Predicting household occupancy for smart heating control: A comparative performance analysis of state-of-the-art approaches

Wilhelm Kleiminger[a,*], Friedemann Mattern[a], Silvia Santini[b]

[a] *Institute for Pervasive Computing, ETH Zurich, Universitätstrasse 6, 8092 Zurich, Switzerland*
[b] *Wireless Sensor Networks Lab, TU Darmstadt, Rundeturmstrasse 10, 64283 Darmstadt, Germany*

Abstract

This paper provides a comparative study of state-of-the-art means of predicting occupancy for smart heating control applications. We focus on approaches that predict the *occupancy state* of a home using occupancy *schedules* – i.e. past records of the occupancy state. We ran our analysis on actual occupancy schedules covering several months for 45 homes. Our results show that state-of-the-art, schedule-based occupancy prediction algorithms achieve an overall prediction accuracy of over 80%. We also show that the performance of these algorithms is close to the theoretical upper bound expressed by the *predictability* of the input schedules. Building upon these results, we used ISO 13790-standard modelling techniques to analyse the energy savings that can be achieved by smart heating controllers that use occupancy predictors. Furthermore, we investigated the tradeoff between achievable savings (typically 6% to 17% on average) and the risk of comfort loss for household residents.

Keywords: occupancy detection, occupancy prediction, smart heating, energy management, smart home, energy efficiency, thermostat strategy, heating setback

*Corresponding author

Email addresses: `wilhelmk@ethz.ch` (Wilhelm Kleiminger), `mattern@ethz.ch` (Friedemann Mattern), `santinis@wsn.tu-darmstadt.de` (Silvia Santini)

1. Introduction

The ability to detect whether or not a house is occupied – i.e. whether residents are at home or not – represents a basic requirement for the operation of many home automation systems. For instance, the presence of at least one resident within a home might trigger the operation of a lighting control system [1]. Similarly, the absence of all residents allows a heating control system to automatically lower the temperature of the home [2, 3], thereby saving energy that would have been unnecessarily used for heating. Since space heating accounts for a large fraction of residential energy use (e.g. 68% in the European Union member states [4]), smart thermostats could thus play an important role in reducing costs and carbon dioxide emissions. Besides the ability to determine whether or not a house is occupied, many home automation systems also need to be able to *predict* when a house is going to be occupied. For instance, a heating control system may require some time to heat a home to a comfortable temperature after its residents have been out for the day. In order to avoid a loss of comfort for the residents – i.e. the house being too cold when they return – the heating needs to be triggered at the right time. However, preheating the house for too long in advance will result in wasted energy.

Both *occupancy detection* and *occupancy prediction* can thus be regarded as basic services upon which many home automation systems need to rely. While such systems[1] enable a large number of applications, this study focuses on the particular scenario in which such services support the operation of "smart" heating control systems. Although several ways of supporting such systems have been presented in the literature, no systematic review of existing techniques has previously been conducted. In particular, notations and terminology are often inconsistent across different contributions, making it hard to compare existing approaches in a qualitative way. Quantitative comparisons are also often impracticable due to the lack of both a common, freely available dataset upon which to base a comparative study as well as the wide variety of scenarios for which different approaches might need to be tested.

In this paper, we address the above-mentioned issues by providing the

[1]Interestingly, a rather general patent "Occupancy pattern detection, estimation and prediction" (US 8510255) has recently been granted to the home automation company Nest – acquired by Google in 2014 and makers of stylishly designed self-learning thermostats.

following contributions: (1) A classification and review of state-of-the-art approaches that predict home occupancy. We outline different techniques used in the literature and identify two main classes (*schedule-based* and *context-aware*) into which existing approaches can be categorised. (2) A quantitative comparison of the performance of selected schedule-based occupancy prediction algorithms. The performance evaluation is based on actual occupancy data for 45 individuals collected over several months. We derived this occupancy data by analysing mobile phone records collected as part of the Lausanne Data Collection Campaign (LDCC) [5].

Several other studies have reviewed the existing literature on occupancy detection and prediction. For instance, Nguyen et al. [6] provide an extensive review of approaches that address the broad topic of "energy intelligent buildings". Guo et al. [7] focus on smart lighting control approaches. While both these studies mention performance figures for the approaches they survey, the numbers in question originate from the papers being surveyed and are thus typically obtained in very different experimental settings. Instead, we provide a quantitative performance analysis based on a common dataset. As all algorithms operate on the same data, the performance figures obtained can be accurately compared.

In order to put our study into its proper context, Sections 2 and 3 provide basic notions regarding smart heating and also occupancy detection and prediction. Our review and classification of existing methods is then presented in Section 4. Section 5 describes the experimental setup. Section 6 discusses the results of our comparative performance analysis and Section 7 mentions some limitations of the modelling technique. Finally, Section 8 summarises the main findings of our study.

2. Smart heating control

The idea of using information and communication technology to automatically and "intelligently" control heating systems has been investigated for several years. Well-known examples of such *smart heating* approaches include the Neurothermostat [1], the GPS Thermostat [8], the Smart Thermostat [2] and several others [3, 9, 10, 11, 12, 13]. The first few commercial products – such as the NEST learning thermostat, tado° and EcoBee's Smart-Si – have recently started to appear.[2]

[2] `www.nest.com`, `www.tado.com/en/`, `www.ecobee.com/solutions/home/smart-si/`.

A smart heating system should meet two main requirements. First, it should significantly reduce the amount of energy spent on heating (compared with conventional room heating systems). Secondly, it must ensure adequate *thermal comfort* – which the ANSI/ASHRAE Standard 55 defines as "the condition of mind that expresses satisfaction with the thermal environment" [2, 14].

The smartness of the system typically lies in its ability to adapt to current environmental conditions, the specific household characteristics and the behaviour of the occupants. The difference between a conventional automatic (or programmable) heating system and a "smart" one is that while the former operates according to a pre-defined and typically deterministic (e.g. timer-based) schedule, the latter typically adapts its control strategy to the user context. In both cases, though, the heating[3] is controlled automatically, i.e. with the aid of a thermostat that does not require explicit human intervention.

An automatic heating control system can be seen as a regulator that ensures that the (average) air temperature measured within a home is sufficiently close to a given target value. To this end, the system controls the activation and deactivation of the heaters available in the home (e.g. heat pumps and/or electrical heaters). Typically, at least two different target temperatures are defined: the *setback temperature* and the *comfort* (or *setpoint*) *temperature*, indicated as Θ_{setb} and Θ_{comf} respectively. Θ_{comf} is typically set by household occupants depending on their personal preferences and indicates the temperature at which they feel comfortable. The value of Θ_{comf} will typically be around $21\,°C$. The setback temperature Θ_{setb} in contrast is defined as the lowest (average) value at which the air temperature of the household is permitted to fall when the occupants are out (or asleep). There are several issues that need to be considered when setting suitable values for the setback temperature. In particular, Θ_{setb} must be sufficiently low to allow for significant energy savings (as the heaters can be – at least temporarily – be deactivated) but still high enough that the time needed to bring the household back up to Θ_{comf} does not exceed a reasonable value. For a

[3]Note that the assessment of thermal comfort according to the ANSI/ASHRAE Standard 55:2010 [14] requires parameters other than air temperature to be additionally considered, e.g. humidity. However, with respect to the discussion of occupancy detection and prediction algorithms upon which this paper focuses, there is no loss of generality in limiting our consideration to air temperature only.

more detailed discussion of this issue, the interested reader is referred to [2] and references therein. We will consider $10\,°C$ as a typical value for a deep setback Θ_{setb} when a house is unoccupied.

An optimal heating system should thus be able to maintain the temperature of a home at Θ_{setb} for as long as possible, so as to reduce the amount of energy spent on heating. At the same time, the system must ensure that the temperature is close to Θ_{comf} whenever at least one occupant is at home (and awake) – so as to avoid any loss of comfort. However, the time needed to bring the home from Θ_{setb} to Θ_{comf} (and vice versa) is typically non-negligible (e.g. > 1 hour). An optimal heating system therefore needs to be able to both immediately detect when the home becomes unoccupied – so as to to turn off the heating – and also reliably predict when it will be occupied again – in order to restore the temperature to Θ_{comf} by the time the occupants return.

Smart heating systems try to approximate this behaviour by putting in place adequate procedures to both detect and predict the household occupancy state. Different approaches can largely be differentiated on the basis of the technique they use to implement such procedures and the sensor data they require to do so. Before discussing state-of-the-art approaches in Section 4 we will therefore briefly summarise in the next section the basic concepts used in the occupancy detection and prediction literature.

3. Occupancy detection and prediction

A house is said to be *occupied* at a time instant t if at least one of its residents is at home; otherwise, it is said to be *unoccupied*. The *occupancy state* of a house can thus be represented as a binary value (1 for occupied and 0 for unoccupied).

The household occupancy state at any given time can be determined by interrogating sensors deployed within the home, such as passive infrared (PIR) or light sensors. Data from electricity meters can also provide clues regarding human activity – and thus the presence of residents – within a home [15, 16]. However, as outlined in [17], each type of sensor has its own advantages and drawbacks and can only guarantee limited confidence in estimating the actual occupancy state. Also, the deployment and maintenance of sensors within a home may generate significant costs and inconvenience for residents.

Another strategy for detecting household occupancy consists of interrogating sensors carried by the residents, such as RFID tags, dedicated wireless transmitters or GPS modules embedded in mobile phones [3, 18]. For the performance analysis presented in Sections 5 and 6, we used occupancy data derived from the analysis of mobile phone records.

To represent the historical occupancy states of a home, it is usually convenient to divide the hours of the day in N_s equally spaced intervals – called *slots*. An *occupancy vector* Γ is then a $1 \times N_s$ vector of binary values in which the ith element indicates whether the home was occupied or unoccupied during slot i. More specifically, we use $\Gamma_{1..96}$ to denote a 24-hour *ground truth* occupancy vector based on 15-minute timeslots and $\gamma_{1..96}$ to refer to a 24 hour *predicted* occupancy vector. Accordingly, an *occupancy schedule* is a $N_d \times N_s$ matrix containing occupancy data for N_d consecutive days. To accommodate slots for which no data is available, occupancy states can also be represented using three – rather than two – symbols, where one symbol is reserved to represent an *unknown* occupancy state.

Conventional programmable thermostats operate according to user-defined schedules. Their settings need to be changed manually as the residents' occupancy schedules vary. Smart heating systems seek to overcome this need for manual re-programming by predicting household occupancy and supplying the control schedules to the thermostat without any direct user involvement. So when the occupants leave the building, the heating may be switched off automatically and the temperature allowed to drop to Θ_{setb}. However, this *reactive* strategy fails when the occupants return, as the thermal properties of the house will result in a certain *time lag* until the comfortable temperature Θ_{comf} is reached again. The time lag describes the time taken by the heating system to reach Θ_{comf} from the current *indoor air temperature* Θ_{air}. The longer the house has been left unoccupied and the temperature has been allowed to drop, the greater the time lag will be. Therefore, at any given time, if the occupants have left the household, the system needs to know how long it would take to re-heat the property to Θ_{comf} and whether the house is likely to be occupied within this time span. We call the time slots involved in this calculation the *prediction horizon* I^*. Amongst other variables, I^* is determined by the current indoor air temperature Θ_{air}, the target comfort temperature Θ_{comf} and the forecast for the *outside temperature* Θ_e. We refer to the process of computing the future occupancy states within I^* as *occupancy prediction*.

4. An overview of occupancy prediction algorithms

Several occupancy detection and prediction algorithms for smart heating control have been proposed in the literature [2, 3, 8, 9, 12, 18, 19, 20]. Occupancy detection algorithms rely on a relatively small number of basic techniques. For instance, detection is performed using sensors installed in the home – such as PIR, RFID or camera sensors [2, 3, 12, 19] or by leveraging GPS modules, which can usually be found in smartphones carried by the home's occupants [8, 18]. For occupancy prediction, different mathematical models – including artificial neural networks [1] and Markov chains [12] – are used. In the following, we focus mainly on occupancy prediction algorithms; a discussion of occupancy detection is provided by [7].

4.1. Schedule-based approaches

Several approaches compute occupancy predictions relying on past occupancy schedules only [2, 3, 18, 21]. Such approaches, which we refer to as *schedule-based* algorithms, take as input *historical* data on the household occupancy state. This data is typically collected over an extended period of time (weeks to months).

The *PreHeat* (PH) algorithm presented by Scott et al. [3] is an example of a schedule-based approach. PreHeat maintains a vector for storing the actual occupancy state registered for the current day starting from midnight. Each element of the vector represents the occupancy state of the home in a 15-minute interval. An element is set to 1 or 0 depending on whether the house is occupied or not during the relevant time interval. To compute an occupancy prediction from a given time of day onwards, PreHeat first computes the Hamming distance between the occupancy pattern thus far observed for the current day and the corresponding segments of past occupancy vectors. The k past vectors with the lowest Hamming distances are then selected (k is fixed and equal to 5 in [3]) and averaged element-by-element. These averages approximate to the probability for the home being occupied during the corresponding time interval. The actual prediction is computed assuming that the house will be occupied during a future time interval if the corresponding probability exceeds a given threshold α, or else unoccupied. In [3], the value of α is fixed and equal to 0.5. Building upon this basic version of the algorithm, Scott et al. introduce two additional features. The first consists of differentiating between weekdays and weekends. The second is to pad the current occupancy vector with data for the 4 hours before and after midnight, taken from the previous and following day respectively. This helps the algorithm to predict past midnight. Once the prediction is computed, the algorithm decides whether to start heating. This control decision depends on a number of factors including the current and desired temperatures as well as the rate (in terms of degrees per hour) at which the house can actually be heated.

The *Presence Probabilities* (PP) approach presented by Krumm and Brush is another well-known schedule-based approach [18]. Household occupancy is detected using a GPS device carried by the residents. The home is assumed to be occupied if the device indicates that a resident is less than 100 meters away from it. Using the GPS data, PP computes the probability for a home being unoccupied – called p_{away} – during any time slot of a day of the week. The values of p_{away} in slots are computed using the ratios between the number of GPS data points that lie outside the 100-meter radius of the home and the total number of GPS data points available for the slot. The value of p_{away} for each time slot is stored in a vector called p_{week} containing 336 elements (7 days a week, 48 slots a day). The probability within each slot is smoothed using the values of the previous and subsequent slots. To adjust the values of p_{away} for weekdays, a generic vector $p_{weekday}$ that contains the average values of p_{away} for a "generic" weekday is used. Using a regularisation factor λ_{wd} this vector can account for "*greater or lesser variability on weekdays*" [18]. The values of p_{away} in each slot of the final *probabilistic schedule* $\widetilde{p}_{week}$ are then computed as the sum of the elements of p_{week} and the relevant elements of $p_{weekday}$. In our paper, we refer to the version of the Presence Probabilities algorithm described above as PP and to a simplified version that does not consider smoothing or the generic weekday schedule as PPS.

The *Smart Thermostat* (ST) by Lu et al. [2] also relies on historical schedules to predict occupancy. The occupancy state of a home is determined using a Hidden Markov Model. The model allows an estimate of whether the home is occupied or not and in the former case also whether the occupants are asleep or active. To compute the estimation, the Hidden Markov Model takes as input both prior information derived from historical schedules and actual data collected by several sensors deployed within the home (e.g. PIR sensors). The model is trained using a set of actual past occupancy schedules and sensor data traces. When the house is classified as unoccupied, ST switches the heating system off and allows the temperature of the household to fall to a "deep" setback temperature. If the occupants were to come back home unexpectedly while the house was at the deep setback temperature they would experience a significant comfort loss. ST thus keeps records of all previously observed arrival times (i.e. the time instants at which the house became occupied again after a period of absence).[4] The minimum of such previous arrival times is set as the time by which the household must be preheated to at least a "shallow" setback temperature. This mechanism makes it possible to reduce the risk of comfort loss. ST also estimates the optimal time instant t^* – called the *preheat time* – at which the heating system must be activated to preheat the house. The preheat time t^* is chosen so as to minimise the average amount of energy wasted to heat the household and maintain it at the comfort temperature when the occupants are out. To identify the preheat time for a given day, ST considers all arrival times $\underline{a} = [a_0, a_1, \dots, a_n]$ observed on previous days. Then it considers all time instants $t \in [max(\underline{a}), min(\underline{a})]$ for the current day as candidate preheat times. For each $t_i \in [max(\underline{a}), min(\underline{a})]$, the system computes the amount of energy waste $w_j(t_i)$ that would occur if t_i were the preheat time and the household were to be occupied again at arrival time a_j. The expected average energy waste that would occur if t_i were the preheat time is then the average: $w(t_i) = \sum_{j=1}^{n} w_j(t_i)$. The preheat time is chosen as the time instant that minimises the expected average energy waste: $t^* = argmin_{t_i \in [max(\underline{a}), min(\underline{a})]} w(t_i)$. The occupancy prediction mechanism of ST thus requires the identification of arrival times based on past schedules. Both the minimum of these arrival times and their weighted average are used to trigger different stages of the heating system. For the computation of the amount of energy waste, ST assumes a three-stage heating system and the availability of knowledge about the energy consumed by each stage.

Our comparative study focuses on schedule-based approaches and includes both the PP (or PPS) and PH algorithms. In place of ST itself we instead considered two heuristic prediction strategies – called *Mean Arrival Time* (MAT) and *Minimum Distance Mean Arrival Time* (MDMAT) – which mimic the occupancy prediction algorithm used by ST. As described above, ST uses the minimum of all previously observed arrival times as the time instant at which the household has to change from deep to shallow setback. ST also heats the house to the comfort temperature using a policy

[4] Although this is not specified explicitly in [2], we assume that only one arrival event per day is considered.

that minimises energy waste. To this end, a three-stage heating system with different efficiencies for each stage is assumed to be in place. In our study, we analysed performance (e.g. efficiency gain) in terms of occupancy prediction separately from that due to the specific heating strategy. Also, we assume a single-stage heating system. Thus, ST would always choose the latest observed arrival time as the preheat time. This is due to the fact that heating reactively guarantees the lowest energy waste when comfort loss is not considered and a single-stage heating system is in place.

We therefore introduce the MAT and MDMAT methods as adaptations of ST's preheating strategy. Like ST, the MDMAT algorithm records all n observed arrival times in a vector $\underline{a}$. For each $a_i \in \underline{a}, i = 1, \ldots, n$, MDMAT calculates the distance to all other arrival times $a_j \in \underline{a}, j \neq i$ as $d(a_i) = \sum_{a_j \in \underline{a}, j \neq i} \min(|a_i - a_j|, |a_i - (a'_j + 24)|)$. The *most likely* arrival time for the current day is then chosen as $a^* = \arg\min_{a \in \underline{a}} d(a)$. MAT instead computes the expected arrival time for each day as the arithmetic mean of the arrival times recorded on all previous days. To this end, only one arrival time per day is considered. This is selected as the first arrival event after 2 p.m. and before 2 a.m. We impose this restriction to limit the effect of outliers (e.g. unusual arrival events in the morning) and to avoid the computation of the arithmetic mean of the arrival times causing misleading results due to the use of a 24-hour interval.[5] In contrast to ST's original strategy, which targets a reduction in energy consumption, MAT and MDMAT trade off energy efficiency against comfort loss.

In this paper, we do not describe any other existing schedule-based occupancy prediction algorithms in detail, but refer the interested reader to [1, 20, 21, 22, 23]. For the sake of completeness, however, we outline in the next section selected approaches that detect and predict occupancy for smart heating control using techniques other than those described above and summarised in Table 1.

[5] For example, given two arrival events – one at 1:00 a.m. and one at 9:00 p.m. (21:00), their arithmetic mean computed over a 24-hour interval (from 00:00 to 24:00) would return the value 11 a.m., although the desired mean value would be 11 p.m.

Table 1: Algorithms considered for the comparative performance analysis.

Acronym	**Name**	**Source**
PH	PreHeat	[3]
PP	Presence Probabilities	[18]
PPS	Presence Probabilities Simplified	[18]
MAT	Mean Arrival Time	Emulating ST [2]
MDMAT	Minimum Distance Mean Arrival Time	Emulating ST [2]

4.2. Other approaches

Several authors have proposed techniques that estimate the future occupancy state of a home by observing the current context of its occupants. We refer to these techniques as *context-aware* approaches, since they depend on the current context (e.g. location or activity) of the user, rather than the home's historical occupancy schedule. One example of this is the algorithm presented by Gupta et al. [8], which estimates the time at at which residents will return home based on their current position and driving trajectory. The position is determined using GPS modules embedded either in dedicated devices or in occupants' mobile phones. A web-based mapping service is used to determine the distance from home and the corresponding remaining *drive time*. The thermostat is then instructed to preheat the home if the remaining drive time is less than a given threshold. In [18], Krumm and Brush show how to combine their Presence Probabilities algorithm with Gupta et al.'s drive time prediction approach. In contrast to [8], Krumm and Brush allow drive times to be pre-computed, thereby increasing efficiency but reducing accuracy, particularly in areas prone to congestion. In an earlier paper [24], Krumm et al. also introduced a method called *Predestination*. This method uses historical data along with information on a user's driving habits to obtain the most likely next destination. A similar system, TherML, is presented by Koehler et al. [25]. TherML utilises a hybrid prediction algorithm that switches between predicting the next destination and static schedules based on the user's mode of travel (stationary, walking or driving). Other approaches such as [21], [26] and [27] also use context information about the user to predict where he/she is likely to go next.

A number of occupancy detection and prediction approaches focus not only on heating but also on ventilation. For instance, Erickson et al. propose a system that controls both ventilation and heating/cooling in an office building [12]. The system estimates the occupancy level of different rooms – i.e. the number of people present in each room at any given time – using

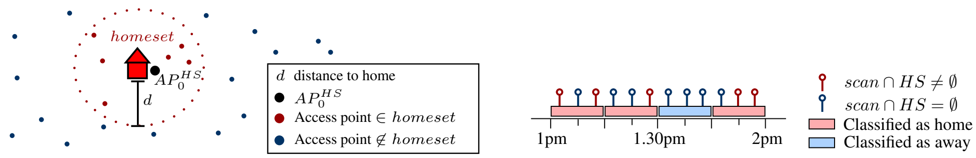

(a) Obtaining the homeset.

(b) Interval classification.

Figure 1: (a) Obtaining the homeset from a set of Wi-Fi access points in the vicinity of the home access point AP_0^{HS} and (b) classifying intervals based on the homeset data.

a Markov Model. The model takes as input both prior occupancy level data and contextual information on movements between rooms. To detect such *transitions*, a network of 16 cameras is used to monitor the so-called *transition boundaries* (e.g. corridors). As the probability of a transition occurring correlates to the time of day, the transition probabilities between different occupancy states are computed on an hourly basis.

5. Setting up the comparative performance analysis

Schedule-based algorithms represent an important category of approaches for predicting occupancy. The goal of our comparative study is to evaluate and discuss the performance of a representative subset of these algorithms. We considered the algorithms listed in Table 1 and conducted our study using the methodology described below.

5.1. Actual occupancy schedules

To compare the performance of different occupancy prediction algorithms in a consistent manner, we evaluated them using a large dataset of actual occupancy schedules. We inferred these schedules using sensor data collected as part of the Nokia Lausanne Data Collection Campaign (LDCC) [5]. To the best of our knowledge, no publicly available data existed on long-term, high-granularity occupancy schedules, making it necessary to build such schedules in order to conduct our evaluation.

The LDCC dataset contains about 18 months' worth of traces of Wi-Fi scans, GPS coordinates, accelerometer readings and several other sensors, as well as demographic information from mobile phone users [5]. However, the dataset does not contain any information concerning user-relevant locations, i.e. it is not known where the user's home, office, etc. are located. We therefore developed a technique, called the *homeset algorithm* [28], to infer this information from the available LDCC data.

The goal of the homeset algorithm is to infer when each user was at home and when they were not during the data collection period. Thus, the algorithm computes the occupancy schedule of each user. To do so, it only uses records of visible Wi-Fi access points. During the LDCC, mobile phones were set to regularly scan for the presence of visible Wi-Fi access points (APs) in the immediate vicinity of the phone (and therefore the user). After each scan, the phone stored information about the detected APs along with a corresponding timestamp. The input data for the homeset algorithm consists of a list of these records, from which only the timestamps and the identifier (MAC address) of the APs are used by the algorithm. A single Wi-Fi scan is thus represented as a tuple $< t_k, AP_0, AP_1, \ldots, AP_{m_k-1} >$ where t_k is the timestamp at which the k-th scan was performed, m_k is the total number of APs detected during the k-th scan and AP_i, $i = 0 \ldots, m_k - 1$ are the MAC addresses of the APs. The homeset algorithm uses these scans to identify a set of access points that are located within, or in the immediate proximity of, the mobile phone owner's home. We call this set the *homeset* (HS) and assume it contains n access points, such that $HS = \{AP_0^{HS}, AP_1^{HS}, ..., AP_{n-1}^{HS}\}$. Given a Wi-Fi scan $< t_k, AP_0, AP_1, \ldots, AP_{m_k-1} >$, the homeset algorithm tests whether $\{AP_0, AP_1, AP_2, ..., AP_{m_k-1}\} \cap HS \neq \emptyset$. If this is the case, the algorithm assumes the home to be occupied in the time slot identified by the time-stamp t_k. Figure 1(b) illustrates the rationale behind the homeset algorithm.

To bootstrap the homeset algorithm, we determine for each user the AP that has the highest empirical probability of being detected at least once between 3 a.m. and 4 a.m. on any particular night. This AP is set to be AP_0^{HS}. This procedure assumes that typical users spend most of their nights at home. Once AP_0^{HS} has been identified, the homeset HS is constructed by adding to HS any other APs that appear in a Wi-Fi scan together with AP_0^{HS}. Simple heuristics are used (e.g. the number n of APs in HS is restricted) to improve the robustness and reliability of the algorithm [28].

5.2. Preparing the schedule

For the study presented in this paper we used only occupancy schedules for users who had collected data for at least 100 days during the LDCC (i.e. $N_d > 100$) and for whom the occupancy state could be inferred in at least 70% of the slots. This was done to ensure sufficiently large training and test

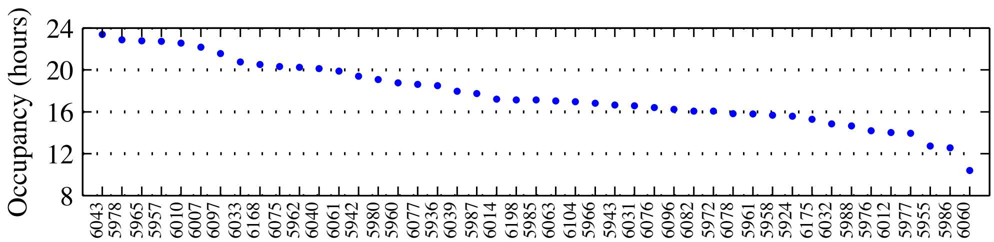

Figure 2: Occupancy in hours for all 45 households in the dataset (identified by the unique LDCC participant number).

sets. We also discarded the schedules of users whose probability of being at home between 3 a.m. and 4 a.m. on weekdays was estimated to be less than 60%. This ensured we considered in the study only users for whom the homeset algorithm could reliably identify the home. This first data cleaning phase enabled us to select 59 occupancy schedules.

The PreHeat algorithm by Krumm et al. imposes additional constraints. For instance, daily schedules need to be padded with four hours from the previous day and four hours from the next day [18]. We consequently discarded from the schedules all days for which this information was not available in order to ensure all algorithms were trained and tested on the same data. This left 45 schedules to be used for our evaluation. Figure 2 shows the average occupancy in hours per day for all the participants in the dataset. On average, these schedules include 74 days' worth of occupancy data, with the participants staying at home for 17 hours and 40 minutes per day on average.

5.3. Building model and simulation setup

The algorithms analysed in this paper aim to predict occupancy for smart heating control systems. The goal of such systems is to reduce the energy consumed by heating, while at the same time avoiding any loss of comfort for the residents. We therefore assessed the suitability of the prediction algorithms in terms of their ability to save energy and ensure comfortable temperatures when required. To this end, we built a predictive controller to control the temperature of a building based on the current occupancy state and the algorithms' predictions of the future occupancy states of the building. In order to analyse the performance of the controller under different conditions, we ran simulations using the 5R1C thermal building model from the ISO 13790 energy performance standard [29] on 32 different scenarios. In particular, we analysed the influence of different weather conditions, building sizes and insulation levels.

The ISO 5R1C model simulates the transient heat conduction between the property and its surroundings using an analogous electrical resistance-capacitance (RC) circuit and thus offers a method of calculating the energy required for heating and cooling while maintaining specified setpoint temperatures. This modelling principle was first introduced by Beuken in 1936 [30] and has since been widely employed in building design [31]. In contrast to simpler models [1], the ISO 5R1C model takes into account the heat transfer by transmission and ventilation as well as solar and internal gains.

The response of the heating system was simulated for 32 different weather and building settings. We considered two different building sizes – a 52 m^2 studio flat (F) and a 176 m^2 house (H). In order to measure the effect of the building envelope on thermal performance, we also simulated the response of the ISO 5R1C model for low and high U-values[6]. The U-value (W/m^2K) denotes the overall heat transfer coefficient of a building element. Elements with high U-values conduct more heat per unit temperature difference between the inside and outside. A building with high U-values is considered poorly insulated and thus leaking a significant amount of heat to the outside. For each of the resulting four building configurations (flat F-$\mathrm{U}_{\mathrm{low}}$, F-$\mathrm{U}_{\mathrm{high}}$; house H-$\mathrm{U}_{\mathrm{low}}$, H-$\mathrm{U}_{\mathrm{high}}$), the design heat load (maximum heat input) in watts $\Phi_{H,max}$ was determined using the DIN EN 12831 standard [33]. The internal gains Φ_{int} were assumed to be 250W and 375W, whenever the house was occupied, equivalent to the metabolic heat rate of two and three residents for the flat and house respectively. Table 2 shows the parameters for the ISO 5R1C model for all the building variants we analysed.

The effect of different weather conditions on the heating load was captured by eight representative weather scenarios synthesised from real weather data[7] for the Lausanne (Switzerland) area where also the data used to derive

[6]The U-values for a well-insulated buildings (F-$\mathrm{U}_{\mathrm{low}}$ and H-$\mathrm{U}_{\mathrm{low}}$) correspond to the maximum allowed U-values for new properties in Germany according to EnEV'14 [32]. For the poorly insulated buildings (F-$\mathrm{U}_{\mathrm{high}}$ and H-$\mathrm{U}_{\mathrm{high}}$), we used a list of high U-values reported in `http://en.wikipedia.org/wiki/Thermal_transmittance` (accessed on May 8, 2014).

[7]Global solar radiation and outdoor temperature (2m above ground) were obtained from MeteoSwiss; the global radiation was split into direct and indirect radiation using the Reindl* method [34].

Table 2: ISO 5R1C building model parameters for different building variants.

	Building variant				
Parameter	**F-U$_{\mathbf{low}}$**	**F-U$_{\mathbf{high}}$**	**H-U$_{\mathbf{low}}$**	**H-U$_{\mathbf{high}}$**	**Units**
Thermal transmission coefficient for opaque building elements – $H_{tr,op}$	47.16	184.57	103.57	379.35	W/K
Thermal transmission coefficient for windows and doors – $H_{tr,w}$	12.68	31.50	33.07	102.06	W/K
Thermal transmission coefficient for ventilation – H_{ve}	47.33	47.33	161.57	161.57	W/K
Internal zone capacitance – C_m	8.51	8.51	29.04	29.04	MJ/K
Floor area – A_f	51.56	51.56	176.00	176.00	m^2
Design heat load according to [33] – $\Phi_{H,max}$	2.80	6.86	7.78	16.75	kW

Table 3: Weather scenarios. For each of the 8 scenarios, the table shows the average daily temperature $\Theta_{e,d}$ and the average daily global radiation I_{avg} for reference.

		$\Theta_{e,d}$ (°C)		I_{avg} (W/m^2)	
Scenario	**Range**	**clear**	**cloudy**	**clear**	**cloudy**
Very low temperature	$-6\,°C \leq \Theta_{e,d} \leq -4\,°C$	-5.4	-4.7	142.9	35.5
Freezing temperature	$-1\,°C \leq \Theta_{e,d} \leq 1\,°C$	0.1	0.0	137.5	30.2
Low temperature	$4\,°C \leq \Theta_{e,d} \leq 6\,°C$	5.1	5.1	148.5	26.1
Moderate temperature	$9\,°C \leq \Theta_{e,d} \leq 11\,°C$	10.1	10.0	180.7	29.7

the occupancy schedules was gathered (cf. Section 5.1). Lausanne is situated within a transition zone between a humid oceanic climate zone and a continental temperate zone.

Table 3 shows the eight weather scenarios used in the evaluation. The scenarios cover four different temperature levels under clear as well as cloudy sky conditions. Each scenario consists of 24-hour vectors of the outside temperature and the direct solar radiation, replicated n times to reflect the number of days in the occupancy data. The vectors are the average of multiple days fitting the temperature ranges shown in Table 3. We have not included a detailed description of the methodology used to define the weather scenarios and refer the interested reader to the supplementary technical report [35].

5.4. Heating controller

We implemented a predictive heating controller to translate the occupancy schedules predicted by the algorithms into actual heating schedules. A heating schedule defines the *target indoor air temperature* $\Theta_{air,set}$ at 15-minute time intervals t. Given the predicted occupancy schedule and the RC model, the heating controller sets $\Theta_{air,set}$ to Θ_{comf} for t if: (1) The house is occupied at time t (reactive policy); (2) The house is expected to become

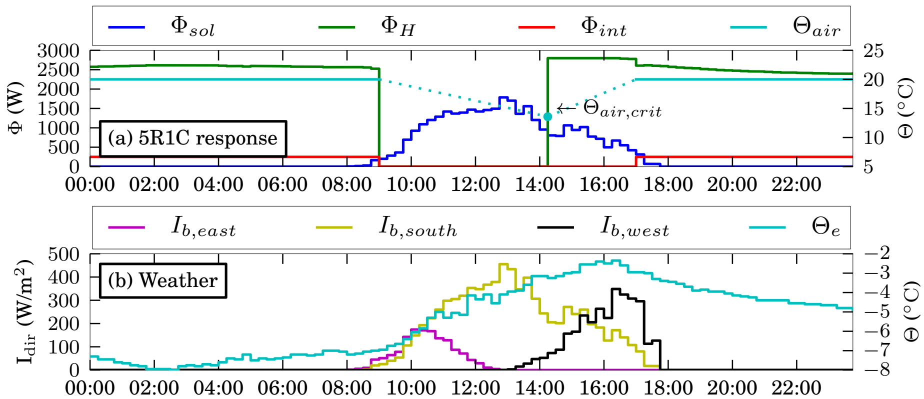

Figure 3: Typical behaviour of a heating system according to the ISO 5R1C model (F-U$_{\text{low}}$, very low temperature, clear sky) for a scenario where the house is unoccupied between 9 a.m. and 5 p.m. The upper part shows the inputs (solar gain Φ_{sol}, heat input Φ_H and internal gain Φ_{int}), the lower part the direct radiation $I_{b,\{east,south,west\}}$ and outside temperature Θ_e. $\Theta_{air,crit}$ denotes the critical temperature at which the preheating starts to reach Θ_{comf} at 5 p.m.

occupied between $t+1$ and $t+I^*$. The prediction horizon I^* (cf. Section 3) is the time needed to raise the indoor air temperature Θ_{air} to Θ_{comf} (predictive policy), starting from the temperature at time $t+1$, using the maximum available heating power $\Phi_{H,max}$ (DIN EN 12831 design heat load) and assuming that the target temperature was Θ_{setb} at time t. If neither of these two conditions is fulfilled, the controller sets the target temperature to Θ_{setb} in order to save energy. The heat input Φ_H at any point in time is directly determined by the current setpoint temperature. In all cases, the controller has perfect knowledge[8] of the future weather.

The predictive heating controller is always in one of three different states: the preheat state, the heating state or the cool down state. If the current air temperature is below the setpoint temperature $\Phi_{air,set}$, the controller is in the *preheat* state where the system heats with $\Phi_{H,max}$, the maximum heating power available. If the current air temperature is equal[9] to the setpoint

[8]The alternative, predicting the future weather in order to determine when to heat, would prevent us from isolating the performance of the occupancy prediction algorithm.

[9]In practice "equal" is often taken with a grain of salt: To avoid excessive switching

temperature, the controller is in the *heating* state. Here the heating power is lower than the maximum value and equivalent to the power needed to maintain the setpoint. Otherwise, if the setpoint is lower than the measured air temperature, the system is in *cool down* state and no heat is added to the system (i.e. $\Phi_H = 0$).

The upper part of Figure 3 shows the behaviour of the controller and the indoor air temperature Θ_{air} for a typical occupancy schedule and the F-U_{low}, freezing temperature, clear sky scenario. The lower part of the figure shows the corresponding weather data ($I_{b,\{east,south,west\}}$ indicating the direct solar radiation and the outside temperature Θ_e) used in this scenario. When the occupants leave at 9 a.m., the indoor air temperature is allowed to drop until 2.15 p.m. (from 20 °C to 13 °C), with no heat being added to the system. The controller then preheats the property such that $\Theta_{air} = \Theta_{comf} = 20\,^{\circ}\text{C}$ when the occupants return home at 5 p.m.

6. Results of the comparative performance analysis

This section presents the results of our study. We first report on the prediction accuracy achieved by the MAT, MDMAT, PP, PPS and PH algorithms for the occupancy schedules derived from the LDCC dataset. We then show that they achieve a prediction accuracy close to the theoretical upper bound defined by the *predictability* of the input schedules. We conclude by highlighting the performance of the algorithms in terms of *efficiency gain* (as a measure of the energy saved) and *comfort loss*.

6.1. Prediction accuracy

We say that a *true positive* prediction occurs when an algorithm predicts a house will be occupied during a time slot k and the house is indeed occupied during that time slot. Likewise, correctly predicting the house to be *unoccupied* corresponds to a *true negative* prediction. *False positive* and *false negative* predictions occur when the household is incorrectly predicted to be *occupied* or *unoccupied*, respectively. If, more formally, tp denotes the number of time slots with a true positive prediction (and likewise for tn, fp and fn), the *prediction accuracy* of an algorithm is defined as $\frac{tp+tn}{tp+tn+fp+fn}$.

and to prevent wear of control equipment, controllers (in particular on-off systems) are typically designed to include hysteresis, effectively substituting the setpoint with a delta interval (the “comfort band”) around the setpoint.

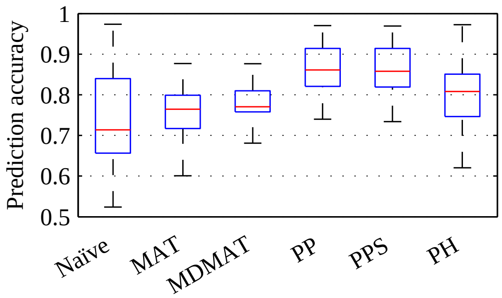

Figure 4: Accuracy of prediction algorithms considered in this study.

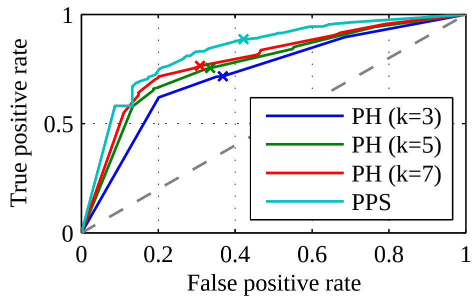

Figure 5: ROC curves of PPS and PH. Crosses indicate $\alpha = 0.5$.

To compare the considered algorithm against a baseline, we introduced a so-called *naïve predictor*. Given the a priori probability p_{occ} of the home being occupied, the *naïve* algorithm always predicts it to be occupied if $p_{occ} \geq 50\%$. If $p_{occ} < 50\%$ the *naïve* predictor always predicts the house to be unoccupied. For our study, we computed p_{occ} from the occupancy schedules as the number of slots containing a 1 in the schedule divided by the total number of slots.[10]

Figure 4 shows the prediction accuracy of all five algorithms considered in this study along with that of the *naïve* predictor for the LDCC occupancy schedules. For each prediction algorithm, the box plot indicates the median as well as the 25th and 75th percentiles of the accuracy across all 45 households. The interquartile range between the top and the bottom of the box thus represents the accuracy achieved in 50% of the homes. The whiskers represent the extreme data points (within $\pm 2.7\sigma$).

The median accuracies in Figure 4 show that all surveyed algorithms improved upon the baseline provided by the naïve predictor. The PP (or PPS) algorithm achieved the highest prediction accuracy. Its median accuracy lies at around 85%, which means that the algorithm achieves at least this accuracy in 50% of the homes in the dataset. It is also the only algorithm for which the accuracy never dropped below 70%, which is the median value of the naïve predictor. We used Tukey's HSD test [37] at the 95% level in conjunction with a one-way balanced ANOVA to establish that the mean accuracy of the PP algorithm was significantly different to the accuracy of

[10] As noted in [36], the naïve predictor was often quite accurate since typical residents spend a significant amount of their time (60% or more) at home.

the other algorithms (except PPS). The ANOVA assumes the distribution of the accuracy for each algorithm to be normal. Confirmation that this assumption holds for the data under analysis was obtained using a two-tailed Shapiro-Wilk test at the 99% confidence level (p-values between 0.23 and 0.75).

The PH algorithm also achieved a good median accuracy around 80% although it exhibits larger deviations to both sides of the median. This shows that for selected homes, PH can achieve a higher accuracy. For "typical" homes, however, PP was the algorithm that performed best. In contrast, the prediction performance of MAT and MDMAT, which are considered here as representative of the basic techniques used by the ST algorithm was noticeably worse. The whiskers indicate that MAT and MDMAT are not suitable for schedules resulting in high values for p_{occ} (i.e. schedules for users who are almost always or almost never at home). This is due to the fact that for every day, MAT and MDMAT assume a period of absence between the computed mean departure and mean arrival times. A single day containing a 9-hour absence may thus result in a predicted schedule with an implied 63% probability of occupancy. In the case of a house otherwise occupied 90% of the time (i.e. $p_{occ} = 90\%$), this results in a drop in accuracy of 27%.

Figure 5 shows the receiver operating characteristic (ROC) curves for the PH and PPS algorithms. The curves highlight the tradeoff between the *true positive rate*, defined as $tp/(tp + fn)$ and the *false positive rate*, defined as $fp/(fp + tn)$. The gray dotted line shows the performance of the random predictor (i.e. tossing a coin). The curves are obtained by varying the value of the threshold α (cf. Section 4.1). The cross markers on the curves show the data points corresponding to $\alpha = 0.5$. For both PH and PPS, setting $\alpha = 0.5$ as done in [3] achieved a good balance between true positive and false positive rates. The figure also shows how the performance of the PH algorithm changes for different values of the parameter k (which represents the number of nearest neighbours taken into account when making the prediction). For $\alpha = 0.5$ and $k = 7$, PH achieved a higher true positive rate and a lower false positive rate than with other parameter configurations. As mentioned above, this is the configuration we used for PH in this study as well as the default choice proposed in [3]. For the PH algorithm we used a prediction horizon of 90 minutes.

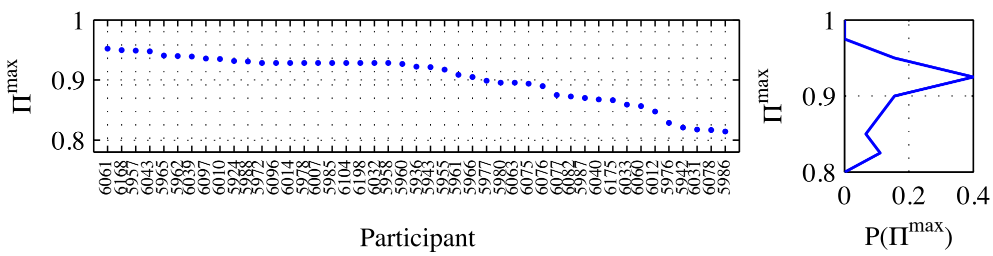

Figure 6: Distribution of predictabilities Π^{max} over all participants.

6.2. Limits of predictability

The results presented above show that among the algorithms considered in this study, the PP predictor achieved the highest median accuracy of 85%. An obvious question to ask would be: Is it possible to do better? In other words, how close is the performance of PP to that of an "optimal" predictor? To answer this question, we built upon the results presented by Song et al. [38]. Their work targets the problem of predicting the next place visited by a person, given that the sequence of places visited thus far – referred to as the *mobility trace* of this person – is known. In this context, they introduce the concept of the *predictability* Π^{max} of a mobility trace $\mathcal{L}$ and show that it represents the "*upper bound that fundamentally limits any mobility prediction algorithm in predicting the next location based on historical records*" [39].

The predictability Π^{max} thus corresponds to the upper limit of the prediction accuracy achievable by schedule-based predictors. If the focus is on occupancy prediction, the next place visited by the participant in the LDCC dataset can either be home or "any place but home." We refer to these two places as L_1 and L_0 respectively. The sequence of places visited by a participant up to a time slot k can then be derived from the schedules. A value of 0 (or 1) in the schedule indicates that the place L_0 (or L_1) has been visited. For instance, assuming 15-minute slots, an excerpt of a schedule indicating a participant is at home for 1 hour and then away from home for 30 minutes corresponds to the sequence $L_1L_1L_1L_1L_0L_0$. In this way, we can derive the mobility trace for each participant and directly apply the method proposed by Song et al. to compute predictability values.

Figure 6 shows the predictability values of the schedules for the 45 participants considered in this study (left) along with the corresponding empirical distribution (right). The predictability is computed for each participant over

Table 4: ISO 13790 average efficiency gain for all experiments with **low U-values (good insulation)**. ☼ and ☁ denote *clear* and *cloudy* scenarios respectively. The rightmost column shows the average total daily energy consumption when no occupancy prediction and setback algorithm is applied.

	Efficiency gain (%)														$\sum$ **kWh**	
	OPT		MAT		MDMAT		PP		PPS		PH		REA		NO SETB.	
Weather	☼	☁	☼	☁	☼	☁	☼	☁	☼	☁	☼	☁	☼	☁	☼	☁
F-U$_{\text{low}}$ (well insulated flat)																
Very low	5	4	4	2	4	2	4	2	4	2	4	3	13	14	51	55
Freezing	8	6	6	5	6	5	6	5	6	5	6	5	10	12	38	44
Low	10	9	8	8	8	8	8	8	8	8	8	8	10	12	27	32
Moderate	11	12	10	11	10	11	10	11	10	11	10	11	11	13	17	20
H-U$_{\text{low}}$ (well insulated house)																
Very low	4	3	3	1	3	1	3	1	3	1	3	2	15	16	155	166
Freezing	6	5	4	4	4	3	5	3	4	3	5	4	10	12	119	134
Low	8	7	6	6	6	6	6	6	6	6	7	6	9	10	84	99
Moderate	9	10	8	8	8	8	8	8	8	8	8	8	9	10	53	65

the whole schedule. The participants are sorted in descending order of Π^{max} from left to right. The maximum value of Π^{max} is 95% while the minimum is 81%. The average of Π^{max} over all homes is 90%. This value is thus an upper bound for the average prediction accuracy achievable by any predictor. In Section 6.1 (see Figure 4) we observed that the median accuracy of the PP algorithm was 85%, which is just 5% below the upper bound of 90%. This indicates that a fairly simple schedule-based approach such as PP can in itself capture most of the predictability intrinsic in typical occupancy schedules. Furthermore, this result indicates that the use of more sophisticated schedule-based algorithms will provide a maximum improvement in accuracy of about 5% only. Note, however, that the use of context-aware algorithms may push the achievable accuracy above the 90% limit, as with such algorithms information other than past occupancy schedules is used to compute predictions.

6.3. Efficiency gain and comfort loss

Having discussed the accuracy of schedule-based occupancy prediction algorithms, we now investigate the performance of a predictive heating controller that uses the MAT, MDMAT, PP(S) and PH algorithms. For reference purposes we have also included OPT, which uses an oracle to provide a perfect prediction of household occupancy. To measure the energy consumption of the heating system, we built a simulation system [35] based on the ISO 5R1C model introduced in Section 5.3. We assumed the heating controller

Table 5: Same as Table 4, but with **high U-values (poor insulation)**.

	Efficiency gain (%)														$\sum$**kWh**	
	OPT		MAT		MDMAT		PP		PPS		PH		REA		NO SETB.	
Weather	☼	☁	☼	☁	☼	☁	☼	☁	☼	☁	☼	☁	☼	☁	☼	☁
F-U$_{\mathbf{high}}$ (poorly insulated flat)																
Very low	10	9	9	9	9	9	9	9	9	9	9	9	11	11	123	124
Freezing	14	13	14	13	14	13	14	13	14	13	14	13	14	14	95	100
Low	16	17	16	17	16	17	16	17	16	17	16	17	16	17	69	74
Moderate	18	19	18	19	18	19	18	19	18	19	18	19	18	19	45	48
H-U$_{\mathbf{high}}$ (poorly insulated house)																
Very low	7	6	6	6	6	5	6	5	6	5	6	5	12	12	328	332
Freezing	11	10	10	9	10	9	10	9	10	9	10	9	13	13	255	269
Low	14	14	13	13	13	13	13	13	13	13	13	13	14	14	186	200
Moderate	15	15	14	15	14	15	14	15	14	15	14	15	15	15	122	133

behaves as described in Section 5.4, irrespective of the algorithm used to predict occupancy. We simulated the response of the controller for the four building variants (F-U$_{\text{low}}$, F-U$_{\text{high}}$, H-U$_{\text{low}}$ and H-U$_{\text{high}}$) and eight weather scenarios introduced in Section 5.3, resulting in 32 different configurations.

We measured the performance of the controller for each algorithm in terms of *efficiency gain*. Let Q_{pred} be the heat injected by a predictive heating controller into the home and $Q_{no\_setback}$ the corresponding heat injected by a controller that maintains the temperature of the home constantly at Θ_{comf} throughout the day. The *efficiency gain* is then defined as $(Q_{no\_setback} - Q_{pred})/Q_{no\_setback}$. Defining and measuring thermal discomfort in an appropriate way is not easy. In 1970, Gupta proposed using "the ratio of the temperature-time curve area outside the specified comfort zone to that area of the comfort zone" as a "degree of discomfort" [40]. We used a discretised variant of that measure which yields absolute values per day. *Discomfort degree hours* as a measure of comfort loss are defined as the average sum of hourly differences between the actual indoor air temperature Θ_{air} and Θ_{comf} for all occupied intervals, formally $1/4(\Theta_{comf}\Gamma_{1..96} - \Theta_{air,1..96}) \cdot \Gamma_{1..96}$. Here, $\Gamma_{1..96}$ denotes the ground truth occupancy vector containing 1's for occupied intervals and 0's for unoccupied intervals. Thus, if $\Theta_{air} = 17\,°\text{C}$ upon the arrival of the occupants at 5 p.m. and the heating system requires 1 hour to heat up to $\Theta_{comf} = 20\,°\text{C}$ (e.g. $\Theta_{air,17:15} = 18\,°\text{C}$, $\Theta_{air,17:30} = 19\,°\text{C}$, $\Theta_{air,17:45} = 19.5\,°\text{C}$ and $\Theta_{air,18.00} = 20\,°\text{C}$), then the discomfort degree hours for this day will be 0.75.

Tables 4, 5 and 6 present the results for all 32 configurations. They show the efficiency gain and discomfort degree hours for all analysed algorithms. It is worth noting that the absolute values for the metrics reported clearly depend on the specific model, data and parameters used in this study. The generalisability of these results is discussed at the end of this section.

A predictive heating system is able to achieve the highest efficiency gain in poorly insulated buildings. The potential efficiency gain as determined by OPT is 9% to 19% for the flat F-$\mathrm{U_{high}}$ and 6% to 15% for the house H-$\mathrm{U_{high}}$ (Table 5). For well insulated buildings (low U-values), the efficiency gain under optimal prediction is reduced to a value of 4% to 12% for the flat and 3% to 10% for the house (Table 4). Higher U-values mean that the buildings' indoor temperature drops more quickly. At the same time, the prediction horizon I^* is reduced due to a higher design heat load $\Phi_{H,max}$ (cf. Table 2 in Section 5.3) and the efficiency gain increases. This happens regardless of the prediction algorithm. As I^* approaches zero, the predictive controller's behaviour approaches that of the *reactive* controller. The reactive controller (REA), which does not predict or preheat (i.e. only heats the building when it is occupied), has the highest efficiency gain for all scenarios – 9% to 19%. However, this also comes at the expense of the highest average discomfort degree hours (i.e. a large loss of comfort). For this reason, REA is clearly not a practical alternative in particular on very cold and freezing days. As the difference between Θ_{comf} and the outside temperature Θ_e becomes smaller, OPT and the reactive strategy converge since it takes less time to heat up the building.

The inability of the analysed algorithms to perfectly predict occupancy has the largest impact on well-insulated buildings (i.e. F-$\mathrm{U_{low}}$ and H-$\mathrm{U_{low}}$) when solar gains and outdoor temperatures are low (i.e. very low temperature, cloudy scenario). In this case, when compared to the perfect prediction OPT, the algorithms typically do not achieve much more than 50% of possible savings. This is due to the fact that this scenario requires prediction over a longer prediction horizon I^*.

As Table 6 shows, none of the prediction algorithms (OPT, MAT, MD-MAT, PP, PPS and PH) produced significant comfort loss in terms of discomfort degree hours. Apart from the very low temperature scenario, where the temperature sometimes dropped below $-6\,^{\circ}\mathrm{C}$ (the design temperature[11] used for the calibration of $\Phi_{H,max}$), the average discomfort degree hours are less than one for all scenarios and prediction algorithms. Moreover, even for

[11] The design temperature is defined as the minimum two-day average temperature that was reached at least 10 times in the last 20 years [33].

Table 6: Average discomfort degree hours per day (as a measure for comfort loss) for all experiments. ☼ and ☁ denote *clear* and *cloudy* scenarios.

Discomfort degree hours per day

Weather	OPT, (MD)MAT, PP(S) ☼	OPT, (MD)MAT, PP(S) ☁	PH ☼	PH ☁	REA ☼	REA ☁	OPT, (MD)MAT, PP(S), PH ☼	OPT, (MD)MAT, PP(S), PH ☁	REA ☼	REA ☁
	F-U$_{\mathbf{low}}$ (well insulated flat)						**F-U$_{\mathbf{high}}$** (poorly insulated flat)			
Very low	0				17	22	0		1	1
Freezing	0				2	7	0			
Low	0				0	1	0			
Moderate	0						0			
	H-U$_{\mathbf{low}}$ (well insulated house)						**H-U$_{\mathbf{high}}$** (poorly insulated house)			
Very low	0		1	1	28	35	0		8	8
Freezing	0				5	12	0		1	2
Low	0				0	2	0			
Moderate	0						0			

the reactive controller (REA) there was no significant comfort loss for the low and moderate temperature scenarios. We will discuss possible reasons for this behaviour in Section 7.1.

One should realise that to achieve significant savings, the response of the "standard" heating controller (cf. Section 5.4) to the algorithms' predictions may be too conservative. Especially for lower temperatures and well-insulated buildings, the additional efficiency gain of the reactive over a predictive controller is substantial. This indicates that with some (negligible or at least acceptable) comfort loss or simply by defining a reasonable temperature comfort bound around the setpoint, higher savings should be obtainable by more "courageous" predictive controllers. A modified controller, which not only optimises for zero miss-time (e.g. $\Theta_{air} = \Theta_{comf} \pm \Delta$) upon the arrival of the occupants) but also assigns a cost to discomfort degree hours and balances this with the actual heating costs, may obtain a higher efficiency gain while incurring only minimal additional discomfort degree hours (and thus comfort loss) per day. This approach has already been suggested by Mozer et al. in [1]. We leave the investigation of controllers that trade comfort loss for efficiency gain to future work.

6.3.1. Annualised savings

So far, the results in this section have shown the efficiency gain for selected weather scenarios. The annual efficiency gain is determined by the number of occurrences of each of these scenarios per year. Thus, they can be computed by weighting the efficiency gain of the weather scenarios by their empirical probability as derived from historical weather data. Table 7 shows

Table 7: ISO 13790 annual efficiency gains.

	Efficiency gain (%)													
	OPT		MAT		MDMAT		PP		PPS		PH		REA	
Building	☼	☁	☼	☁	☼	☁	☼	☁	☼	☁	☼	☁	☼	☁
H-U$_{\mathbf{low}}$	8	8	7 / 6										9	11
F-U$_{\mathbf{low}}$	10	10	8	9	8	8	9	8	9	8	9	9	11	12
H-U$_{\mathbf{high}}$	13	14	13										14	15
F-U$_{\mathbf{high}}$	16	17	16 / 17										16	17

Table 8: Average outside temperatures for selected cities and simulated efficiency gain for January to March (F-U$_{low}$).

	Average temperature (°C)			**Efficiency gain *OPT* (%)**		
City	*Jan*	*Feb*	*Mar*	*Jan*	*Feb*	*Mar*
Moscow	-8.0	-7.0	-2.0	6	7	9
Toronto	-5.8	-5.6	-0.4	7	7	9
Beijing	-4.0	-1.0	6.0	5	6	11
Stockholm	-2.8	-3	0.1	7	7	9
New York	0.5	1.8	5.7	8	8	11
Lausanne	1.3	2.8	5.5	6	7	9
Brussels	3.3	3.7	6.8	8	8	10
London	4.3	4.5	6.9	8	8	10
Seattle	5.6	6.3	8.1	10	11	12

the annualised efficiency gain for all four building scenarios. The weightings for the weather scenarios were determined using the historical weather distribution of the 20 years from 1994 to 2014. The table shows that all the prediction algorithms (MAT, MDMAT, PP(S) and PH) achieved the same annual efficiency gain, close to OPT, ranging from 6% (well insulated house) to 17% (poorly insulated flat).

6.3.2. Impact of climate conditions

Different climate zones may offer varying potential for energy savings. To indicate how well our findings for Lausanne can be generalised to other locations, Table 8 shows the efficiency gain achievable by OPT for the average weather conditions from January to March for selected cities[12]. For these simulations, a simplified model of F-U$_{low}$ with no solar gains and constant outside temperatures was applied, and the outside temperature equaled the average temperature for the month in question. Further details can be obtained from [35].

[12]Temperature data obtained from `wikipedia.org`, if available, otherwise from `weatherbase.com`.

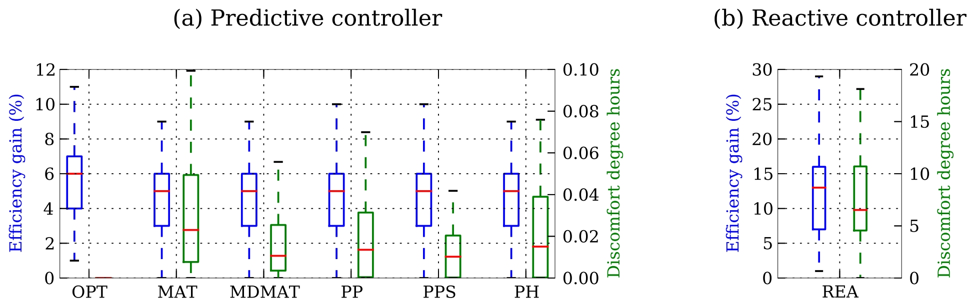

Figure 7: Efficiency gain and comfort loss measured in discomfort degree hours per day according to the ISO 5R1C model (F-U_{low}, freezing temperature, cloudy).

Table 8 shows an increase in the efficiency gain of between 5% (Beijing) and 10% (Seattle) in January to a range between 9% (Toronto) and 12% (Seattle) in March. This pegs the efficiency gain closely to the annualised figures obtained for the more detailed Lausanne simulation shown in Table 7. Cities with larger differences in the average outside temperature (e.g. Beijing has a difference of 10 °C between January and March), generally also have a larger variance in efficiency gain. This is due to the fact that the heating system is designed for the lowest temperatures. As the temperatures increase, the additional power of the heating system can be used to heat up the building more quickly.

6.3.3. Impact of the occupancy schedules

As one might expect, the potential for energy savings is highly correlated to a home's occupancy schedule. We analysed the impact of occupancy in the freezing temperature, cloudy sky scenario weather scenario. Figure 7 shows that for the well insulated flat F-U_{low}, efficiency gain and discomfort degree hours vary considerably between the participants. The bar plot shows the median, quartiles and extreme values of metrics for each algorithm (outliers have been removed). The left side of the figure shows the results for the predictive controller in conjunction with the assessed prediction algorithms. The right side shows the results for the reactive controller for comparison. As noted previously, the discomfort degree hours induced by the prediction algorithms are negligible. Overall, there are no significant differences between the algorithms and the distribution of their efficiency gain across the participants.

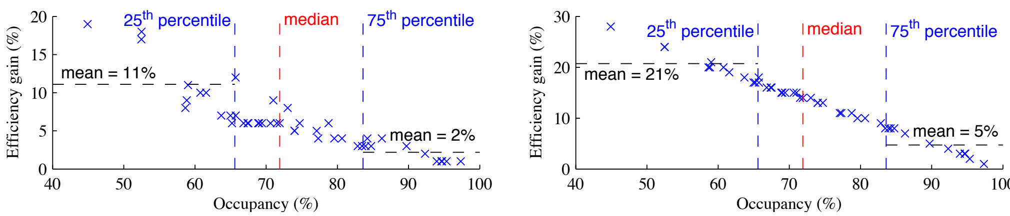

(a) Low U-values (well insulated)

(b) High U-values (poorly insulated)

Figure 8: Efficiency gain / occupancy correlation: Freezing temperature, cloudy.

Figure 8 shows the correlation between average occupancy and the efficiency gain that may be obtained by OPT for all 45 participants. Figures 8(a) and 8(b) contrast this relationship between F-U_{low} (good insulation) and F-U_{high} (poor insulation). The figures show that the quarter of homes that are least-occupied (25^{th} percentile) outperformed the most-occupied homes (75^{th} percentile) by a factor of 4-5. Low occupancy houses are clearly much better suited for installing smart heating systems than those with high occupancy.

The figures also show that for the 25% of homes with the lowest occupancy, the efficiency gain almost doubled from 11% to 21% from the well insulated to the poorly insulated flat. Not surprisingly, one can thus conclude that smart heating systems yield the highest benefits in poorly insulated buildings.

Figure 8(b) shows an almost linear relationship between occupancy and efficiency gain. This relationship is less pronounced in Figure 8(a). Here, the efficiency gain for the quarter of participants between the 25% quantile and the median is almost constant. As OPT's prediction is perfect, the reason for this effect lies in the structure of the occupancy schedules in conjunction with the increased prediction horizon due to the better insulation. The more arrival and departure events a schedule contains, the more difficult it is for the heating system to lower the temperature to a setback temperature.

7. Modelling limitations

Due to their novel nature, performance data from smart heating installations in domestic buildings is still sparse. However, to make substantiated claims regarding the impact of different variables such as the building's occupancy and insulation on the efficiency gain and comfort loss of a predictive heating system, one must analyse each variable *ceteris paribus*. Thus, for the time being, in order to analyse the specific impact of different variables, one must resort to simulations. Simulation and modelling naturally involve a trade-off between model complexity and simulation accuracy. In the following, we will briefly discuss some of the shortcomings of the ISO 5R1C model used in this report and analyse our choice of baseline strategy for computing efficiency gain.

7.1. Building model

To simulate the heating system, we used the 5R1C model from the ISO 13790 standard [29]. In this model, the heat source is connected via the node for the indoor air temperature. As such, even though it has been widely adopted for building design in Europe [41, 42], the ISO 5R1C model more closely resembles a forced-air heating system common in the US, rather than the hydronic systems more typically encountered in Europe. A forced-air heating system typically reduces the preheat time and lowers the penalty for false predictions, thereby resulting in the low comfort loss exhibited by the simulation results (cf. Table 6). From the variations between different insulation levels (cf. Figure 8), we have already seen that shorter preheat times induced by more powerful heating systems result in an almost reactive strategy and thus in higher energy savings. As such, our evaluation hints at an upper bound on the savings that can be achieved using predictive heating systems and may lead to an underestimation of comfort loss.

7.2. Baseline metrics

We employed an *always-on* strategy as the baseline for evaluating the predictive controller and the occupancy prediction algorithms. In practice, households often use a (static) *night-time setback*. Allowing the temperature to drop during the night by 4 °C to 6 °C has been shown to result in savings between 4% and 7% [43, 44]. A baseline strategy using a night-time setback thus lowers the overall energy consumption, thereby – assuming the predictive setback generally occurs during the day – slightly increasing the efficiency gain of the predictive controller. Using a night-time setback strategy as the baseline, however, necessitates a clear separation between the efficiency gain achieved by this setback and the predictive strategy.

Substituting the ISO 5R1C model with a more suitable and possibly more detailed building model, and also considering night-time setback, could be a task for future work. While an even more realistic simulation model would increase the confidence in the simulation results, we do not expect this would significantly affect the outcome.

8. Conclusions and summary of results

The insights gained through our simulation-based performance analysis of occupancy-based approaches for smart heating control, based on real-world weather data and established building standards, can be summarised as follows:

- Among the considered algorithms, the *Presence Probabilities* (PP, PPS) approach by Krumm and Brush [18] provides for the best overall performance in terms of prediction accuracy for the dataset considered in this paper. The approaches suggested by Lu et al. [2] and Scott et al. [3] (MAT, MDMAT, PH) perform slightly worse, albeit not by a large margin.

- The prediction accuracy of existing schedule-based algorithms is close to the achievable *theoretical upper limit*; this limit is expressed by the predictability of the underlying occupancy schedules. Further performance improvements can thus only be achieved by context-aware approaches that consider additional input information rather than occupancy schedules only.

- Actual *comfort loss* in terms of *discomfort degree hours* is lower than the values implied by the accuracy of the prediction algorithm. A prediction accuracy of around 80% does not necessarily result in an uncomfortable thermal environment for 20% of the time. This is mainly due to the reactive nature of the heating scenario (e.g. heating is not turned off prematurely based on a predicted state if the occupants are still present). Moreover, the comfort loss is bounded by the time it takes to heat from the current temperature to the comfort temperature.

- The *efficiency gain* achievable by occupancy prediction is dependent on the structure of the building, its occupancy and the weather conditions. Annual savings range from 6% to 17% depending on the type of building (cf. Table 7). Savings are almost doubled for poorly insulated buildings. The 25% of households with the lowest occupancy have a

4-5 times higher potential for efficiency gains than the quarter of homes with the highest occupancy. Lower temperatures and cloudy skies reduce efficiency gain and increase comfort loss as it takes longer to heat the building. Our data confirms similar results by [43] and [44] which showed energy savings of between 6% and 10% for cool and temperate climates using setback thermostats.

- The algorithms' inherent difficulty in correctly predicting the arrival time of the occupants imposes a penalty on the efficiency gain. To save more energy, *additional intelligence* could thus be incorporated into the controller. One example would be to forgo heating if only a short period of occupancy is predicted that would nevertheless result in significant energy expenditure to heat up the property. A mobile application or simple "override" button on the thermostat to enable the occupants to control the smart thermostat in a simple and easy manner could deal with exceptional cases and increase user acceptance.

Acknowledgements

This work has been partially supported by the Collaborative Research Center 1053 funded by the German Research Foundation and by the LOEWE Priority Program Cocoon funded by the LOEWE research initiative of the state of Hesse, Germany.

References

[1] M. C. Mozer, L. Vidmar, R. H. Dodier, The Neurothermostat: Predictive optimal control of residential heating systems, in: M. C. Mozer, M. I. Jordan, T. Petsche (Eds.), Advances in Neural Information Processing Systems, Vol. 9, The MIT Press, 1997, pp. 953–959.

[2] 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, in: Proc. SenSys'10, Zurich, Switzerland, ACM, 2010, pp. 211–224.

[3] J. Scott, A. J. Brush, J. Krumm, B. Meyers, M. Hazas, S. Hodges, N. Villar, Preheat: Controlling home heating using occupancy prediction, in: Proc. UbiComp'11, Beijing, PRC, ACM, 2011, pp. 281–290.

[4] European Environment Agency, Household energy consumption by end-use in the EU-27, www.eea.europa.eu/data-and-maps/figures/households-energy-consumption-by-end-uses-4 (2011).

[5] N. Kiukkonen, J. Blom, O. Dousse, D. Gatica-Perez, J. Laurila, Towards rich mobile phone datasets: Lausanne data collection campaign, in: Proc. ICPS'10, Berlin, Germany, ACM, 2010.

[6] T. A. Nguyen, M. Aiello, Energy intelligent buildings based on user activity: A survey, Energy and Buildings 56 (2013) 244–257.

[7] X. Guo, D. Tiller, G. Henze, C. Waters, The performance of occupancy-based lighting control systems: A review, Lighting Research and Technology 42 (4) (2010) 415–431.

[8] M. Gupta, S. S. Intille, K. Larson, Adding GPS-control to traditional thermostats: An exploration of potential energy savings and design challenges, in: Proc. Pervasive'09, Nara, Japan, Springer, 2009, pp. 1–18.

[9] Y. Agarwal, B. Balaji, S. Dutta, R. Gupta, T. Weng, Duty-cycling buildings aggressively: The next frontier in HVAC control, in: Proc. IPSN'11, Chicago, IL, USA, IEEE, 2011, pp. 246–257.

[10] F. Oldewurtel, A. Parisio, C. Jones, M. Morari, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, K. Wirth, Energy efficient building climate control using stochastic model predictive control and weather predictions, in: Proc. ACC'10, Baltimore, MD, USA, IEEE, 2010, pp. 5100–5105.

[11] C. Ellis, J. Scott, M. Hazas, J. Krumm, Earlyoff: Using house cooling rates to save energy, in: Proc. BuildSys'12, Toronto, Canada, ACM, 2012, pp. 39–41.

[12] V. Erickson, M. Carreira-Perpiñán, A. Cerpa, OBSERVE: Occupancy-based system for efficient reduction of HVAC energy, in: Proc. IPSN'11, Chicago, IL, USA, IEEE, 2011, pp. 258–269.

[13] M. Fountain, G. Brager, E. Arens, F. Bauman, C. Benton, Comfort control for short-term occupancy, Energy and Buildings 21 (1994) 1–13.

[14] ANSI/ASHRAE standard 55-2010 – thermal environmental conditions for human occupancy, www.techstreet.com/standards/ashrae/55_2010?product_id=1741646 (2010).

[15] W. Kleiminger, C. Beckel, T. Staake, S. Santini, Occupancy detection from electricity consumption data, in: Proc. BuildSys'13, ACM, Rome, Italy, 2013.

[16] D. Chen, S. Barker, A. Subbaswamy, D. Irwin, P. Shenoy, Non-intrusive occupancy monitoring using smart meters, in: Proc. BuildSys'13, ACM, Rome, Italy, 2013.

[17] T. W. Hnat, V. Srinivasan, J. Lu, T. I. Sookoor, R. Dawson, J. Stankovic, K. Whitehouse, The hitchhiker's guide to successful residential sensing deployments, in: Proc. SenSys'11, Seattle, WA, USA, ACM, 2011, pp. 232–245.

[18] J. Krumm, A. J. Brush, Learning time-based presence probabilities, in: Proc. Pervasive'11, San Francisco, CA, USA, IEEE, 2011, pp. 79–96.

[19] R. Dodier, G. Henze, D. Tiller, X. Guo, Building occupancy detection through sensor belief networks, Energy and Buildings 38 (9) (2006) 1033–1043.

[20] S. Tominaga, M. Shimosaka, R. Fukui, T. Sato, A unified framework for modeling and predicting going-out behavior, in: Proc. Pervasive'12, Newcastle, UK, Springer, 2012, pp. 73–90.

[21] S. Scellato, M. Musolesi, C. Mascolo, V. Latora, A. T. Campbell, NextPlace: A spatio-temporal prediction framework for pervasive systems, in: Proc. Pervasive'11, San Francisco, CA, USA, IEEE, 2011, pp. 152–169.

[22] K. Farrahi, D. Gatica-Perez, Discovering routines from large-scale human locations using probabilistic topic models, ACM Trans. Intell. Syst. Technol. 2 (1) (2011) 3:1–3:27.

[23] F. Vazquez, W. Kastner, Clustering methods for occupancy prediction in smart home control, in: Proc. ISIE'11, Gdansk, Poland, IEEE, 2011, pp. 1321–1328.

[24] J. Krumm, E. Horvitz, Predestination: Inferring destinations from partial trajectories, in: Proc. UbiComp'06, Orange County, CA, USA, Springer, 2006, pp. 243–260.

[25] C. Koehler, B. D. Ziebart, J. Mankoff, A. K. Dey, TherML: Occupancy prediction for thermostat control, in: Proc. UbiComp'13, Zurich, Switzerland, ACM, 2013, pp. 103–112.

[26] 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.

[27] Y. Ye, Y. Zheng, Y. Chen, J. Feng, X. Xie, Mining individual life pattern based on location history, in: Proc. MDM'09, Washington, DC, USA, IEEE, 2009, pp. 1–10.

[28] W. Kleiminger, C. Beckel, A. Dey, S. Santini, Using unlabeled Wi-Fi scan data to discover occupancy patterns of private households, in: Proc. SenSys'13, Rome, Italy, ACM, 2013.

[29] ISO, Energy performance of buildings - calculation of energy use for space heating and cooling, ISO 13790-1:2008, International Organization for Standardization, Geneva, Switzerland (2008).

[30] C. L. Beuken, Wärmeverluste bei periodisch betriebenen elektrischen Öfen: Eine neue Methode zur Vorausbestimmung nicht-stationärer Wärmeströmungen, Ph.D. thesis, Sächsische Bergakademie Freiberg (1936).

[31] E. Mathews, P. Richards, C. Lombard, A first-order thermal model for building design, Energy and Buildings 21 (2) (1994) 133–145.

[32] Verordnung über energiesparenden Wärmeschutz und energiesparende Anlagentechnik bei Gebäuden in der Fassung der Bekanntmachung vom 24. Juli 2007 (BGBl I S. 1519), zuletzt geändert am 18. November 2013 (BGBl I S. 3951–3990).

[33] DIN, Heating systems in buildings - method for calculation of the design heat load, DIN EN 12831-03:2008, Deutsches Institut für Normung, Berlin, Germany (2008).

[34] N. Helbig, Application of the radiosity approach to the radiation balance in complex terrain, Ph.D. thesis, University of Zurich, Zurich, Switzerland (February 2009).

[35] W. Kleiminger, F. Mattern, S. Santini, Simulating the energy savings potential in domestic heating scenarios in Switzerland, Tech. rep., ETH Zurich, Department of Computer Science, `http://www.vs.inf.ethz.ch/publ/papers/wilhelmk_tr_2014.pdf` (August 2014).

[36] P. Baumann, W. Kleiminger, S. Santini, The influence of temporal and spatial features on the performance of next-place prediction algorithms, in: Proc. UbiComp'13, Zurich, Switzerland, ACM, 2013, pp. 449–458.

[37] T. A. Reddy, Applied Data Analysis and Modeling for Energy Engineers and Scientists, Springer, 2011.

[38] C. Song, Z. Qu, N. Blumm, A.-L. Barabási, Limits of predictability in human mobility, Science 327 (5968) (2010) 1018–1021.

[39] M. Lin, W.-J. Hsu, Z. Q. Lee, Predictability of individuals' mobility with high-resolution positioning data, in: Proc. UbiComp'12, Pittsburgh, PA, USA, ACM, 2012, pp. 381–390.

[40] C. Gupta, A systematic approach to optimum thermal design, Building Science 5 (3) (1970) 165–173.

[41] J. Jokisalo, J. Kurnitski, Performance of EN ISO 13790 utilisation factor heat demand calculation method in a cold climate, Energy and Buildings 39 (2) (2007) 236–247.

[42] J. A. Orosa, A. C. Oliveira, Implementation of a method in EN ISO 13790 for calculating the utilisation factor taking into account different permeability levels of internal coverings, Energy and Buildings 42 (5) (2010) 598–604.

[43] J. Ingersoll, J. Huang, Heating energy use management in residential buildings by temperature control, Energy and Buildings 8 (1) (1985) 27–35.

[44] M. M. Manning, M. C. Swinton, F. Szadkowski, J. Gusdorf, K. Ruest, The effects of thermostat setting on seasonal energy consumption at the CCHT twin house facility, ASHRAE Trans. 113 (1) (2007) 1–12.