it was empirically found that the rooms could be heated to a reasonable temperature within 15 min. Thus requiring predictions 15 min ahead of time. Furthermore, in case the resident is present at an unexpected moment in time, it was deemed reasonable to endure a certain level of discomfort for 15 min. Nevertheless, the effect of more fine-grained data on system performance demands further investigation. Finally, heat transfers from adjoining rooms that were not part of the study could not be incorporated in the energy savings computations. However, since this applies to both the learning and control period, its impact is believed to be very limited. Thus, the computation stands when the savings of the studied rooms are considered in isolation. However, when assessing the reported savings in terms of the entire building, they may be slightly overestimated as diminished consumption in the controlled rooms might elicit an increase in consumption in the adjoining rooms.

7. Conclusions and future work

This paper investigates the energy saving potential of adopting the proposed intelligent heating control in student rooms. This type of application is particularly interesting as it is fairly straightforward (only 1 room), and the rooms are vacant on a regular, though non-deterministic, basis (lectures, holidays, weekends, etc.). Moreover, in this case a continuous heating strategy is implemented.

As was also concluded in [7], the temporal and spatial distribution of vacancy events and the influence of heated adjacent rooms limits the assessment to the complete building. The extent to which energy consumption was reduced in individual rooms can thus, unfortunately, not be determined since a lowered consumption in one room could induce elevated consumption in the neighbouring room(s). The way thermal energy flows within the building, as a result of the vacancy events, also complicates the heater simulation. Further research is required to improve the simplified model.

An experiment was conducted in which the heating was automatically controlled for 1 week. The corresponding consumption was compared to the consumption of four weeks without the intelligent control, while normalising for occupancy and climatological differences. Total savings range between 26.9%, or 146.3 kWh, and 59.5%, or 586.7 kWh. Additionally, the system was compared with and significantly outperforms three baselines. Furthermore, the environmental impact reduction was estimated to range from 3.2 to 12.9 EcoPoints or 37.5 to 150.2 $kgCO_2eq$. In addition, the potential cost savings were estimated. For the control week, these amount to 13.81 €/, leading to annual savings of 479.45 € and a payback period of 3.3–5.1 years. Moreover, calculated thermal comfort was not affected by the energy conservation endeavour and few system overrides were registered.

However, more and longer of these real-life experiments, in varying building types and regions, must be performed to establish an elaborate view on the energy conservation capacity of smart thermostats in general. Furthermore, savings should be compared to realistic benchmarks. For example, many modern non-intelligent thermostats can be programmed and controlled remotely over the internet, which, if used properly, allows to optimise consumption. In contrast to regular programmable, offline, thermostats schedules and settings can be modified in a very easy, intuitive way. Therefore, the potential consumption reduction due to intelligent thermostats should also be compared with reference to this case.

CRediT authorship contribution statement

Yannick De Bock: Conceptualization, Methodology, Software, Formal analysis, Investigation, Resources, Data curation, Writing - original draft. **Andres Auquilla:** Conceptualization, Writing - review & editing. **Ellen Bracquené:** Formal analysis, Writing - review & editing. **Ann Nowé:** Supervision. **Joost R. Duflou:** Conceptualization, Writing - review & editing, Supervision.

Declaration of competing interest

No author associated with this paper has disclosed any potential or pertinent conflicts which may be perceived to have impending conflict with this work. For full disclosure statements refer to https://doi.org/10.1016/j.suscom.2021.100585.

Acknowledgement

The authors would like to recognise the financial support from Flanders Innovation and Entrepreneurship (VLAIO).

References

[1] L.H. Shu, J.R. Duflou, C. Herrmann, T. Sakao, Y. Shimomura, Y. De Bock, J. Srivastava, Desgin for reduced resource consumption during the use phase of products, CIRP Annals 66 (2) (2017) 635–658.
[2] S.-L. Lai, L.H. Shu, Do-it-yourselfers as lead users for environmentally conscious behaviour, Procedia CIRP 15 (2014) 431–436.

[3] J. Srivastava, L.H. Shu, Affordances and product design to support environmentally conscious behaviour, Mech. Des. 135 (10) (2013) 8.
[4] Energy Information Administration, Residential Energy Consumption Survey, 2009, https://www.eia.gov/consumption/residential/data/2009/index.php?view=consumption#end-use. (Accessed 1 2018).
[5] Energy information administration, Commercial Buildings Energy Consumption Survey, 2012. https://www.eia.gov/consumption/commercial/data/2012/c&e/cfm/e1.php. (Accessed 1 2018).
[6] European Environment Agency, Household energy consumption by end-use in the EU-27, 2012. www.eea.europa.eu/data-and-maps/figures/households-energy-consumption-by-end-uses-4. (Accessed 1 2017).
[7] M. Pritoni, J.M. Woolley, M.P. Modera, Do occupancy-responsive learning thermostats save energy? A field study in university residence halls, Energy Build. 127 (2016) 469–478.
[8] T. Peffer, M. Pritoni, A. Meier, C. Aragon, D. Perry, How people use thermostats in homes: a review, Build. Environ 43 (12) (2011) 2529–2541.
[9] R. Yang, M.W. Newman, Learning from a learning thermostat: lessons for intelligent systems for the home, Proc. UbiComp (2013) 93–102.
[10] W. Kleiminger, F. Mattern, S. Santini, Predicting household occupancy for smart heating control: a comparative performance analysis of state-of-the-art approaches, Energy Build 85 (2014) 493–505.
[11] S. Mamidi, Y.H. Chang, R. Maheswaran, Improving building energy efficiency with a network of sensing, learning and prediction agents, Proc. AAMAS 1 (2012) 45–52.
[12] V.L. Erickson, S. Achleitner, A.E. Cera, POEM: Power-efficient occupancy-based energy management system, in: Proc. ACM/IEEE IPSN, 2013, pp. 203-216.
[13] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, K. Whitehouse, The smart thermostat: using occupancy sensors to save energy in homes, Proc. ACM Sensys (2010) 211–224.
[14] M.C. Mozer, L. Vidmar, R.H. Dodier, The neurothermostat: predictive optimal control of residential heating systems, Proc. NIPS (1997) 953–959.
[15] J. Scott, A.B. Brush, J. Krumm, B. Meyers, M. Hazas, S. Hodges, N. Villar, PreHeat: controlling home heating using occupancy prediction, Proc. UbiComp (2011) 281–290.
[16] M. Gupta, S.S. Intille, K. Larson, Adding GPS-control to traditional thermostats: an exploration of potential energy savings and design challenges, Proc PerCom (2009) 95–114.
[17] C. Koehler, B.D. Ziebart, J. Mankoff, A.K. Dey, Therml: occupancy prediction for thermostat control, Proc. UbiComp (2013) 103–112.
[18] S. Lee, Y. Chon, Y. Kim, R. Ha, H. Cha, Occupancy prediction algorithms for thermostat control systems using mobile devices, IEEE Trans. Smart Grid. 4 (3) (2013) 1332–1340.
[19] J. Krumm, A.B. Bush, Learning time-based presence probabilities, Proc. PerCom (2011) 79–96.
[20] T.A. Nguyen, M. Aiello, Energy intelligent buildings based on user activity: A survey, Energy Build. 56 (2013) 244–257.
[21] K. Padmanabh, A. Malikarjuna, V.S. Sen, S.P. Katru, A. Kumar, S.P.C.S.K. Vuppala, S. Paul, Isense: A wireless sensor network based conference room management system, Proc. Buildsys (2009) 37–45.
[22] ASHRAE, ASHRAE Guideline 14-2014: measurement of energy, Demand Water Sav. (2014).
[23] P.O. Fanger, Thermal Comfort. Analysis and Applications in Environmental Engineering, Danish Technical Press, Copenhagen, Denmark, 1970.
[24] T. Teixeira, G. Dublon, A. Savvides, A survey of human-sensing: Methods for detecting presence, count, location, track and identity, AXM Comput. Surv. 5 (1) (2010) 59.
[25] Y. De Bock, A. Auquilla, A. Nowé, J.R. Duflou, Non-parametric user activity modelling and prediction, User Model User-Adap Inter 30 (2020) 803–831.
[26] T.S. Ferguson, A Bayesian analysis of some nonparametric problems, Ann. Statist. 1 (1973) 209–230.
[27] C.E. Antoniak, Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems, Ann. Statist. 2 (1974) 1152–1174.
[28] H. Ishwaran, L.F. James, Gibbs sampling methods for stick-breaking priors, J. Amer. Statist. Assoc. 96 (2001) 161–173.
[29] J. Sethuraman, A constructive definition of Dirichlet priors, Statist. Sinica 4 (1994) 639–650.
[30] N.C. Truong, J. McInerney, L. Tran-Thanh, E. Costanza, S.D. Ramchurn, Forecasting multi-appliance usage for smart home energy management, Proc. IJCAI (2013) 2908–2914.
[31] Engineering ToolBox, Thermal Conductivity of common Materials and Gases, 2003. https://www.engineeringtoolbox.com/thermal-conductivity-d_429.html. (Accessed 6 2018).
[32] J. Duflou, A. Auquilla, Y. De Bock, A. Nowé, K. Kellens, Impact reduction potential by usage anticipation under comfort trade-off conditions, CIRP Annals 65 (1) (2016) 33–36.
[33] T. Day, Degree-Days: Theory and Application, The Chartered Institution of Building Services Engineers, London, 2006, p. 106.
[34] V. Muggeo, Segmented an R package to fit regression models with broken-line relationships, R News 8 (1) (2008) 20–25.
[35] B.P. Weidema, C. Bauer, R. Hischier, Mutel, C, T. Nemecek, J. Reinhard, C. Vandenbo, G. Wernet, Overview and Methodology. Data Quality Guideline for the Ecoinvent Database Version 3, Ecoinvent Report 1(v3), St. Gallen, The ecoinvent Centre, 2013.
[36] M. Goedkoop, R. Heijungs, M. Huijbregts, A. De Schryver, J. Struijs, R. van Zelm, (Update 2013) ReCiPe 2008, Report 1: Characterisation, http://www.lcia-recipe.net/.
[37] Y. De Bock, A. Auquilla, K. Kellens, D. Vandevenne, A. Nowé, J.R. Duflou, User adapting system design for improved energy efficiency during the use phase of products: case study of an occupancy-driven, self-learning thermostat, Sustain. Through Innov. Product Life Cycle Design (2016) 883–898.
[38] ASHRAE. Standard 55-2013, in: Thermal environmental conditions for human occupancy, 2013.
[39] S. Hoyt T. Schiavon, A. Piccioli, T. Cheung, D. Moon, K. Steinfeld, CBE Thermal Comfort Tool, Center for the Built Environment, University of California Berkeley, 2017, http://comfort.cbe.berkeley.edu/.
[40] P. Dongmei, L. Zhongping, N. Mao, S. Mengjie, A study on the effects of different bedding systems on thermal comfort–quantifying the sensitivity coefficient used for calculating predicted mean vote (PMV) in sleeping environments, Energy Procedia 142 (2017) 1939–1946.
[41] N. Kántor, J. Unger, The most problematic variable in the course of human-biometeorological comfort assessment - The mean radiant temperature, Open Geosci. 3 (1) (2011) 90–100.
[42] M. Wenninger, J. Schmidt, T. Goeller, Appliance usage prediction for the smart home with an application to energy demand side management. And why accuracy is not a good performance metric for this problem, Proc. Smartgreens (2017) 143–150.
[43] Y. De Bock, A. Auquilla, K. Kellens, A. Nowé, J.R. Duflou, Intelligent occupancy-driven thermostat by dynamic user profiling, in: Electronics Goes Green 2016+ (EGG), IEEE, 2016, pp. 1–8.