Energy and Buildings 85 (2014) 493–505

Contents lists available at ScienceDirect

Energy and Buildings

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

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

ARTICLE INFO

Article history:
Received 24 August 2013
Received in revised form 29 August 2014
Accepted 25 September 2014
Available online 7 October 2014

Keywords:
Occupancy detection
Occupancy prediction
Smart heating
Energy management
Smart home
Energy efficiency
Thermostat strategy
Heating setback

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* – that is, 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 trade-off between achievable savings (typically 6–17% on average) and the risk of comfort loss for household residents.

1. Introduction

The ability to detect whether or not a house is occupied – that is, 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 – that is, 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 following contributions: (1) A classification and review of state-of-the-art approaches that predict home occupancy. We

* Corresponding author. Tel.: +41 0446328463.
E-mail addresses: wilhelmk@ethz.ch (W. Kleiminger), mattern@ethz.ch (F. Mattern), santinis@wsn.tu-darmstadt.de (S. Santini).

[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.

http://dx.doi.org/10.1016/j.enbuild.2014.09.046
0378-7788/

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 [3,9–13]. The first few commercial products – such as the NEST learning thermostat, tado° and EcoBee's Smart-Si – have recently started to appear.[2]

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, that is, 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 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 h). 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

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

[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, for example, 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.