Occupancy Detection from Electricity Consumption Data

Wilhelm Kleiminger,
Christian Beckel
Institute for Pervasive
Computing
ETH Zurich, Switzerland
{wilhelmk,beckelc}@ethz.ch

Thorsten Staake
Energy Efficient Systems
University of Bamberg,
Germany
thorsten.staake@uni-
bamberg.de

Silvia Santini
WSN Lab
TU Darmstadt, Germany
santinis@wsn.tu-
darmstadt.de

ABSTRACT

Detecting when a household is occupied by its residents is fundamental to enable a number of home automation applications. Current systems for occupancy detection usually require the installation of dedicated sensors, like passive infrared sensors, magnetic reed switches or cameras. In this paper, we investigate the suitability of digital electricity meters – which are already available in millions of households worldwide – to be used as occupancy sensors. To this end, we have collected fine-grained electricity consumption data along with ground-truth occupancy information for 5 households during a period of about 8 months. Our results show that using common classification methods it is possible to achieve occupancy detection accuracies of more than 80%.

Categories and Subject Descriptors

H.1.2 [**Models and Principles**]: User/Machine Systems–Human Information Processing; H.4 [**Information Systems Applications**]: Miscellaneous

General Terms

Design, Experimentation, Measurement, Human Factors

Keywords

Smart Meter, Electricity, Occupancy, Context-Aware

1. INTRODUCTION

Occupancy detection is a main building block of commercial and residential building automation systems. For instance, systems that regulate heating, ventilation and cooling (HVAC) rely on estimated occupancy information to control buildings' temperature and air-flow [14, 21]. Similarly, many lighting systems rely on the detection of presence (or absence) of people in doorways or meeting rooms to switch lights on (or off) [8]. Dickerson et al. showed that the detection of changes in occupancy patterns can even help revealing clinical diseases such as depression [7].

Despite the large number of potential application scenarios the, detection of buildings occupancy is still a cumbersome, error-prone and expensive process [18]. Occupancy is typically sensed using dedicated devices such as passive infrared (PIR) sensors, magnetic reed switches or cameras [18]. Such sensors need to be purchased, installed, calibrated, powered and maintained. This poses a number of critical constraints, especially in domestic environments. First, the overall costs of the occupancy sensing infrastructure must be kept low. This often implies that only few, cheap and possibly imprecise sensors are available. Also, battery-powered sensors are often used to avoid the need of power cables to be deployed. The availability and reliability of the sensors might thus be affected by depleted batteries waiting to be replaced. Further, in a domestic setting one of the (often technically unexperienced) residents takes over the role of the "building administrator" who installs and maintains the system. Faulty installations and lack of maintenance are a frequent consequence. Taken together these constraints might cause occupancy detection systems to become unreliable and induce faulty behaviours in the home automation systems relying on them. This can in turn cause inconvenience for the residents and hamper their acceptance of the systems. At the same time, many sensor devices are already available in typical households and can be used to perform occupancy detection in an opportunistic manner. These devices can contribute to improve the overall reliability of the system or reduce its cost. For instance, smartphones can be used to detect the presence of one or more of the residents within the household, as we have also discussed in our previous work [12].

In this paper, we investigate the possibility of using digital electricity meters as part of an opportunistic occupancy sensing infrastructure in domestic settings. To this end, we have run an extensive experiment to collect electricity consumption data of 5 households over 8 months. Household residents have recorded ground truth occupancy data using a custom Android application that we have developed for this study. For later analysis as well as indirect validation of our measurements we have also collected data from PIR sensors and the device-level electricity consumption of selected electrical appliances. We have analysed this data set using standard classification techniques to evaluate the occupancy detection accuracy achievable by using only electricity meters as occupancy sensors. Our results show that occupancy classification from electricity consumption data is feasible. In particular, an average detection accuracy of over 80% can be obtained in most settings. To the best of our

knowledge, this is the first study that provides a quantitative analysis of the possibility to detect occupancy from electricity consumption data. This is also due to the fact that large data sets including both electricity consumption data and ground-truth occupancy information have been missing.

We discuss related work in section 2, the setup of our data collection and the data cleaning we have performed on the raw data in section 3 and our results in section 4.

2. RELATED WORK

We see our work at the intersection of two main areas: (1) sensing deployments to detect occupancy and improve energy efficiency in residential and commercial buildings; (2) analysis of electricity consumption data to observe and influence users' electricity consumption behaviour.

2.1 Occupancy sensing to increase energy efficiency

Many authors report about occupancy sensing systems to save energy [18]. In residential environments, most of these approaches aim at controlling the HVAC system more efficiently, as heating and cooling account for most of the energy expenditure in an average household. To optimise heating or cooling of the building, the authors typically utilise various sensors to determine the occupancy state of the building or of individual rooms. For instance, Lu et al. instrumented households with passive infrared (PIR) sensors and reed switches on entrance doors to detect when household occupants are at home and when not [14]. They use this data as ground truth information to evaluate the performance of an occupancy prediction algorithm, which is in turn used to drive a smart heating control system. The occupancy state of the household is computed every five minutes using a Markov model. The model takes as input the hour of the day and the number of firings of the PIR sensors and of the reed switches. Similarly, Scott et al. developed *Pre-Heat* [21], a system that senses and predicts occupancy to efficiently control the heating system. In their deployment, the authors used active RFID tags in three US households as well as motion sensors in two UK homes to monitor per room occupancy. While in our deployment we also use PIR sensors to monitor if people are at home, the focus of our work is on exploring how the electricity consumption curve can be used to estimate the occupancy state of a household.

In commercial buildings, detailed knowledge about the current (and future) state of the building can be used to optimise HVAC, lighting, and the use of individual appliances. To this end, researchers explored many types of sensing systems to obtain the occupancy state of the building and the activities of its occupants [8, 18, 22]. Dodier et al., for instance, equipped two offices with a sensor at the telephone handset and three PIR sensors each. By feeding past and current sensor readings into belief networks (i.e. a class of graphical probability models), the authors estimate the number of persons in the offices as well as their location. Similarly, Milenkovic et al. equipped three offices with PIR sensors and plug-in power meters, which measured the energy consumption of the computer screens [16]. Using layered hidden Markov models (LHMMs), the authors estimated the number of persons in the office as well as their current activity (e.g. desk work). Monitoring activities of occupants at such a fine-grained level requires a large number of sensors. Our work focuses on residential settings, however, where deploying sensors is often infeasible for cost reasons and due to deployment issues. Thus, in our work, we focus on estimating occupancy state of the household solely by analysing its electricity consumption.

2.2 Analysis of electricity consumption data

Measuring and analysing the electricity consumption of households has been addressed by many researchers in the past for different applications. Through the analysis of coarse-grained consumption data (e.g. in the order of 1 measurement per 15 or 30 minutes), for instance, it has been shown that energy providers can identify usage patterns in the electricity consumption data to predict future electricity consumption [6] or model daily routines to improve a providers's supply management [1]. Other researchers have proposed approaches that can cluster hundreds of households into groups of consumers according to their load profile [20, 24] or estimate socio-economic characteristics of a household [3]. Using these techniques, it is possible for an energy provider to identify the households that are unoccupied during the day. To those particular households, energy providers could offer a special tariff or provide them with a system to automatically switch off their heating system when they are not at home. In [17], Molina et al. suggest that such occupancy patterns can be detected from the electricity consumption data. However, the authors made the observation by visually inspecting the load curves and do not perform a data analysis based on the real occupancy state of the households.

By analysing the fine-grained electricity consumption of a household (e.g. in the order of 1 Hz), many researchers have tackled the problem of inferring which appliances are running at what time. Zoha et al. [26] and Zeifman et al. [25] provide two good reviews of related work in "non-intrusive load monitoring" (NILM). One of the first NILM approaches has been proposed by George Hart in 1992 [10]. Hart's method identifies characteristic step changes in the electricity consumption. By comparing these step changes monitored in the electricity consumption with a previously recorded signature database, Hart claims to detect when appliances are being switched on or off. More recent approaches, such as the one from Kim et al. [11], pursue unsupervised disaggregation. These unsupervised approaches do not require a training phase, but require only an explicit labelling of those appliances detected in the load curve.

As some devices are (typically) only used when the occupants are at home, NILM would implicitly provide occupancy detection as required by many energy efficiency applications. However, if the electricity consumption is measured at a granularity of at most 1 Hz, only a few appliances (e.g. the refrigerator, or the washing machine) can be detected reliably from the data [4]. Increasing the accuracy of detecting individual appliances in the electricity consumption data requires a more characteristic signature of each appliance, which can be achieved by increasing the measurement granularity. As Gupta et al. show, this approach can identify and classify most consumer electronic and fluorescent lighting devices correctly with a mean accuracy of more than 93% [9]. However, while it requires special hardware to measure the electricity consumption at multiple kilohertz, our approach relies only on 1 Hz consumption data, which can be obtained from an off-the-shelf electricity meter.

To evaluate the performance of a NILM algorithm, researchers typically rely on a publicly available data set. A

data set often used is the *Reference Energy Disaggregation Dataset (REDD)*, which was collected by Kolter et at. [13]. It contains electricity consumption data measured in 5 homes in the US along with plug-level consumption measurements of individual circuits or appliances. More recently, Barker et al. published the *UMASS Smart* Home Data Set* [2], which contains very detailed submeter measurements as the authors deployed 21 − 26 circuit meters into three homes each. However, since neither of these two data sets contains ground truth information about the occupancy patterns of the inhabitants, we collected our own data set in 5 Swiss households over the course of 8 months.

In contrast to all approaches described above, our work estimates the occupancy pattern of a household solely by analysing its electricity consumption. To this end, we performed an extensive data collection, because there is – to the best of our knowledge – no data set available that contains both electricity consumption and ground truth occupancy information of households.

3. DATA COLLECTION

To estimate the occupancy state of a household based on its electricity consumption, we performed an extensive data collection in collaboration with a utility company in Switzerland. We collected a multi-modal data set in 5 households over the course of 8 months. In addition to the electricity consumption of a household the data set contains sensor information collected from passive infrared (PIR) sensors and smart power outlets. All households also recorded ground truth occupancy data through a tablet computer. This section describes the selection of households, our measurement infrastructure, an overview over the data set, and the steps that were necessary for data cleaning.

3.1 Selection of households

For the data collection we chose the participating households among employees of a utility company in Switzerland. Prospective participants were required to fill in a questionnaire. The questionnaire contained 12 questions targeting the number, age and occupation of the occupants, type of property, number of entry doors, typical occupancy, type of heating, pet ownership as well as the level of affinity for technology of the respondent. The affinity for technology was requested through a 7-point Likert scale (1: low, 4: medium, 7: very high). The purpose of the questionnaire was to ensure households have a reasonable size (i.e. 1-4 occupants) and participants are well-disposed to technical equipment. Also, we avoided to include households in which occupants used more than one entrance, because we wanted each participant to log occupancy through a tablet computer that is located next to the main entrance. To each participant we handed a privacy statement that described in detail the data gathered and their ability to opt out at any time during data collection.

Table 1 shows an overview of the households which we ultimately selected to participate in the data collection. Three of the households consist of two occupants, while two of the households are occupied by four persons. Four out of the five respondents live in detached houses, only the occupants of household 2 live in a flat. All respondents except for one classified themselves as tech-savvy.

3.2 Measurement infrastructure

Table 1: Overview of the participants.

House-hold	No. of occupants	Type of property	Affinity for technology
1	2 adults, 2 children	House	7/7
2	2 adults	Flat	7/7
3	2 adults	House	7/7
4	2 adults, 2 children	House	4/7
5	2 adults	House	6/7
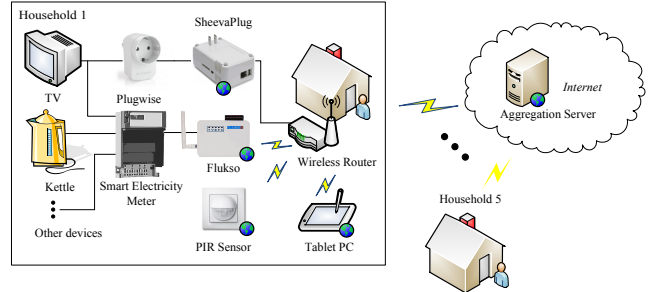
Figure 1: Schematic overview of whole deployment.

Figure 1 shows a schematic view of the components deployed during the data collection. We installed a smart electricity meter (model E750) from Landis+Gyr[1] in each of the participating households. The meters were connected behind the electricity meter and not used for billing purposes. The E750 offers an Ethernet interface to access measurements at 1 Hz granularity through the Smart Message Language (SML) protocol[2]. SML is a request-response protocol and allows the client to send a request specifying the variables to be read. To query the measurements from the electricity meter through SML, we wrote a C-based program called *Pylon*[3] and deployed it on a *Fluksometer*[4] into each household. The Fluksometer is a commercial device with an Ethernet and a Wi-Fi interface that runs the open source router software OpenWrt[5]. Pylon sent the latest measurements to a server once per second using the Flukso's Wi-Fi interface and the Internet connection of the households.

We also measured the electricity consumption of selected appliances (e.g. refrigerator, tumble dryer, router, kettle, washing machine) using smart power outlets (smart plugs) from Plugwise[6]. The smart plugs communicate via Zigbee[7] and automatically establish a mesh network among themselves. To acquire the smart plugs' consumption in real time we utilise the open source library *python-plugwise*[8]. This setup allowed us to query each smart plug at a frequency of 1 Hz. The queries are managed by a python script on a mini computer (*Sheeva Plug*[9]), which communicates the measurements to our server and also buffers the data to avoid data loss in case the network is unavailable.

In order to gather occupancy data, we deployed Roving

[1] www.landisgyr.com

[2] v1.03: www.emsycon.de/downloads/SML_081112_103.pdf

[3] www.github.com/wkleiminger/pylon

[4] www.flukso.net

[5] www.openwrt.org

[6] www.plugwise.com

[7] www.zigbee.org

[8] www.bitbucket.org/hadara/python-plugwise/wiki/Home

[9] www.globalscaletechnologies.com/t-sheevaplugs.aspx

Table 2: Total number of records collected between June 2012 and January 2013.

Sensor (# devices)	Description	# Records
Landis+Gyr E750 (6)	Smart Electricity Meters	106,337,104
Plugwise Sting (39)	Smart Power Outlets	566,407,440
Roving RN-134 (5)	PIR sensors	549,702
Galaxy Tab P7510 (5)	Occupancy ground truth	6,366

RN-134 low-power Wi-Fi modules with passive infrared sensors attached in 4 of the participating households. The RN-134 modules consume very little power when asleep and can sense while the radio is switched off [19]. When a request needs to be transmitted, the radio is switched on for a brief period of time. As soon as the request has been completed, the module goes back to sleep. This allows each module to run for approximately 3 months on two AA batteries.

To obtain ground truth data for our occupancy classification algorithms, we installed a Samsung Galaxy Tab P7510 tablet computer as an in-home display in each household. The tablet was configured to display an interface (UI) for users to record the occupancy state of the residence. For each occupant there is a toggle button that may be pressed to change the status from *present* to *absent* and vice versa. The tablet computer was installed near the main entrance and the occupants were instructed not to move it during the course of the data collection. In order to increase the visibility of the application and to remind the participants to record their occupancy, the tablet's display was automatically switched on whenever the passive infrared sensor sensed a movement. The visualisation of the current electricity consumption, 7-day historical consumption, aggregate consumption and a historical chart with smooth zooming on the in-home display provided further utility for the participants.

Figure 1 shows all the data collected in this deployment was transferred to an aggregation server located at ETH Zurich via HTTP. The raw measurements were stored in a database, while the consumption information was also held in a 3-level cache to facilitate smooth zooming on the tablet. The centralised collection helped us to quickly identify if one of the sensors was faulty. Table 2 shows the total number of measurements gathered from June 2012 to January 2013. Overall, we collected 106,337,104 values from the 5 smart meters, 566,407,440 values from the 39 smart plugs, 549,702 PIR sensor readings, and 6,366 occupancy events (i.e. when an occupant enters or leaves the home).

3.3 Data pre-processing

In this section we will discuss how we prepared the raw data for analysis and filtered erroneous ground truth data. The occupancy analysis focuses on a subset of the data collected. Due to the importance – and difficulty – to record reliable ground truth occupancy data, we instructed households to particularly pay attention to reliably specify their occupancy during two phases in *summer* (July to September) and *winter* (November to January). During these two collection phases every participant was instructed to click on a button bearing his or her name to indicate presence and absence.

Providing data at 1 Hz, the smart electricity meters produce 86,400 measurements per day. In order to be able to directly compare the electricity consumption to the other sensor data, we converted all other data to 86,400 element vectors as well. The smart plugs from Plugwise must be read sequentially [23]. Queries have a round trip time of 80–120 ms for each plug, depending on the network infrastructure. As there are 6–9 plugs per household it takes about one second to obtain the consumption data of each plug. However, problems that occur for one of the plugs (e.g. a slow reply or a timeout due to network interference) can lead to short time periods of $5 - 10$ seconds during which no data from any of the plugs can be obtained. Ultimately, the consumption measurements for each plug are re-sampled to 86,400 measurements a day (i.e. 1 Hz). For each day d, the occupancy states of a household h are captured by $O_{h,d}$. $O_{h,d}$ is a $86,400 \times N_p$ matrix containing the occupancy state for each member p of the household at every second of the day. The element (i,j) of this matrix is set to 1 if – according to the data entered using the tablet – the jth resident is at home at second i. The element is set to 0 otherwise. Following this notation, we compute the binary occupancy schedule $B_{h,d}$, a $86,400 \times 1$ vector by computing the bitwise OR among the rows of the matrix. The resulting vector contains 1s to indicate occupancy and 0s to indicate that none of the occupants are present. For the PIR sensors, the matrix contains a sequence of 1s for the next 30 seconds after a sensor event has been triggered.

In case of the electricity consumption data from the smart meters, we distinguish between two types of data loss. First, if measurements are missing for up to 10 seconds, the corresponding positions in the vector are filled with the last existing measurement (typically only few seconds are lost each day). Second, in case more than 10 consecutive seconds of data are lost – for example in (rare) cases where the Flukso crashed or was switched off – the values are set to -1. For the smart plugs, data loss is dealt with similarly. In this case, we chose 100 seconds instead of 10 seconds as a threshold. This is due to the fact that a data loss of 10 seconds is more common for the reasons described above.

From the thus processed data, we prepare our test and training sets for the occupancy analysis. Even though the participants noted their occupancy diligently, some mistakes still occurred (e.g. one or more occupants occasionally forgot to record their absence or presence). We have therefore manually removed days where:

1. all occupants have indicated "absence" but a firing of the PIR sensor indicated movement in the household,
2. a switch operated device (e.g. kettle, TV, oven) has been operated during the period of "absence",
3. no occupancy information was collected.

Table 3 shows the data gathered by the participants and used in the evaluation. The table shows, for each household, the number of days in both summer and winter phases after erroneous days have been removed from the data set. This results in an average of 52 days for the summer period and 38 days for the winter period.

Figure 2 shows a representative day of data collected for household 2. Figure 2a shows the total electricity consumption of the household, augmented with the binary occupancy state as indicated by the occupants on the tablet interface. The electrical load curve shows a small increase in the electricity consumption when the occupants wake up and prepare breakfast. As the occupants leave the household, the passive infrared sensor (PIR) near the doorway fires (see

Table 3: Number of days for each household used in the evaluation after data cleaning.

Household	Number of days	
	Summer	*Winter*
1	39	46
2	83	45
3	57	21
4	38	48
5	43	31
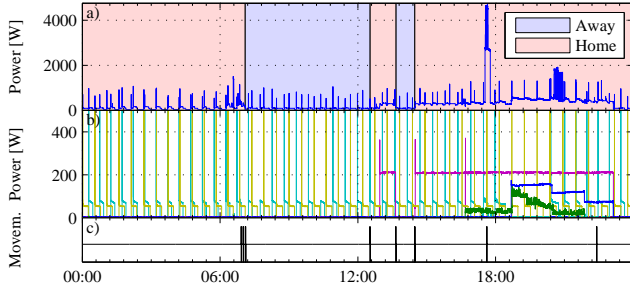
Figure 2: A typical day in household 2. (a) total electricity consumption, (b) Appliance-level consumption (c) Movement sensed by PIR sensor.

figure 2c). As the occupants return again, the PIR sensor fires again and the home entertainment is switched on (see figure 2b). After another short period of absence after lunchtime, the household becomes occupied again and the home entertainment system is in operation. From the total electricity consumption it can be seen that around 6 p.m. the occupants prepare dinner. Shortly before midnight, the electricity consumption falls to the night time mean and all the home entertainment system is switched off.

4. OCCUPANCY CLASSIFICATION

In order to derive occupancy information from the electrical load curve, it is necessary to identify features that may be indicative of occupants being present in the household. A clear indicator for occupancy are switch events in the load curve that require the interaction of an occupant (e.g. television, stove or kettle). The electricity consumption induced by appliances such as fridges, freezers or the standby consumption of electric devices (e.g. the consumption of the digital video recorder) on the other hand does not give any indication about the occupancy state of the household. As introduced in section 2, a number of authors have looked at non-intrusive load monitoring approaches to detect the consumption of individual appliances from the electric load curve. However, such approaches require extensive training periods and are susceptible to changes in the number of appliances installed in the household. In the following section, we therefore identify a set of features of the electrical load curve that relate to the operation of occupancy-relevant appliances.

4.1 Features used for classification

Such features can be found by comparing the day-time electricity consumption during periods of occupancy to times when the household is unoccupied. Since we are concerned with classifying occupancy, we consider the intervals from 6 a.m. to 10 p.m. in our analysis and leave the detection
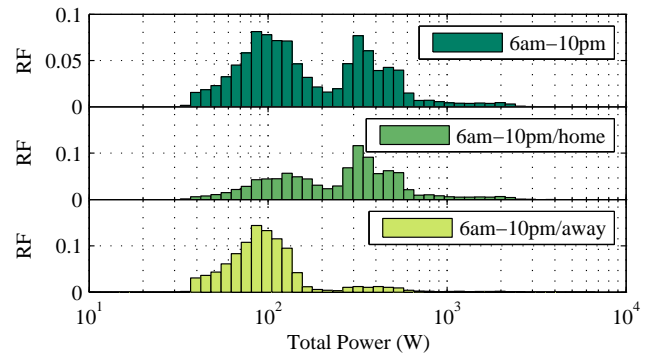

Figure 3: Relative frequencies of various total power consumption (sum of all phases) values over the whole day and divided into presence and absence respectively.

of sleep patterns for future work. Figure 3 shows the relative frequency (empirical probability) of the logarithmically binned total power consumption measurements over summer and winter periods for household 2. The figure shows the probability to see a particular measurement during presence or absence of the occupants. The top graph shows the total distribution of power consumption measurements during daytime. The two graphs below show the distribution during presence and absence periods, respectively. From figure 3 we can see that the power consumption is likely to have a higher mean and standard deviation when the household is occupied. During periods of absence, the electricity consumption is centered around 100 Watt and may be clearly distinguished from the overall day-time curve. While the household is occupied, the probability of higher consumption figures increases. However, there is still a significant probability to see lower consumption values even when the household is occupied. This is due to the fact that occupants may be at home but not using any electrical devices.

Table 4 shows the features selected to represent these observations. The suffixes `_l1` to `_l3` denote the three electrical phases. We computed the mean and standard deviation over 15-minute intervals. Given the sampling frequency of 1 Hz, each feature is thus computed from a 900-element vector. As the smart electricity meter provides a breakdown of the electricity consumption for all three individual phases, we have three features for the mean and standard deviation, each. The boilers in our participating households were programmed to operate during the night. Therefore, a high mean consumption is likely to be caused by presence in the household. A high standard deviation furthermore indicates that there have been significant changes in the electricity consumption during the observed interval. Such changes may have been caused by human involvement (e.g. by operating the stove or the kettle) or by appliances with varying consumption patterns. As the standard deviation only measures the distance to the mean of the data we have introduced another measure – the sum of absolute differences (SAD, features 7–9). The SAD computes the absolute difference between adjacent power measurements and adds them up, giving another measure of the variability of the data. Feature 10 is the prior probability of the household being occupied during a particular 15-minute interval of the day. This probability is computed as the average occupancy dur-

Table 4: Features used in the classification – all features computed over a 15-minute interval (σ: Standard deviation, SAD: Sum of absolute differences).

#	Feature	Description	Sensor
1	p_mean_l1	mean of power phase 1	Landis+Gyr E750
2	p_mean_l2	mean of power phase 2	Landis+Gyr E750
3	p_mean_l3	mean of power phase 3	Landis+Gyr E750
4	p_sd_l1	σ of power phase 1	Landis+Gyr E750
5	p_sd_l2	σ of power phase 2	Landis+Gyr E750
6	p_sd_l3	σ of power phase 3	Landis+Gyr E750
7	p_sad_l1	SAD of power phase 1	Landis+Gyr E750
8	p_sad_l2	SAD of power phase 2	Landis+Gyr E750
9	p_sad_l3	SAD of power phase 3	Landis+Gyr E750
10	prior	mean 24h occupancy	ground truth
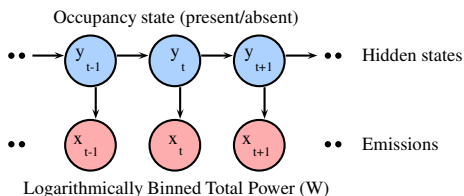
Figure 4: The HMM classifier uses a probabilistic model based on the distribution of the power measurements to switch between states.

ing this time interval over the training set.

4.2 Classification algorithms

In the following analysis we discuss the results obtained from testing three stateless and one stateful classifier on our winter and summer data sets. The stateless classifiers are support vector machines (SVM), K-nearest neighbour (KNN) and thresholding (THR). To evaluate stateful classification we use a hidden markov model (HMM).

For the KNN classifier we used the ClassificationKNN classes from the Matlab Statistics Toolbox. To implement the SVM classifier we used the LIBSVM library by Chang and Lin [5]. In addition, we introduced a simple classifier based on thresholding. The THR classifier computes the mean over the features during all unoccupied intervals. For each feature it thus computes a threshold above which it labels the interval as occupied. The final classification of an interval is based on a majority vote of the thresholding applied to all 10 features. The thresholding classifier implicitly assumes that a higher mean electricity consumption and a higher variance and sum of absolute differences between subsequent measurements relate are positively correlated with occupancy.

In contrast to the KNN, SVM and THR classifiers, the HMM only uses features 1-3 (see table 4) for its classification. A HMM (see figure 4) relates its hidden states (e.g. occupied, unoccupied) to emission (e.g. the observed electricity consumption) using emission and transition probabilities. The training of a HMM requires a discrete set of emissions. Since our features are continuous real values, we follow the approach shown in the figure 3 and obtain a set of possible emissions by logarithmically binning the training data into 20 bins. For this purpose we compute the mean power total power (sum of all phases). From this and the known occupancy states we estimate the emission and transition probabilities.
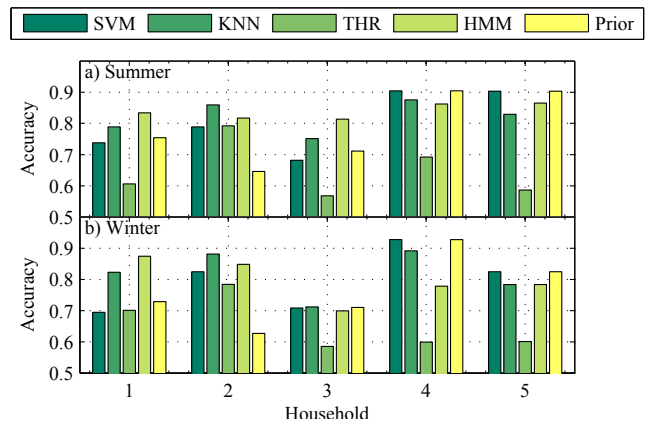
Figure 5: Accuracy of the SVM, KNN, THR, HMM classifiers compared to the baseline Prior.

4.3 Classification performance

In this section we discuss the performance of the classification algorithms. For all experiments we have chosen a two-fold cross validation, which randomly splits the data into two equally sized subsets. These two data sets are then used in turn for training and testing the classifiers.

In order to analyse the performance of the algorithms we follow the receiver operation characteristic notation. An instance of correctly assessing the household's occupancy from the electrical consumption data to be *occupied* during an interval is thereby named a *true positive* (TP). Likewise, correctly labelling the data to be *unoccupied* we will call a *true negative* (TN). Then *false positive* (FP) and *false negative* (FN) denote the instances of incorrectly labelling the household *occupied* or *unoccupied*, respectively. The *accuracy* of a classifier c is then computed as $A_c = \frac{TP+TN}{TP+TN+FP+FN}$.

4.3.1 Classification accuracy

Figures 5a and 5b show the accuracy obtained by the four classifiers for the summer and winter data sets, respectively. The same data is provided in table 5. *Prior* is a maximum likelihood classifier that always assigns an input data to the class of the majority of data points in the training set. Since in our data set the households are occupied more than 50% of the time, *Prior* always classifies the households as occupied. We use the accuracy of the *Prior* classifier as a baseline for the other methods. Applied on our data set, the baseline returns accuracies between 63% (household 2, winter) and 95% (household 4, winter). Thus, for households 1-3, the classification algorithms return better or equal results with respect to the baseline. Households 4 and 5 must be treated differently as they have at least one occupant who stays at home most of the time.

For both winter and summer data sets and all households except for household 2, the SVM classifier performs at or just below the accuracy of *Prior*. The SVM always labels the household to be occupied. This means that it correctly classifies all those states during which the household was actually occupied, but incorrectly classifies all states in which the house was unoccupied. This results in an accuracy equal to the accuracy of *Prior*, but means the results obtained from the classifier are not useful in practical settings.

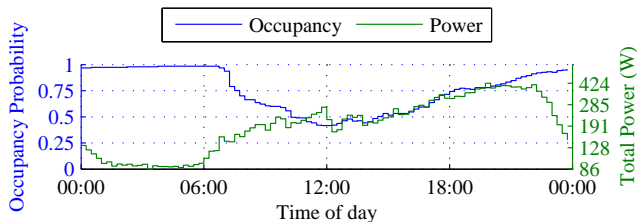The KNN classifier performs best on household 2 achiev-

Figure 6: Mean binary occupancy and mean electricity consumption (household 2, all data).

Table 5: Accuracy of the classifiers.

Household	SVM	KNN	THR	HMM	Prior
	summer				
1	0.74	0.79	0.61	**0.83**	0.75
2	0.79	**0.86**	0.79	0.82	0.65
3	0.68	0.75	0.57	**0.81**	0.71
4	0.90	0.88	0.69	0.86	0.90
5	0.90	0.83	0.59	0.87	0.90
	winter				
1	0.69	0.82	0.70	**0.87**	0.73
2	0.82	**0.88**	0.78	0.85	0.63
3	0.71	0.71	0.59	0.70	0.71
4	0.93	0.89	0.60	0.78	0.93
5	0.82	0.78	0.60	0.78	0.82

Table 6: Matthews correlation coefficients.

Household	SVM	KNN	THR	HMM
	summer			
1	0.22	0.44	0.40	**0.63**
2	0.57	**0.70**	0.55	0.67
3	0.12	0.42	0.38	**0.60**
4	/	0.31	0.37	**0.42**
5	/	0.02	**0.19**	0.13
	winter			
1	0.09	0.57	0.47	**0.72**
2	0.63	**0.75**	0.53	0.71
3	-0.02	0.26	**0.31**	0.19
4	/	**0.21**	0.17	0.07
5	/	**0.24**	0.18	0.03

ing an accuracy of 86% on the summer data and 88% on the winter data. The performance can be explained by the strong linear correlation between occupancy and power consumption. Figure 6 shows the mean binary occupancy over 24 hours plotted against the mean electricity consumption averaged at 15-minute intervals. During our classification period from 6 a.m. to 10 p.m. the two curves are almost moving in lockstep. The KNN classifier beats the prior accuracy for households 1-3 in both summer and winter. Only for households 4 and 5 it comes second to the SVM classifier.

Finally, the thresholding (THR) classifier performs consistently worse than the other classifiers and exceeds the baseline only for household 2. Even though it does not use the prior occupancy as an input feature, the HMM classifier performs best amongst households 1-3, achieving accuracies over 80% in 5 out of 6 cases.

4.3.2 Limitations of the accuracy

The classification accuracy describes only partially the performance of a classifier. Especially for data with unbalanced classes as witnessed in our example – households 4 and 5 have occupancy figures exceeding 90% – a high accuracy may be achieved by always predicting the household to be occupied. In the case of occupancy detection, correctly classifying both occupied (true positive) and unoccupied (true negative) states is paramount. For this reason we computed the Matthews correlation coefficient (MCC) over the results of our classifiers [15]. A perfect prediction is represented by an coefficient of +1. On the other hand, a value of −1 indicates that no single instance was classified correctly. A coefficient of 0 represents a classification, which is no better than a random guess. The MCC of a classifier c is calculated as: $\mathrm{MCC}_c = \frac{TP \times TN - FP \times FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}}$. The MCC provides a balanced measure even if the input data are heavily skewed towards one class.

Table 6 shows MCC_c for the SVM, KNN, THR and HMM classifiers. The classifier achieving the highest MCC in each case is highlighted in bold print. The coefficient for the SVM classifier on households 4 and 5 could not be calculated since these are never classified as unoccupied and therefore $TN + FN = 0$. For both summer and winter, the performance of KNN and HMM reflect the results obtained in terms of accuracy. The KNN classifier performs best on household 2 with coefficients of 0.7 and 0.75, respectively. The HMM classifier performs best on household 1 with coefficients of 0.63 (summer) and 0.73 (winter). The THR classifier generally performs better than the accuracies in figure 5 indicate. For household 3, it gives the best classification during winter with a coefficient of 0.31, outperforming both KNN (0.26) and HMM (0.19).

Both accuracy and MCC assess the performance of a classifier based on the correct classification of individual intervals, independently of each other. Any correct or incorrect classification of an interval contributes with the same weight to the metric. The ability to detect *occupancy transitions* – i.e. changes in the occupancy state (from occupied to unoccupied and vice versa) – is however crucial to many systems. For instance, when a smart heating system detects that the household has become occupied, it may decide to start heating immediately. As every transition corresponds to a switch event in the controller, correctly identifying the number of transitions is of equal importance to the accuracy of the classification itself. Figure 7 shows the classification for the first 200 15-minute intervals of household 2. The ground truth occupancy data shows only 6 transitions. Due to their inherent statelessness, the three classifiers (SVM, KNN and THR) identify a number of additional transitions. Such transitions must be filtered out before the current occupancy state is passed on to a controller. The HMM incorporates this step by taking into account the transition probability between occupied and unoccupied states at different consumption levels.

5. CONCLUSIONS AND FUTURE WORK

In this paper we presented and evaluated an approach that leverages electricity meters as occupancy sensors. We performed our analysis using a data set that we collected during an 8-month long experiment run in 5 households. During this period we gathered data from digital electricity meters, smart plugs and PIR sensors as well as ground-truth occupancy data. Our results show that occupancy detection accuracies over 80% are feasible in most scenarios. Our next steps include the investigation of sensor fusion methods to incorporate the other sensor data gathered in this deployment to improve the overall occupancy detection accuracy.
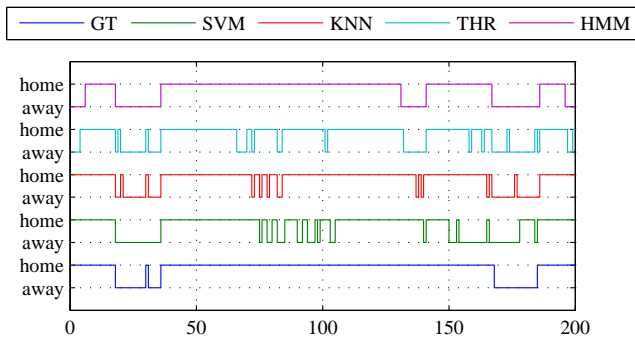
Figure 7: Ground truth (GT) and classification results for the first 200 intervals of household 2.

Acknowledgments

6. REFERENCES

[1] J. M. Abreu, F. P. Câmara, and P. Ferrão. Using pattern recognition to identify habitual behavior in residential electricity consumption. *Energy and Buildings*, 49:479–487, 2012.

[2] S. Barker, A. Mishra, D. Irwin, E. Cecchet, P. Shenoy, and J. Albrecht. Smart*: An open data set and tools for enabling research in sustainable homes. In *Proc. SustKDD'12*. ACM, Aug. 2012.

[3] C. Beckel, L. Sadamori, and S. Santini. Automatic socio-economic classification of households using electricity consumption data. In *Proc. e-Energy'13*. ACM, May 2013.

[4] K. Carrie Armel, A. Gupta, G. Shrimali, and A. Albert. Is disaggregation the holy grail of energy efficiency? The case of electricity. *Energy Policy*, 52(C):213 – 234, 2013.

[5] C.-C. Chang and C.-J. Lin. LIBSVM: A library for support vector machines. *ACM Trans. on Intelligent Systems and Technology*, 2:27:1–27:27, 2011.

[6] D. De Silva, X. Yu, D. Alahakoon, and G. Holmes. A data mining framework for electricity consumption analysis from meter data. *IEEE Trans. on Industrial Informatics*, 7:399–407, 2011.

[7] R. Dickerson, E. Gorlin, and J. Stankovic. Empath: A continuous remote emotional health monitoring system for depressive illness. In *Proc. Wireless Health'11*. ACM, Oct. 2011.

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

[9] S. Gupta, M. Reynolds, and S. Patel. ElectriSense: Single-point sensing using EMI for electrical event detection and classification in the home. In *Proc. UbiComp'10*. ACM, Sept. 2010.

[10] G. Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870–1891, 1992.

[11] H. Kim, M. Marwah, M. Arlitt, G. Lyon, and J. Han. Unsupervised disaggregation of low frequency power measurements. In *Proc. SDM'11*. SIAM, Apr. 2011.

[12] W. Kleiminger, C. Beckel, and S. Santini. Opportunistic sensing for efficient energy usage in private households. In *Proc. SES'11*, Sept. 2011.

[13] J. Z. Kolter and M. J. Johnson. REDD: A public data set for energy disaggregation research. In *Proc. SustKDD'11*. ACM, Aug. 2011.

[14] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse. The smart thermostat: Using occupancy sensors to save energy in homes. In *Proc. SenSys'10*. ACM, Nov. 2010.

[15] B. W. Matthews. Comparison of the predicted and observed secondary structure of t4 phage lysozyme. *Biochimica et Biophysica Acta (BBA)-Protein Structure*, 405(2):442–451, 1975.

[16] M. Milenkovic and O. Amft. An opportunistic activity-sensing approach to save energy in office buildings. In *Proc. e-Energy'13*. ACM, May 2013.

[17] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin. Private memoirs of a smart meter. In *Proc. BuildSys'10*. ACM, Nov. 2010.

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

[19] B. Ostermaier, M. Kovatsch, and S. Santini. Connecting things to the web using programmable low-power wifi modules. In *Proc. WoT'11*, June 2011.

[20] I. Sánchez, I. Espinós, L. Moreno Sarrion, A. López, and I. Burgos. Clients segmentation according to their domestic energy consumption by the use of self-organizing maps. In *Proc. EEM'09*. IEEE, May 2009.

[21] J. Scott, A. Bernheim Brush, J. Krumm, B. Meyers, M. Hazas, S. Hodges, and N. Villar. Preheat: Controlling home heating using occupancy prediction. In *Proc. UbiComp'11*. ACM, Sept. 2011.

[22] J. Taneja, A. Krioukov, S. Dawson-Haggerty, and D. Culler. Enabling advanced environmental conditioning with a building application stack. Technical Report UCB/EECS-2013-14, Electrical Engineering and Computer Sciences, University of California at Berkeley, 2013.

[23] B. Vande Meerssche, G. Van Ham, G. Deconinck, J. Reynders, M. Spelier, and N. Maes. Practical use of energy management systems. In *Proc. AmiEs'11*, Sept. 2011.

[24] S. V. Verdú, M. O. Garcia, C. Senabre, A. G. Marín, and F. J. G. Franco. Classification, filtering, and identification of electrical customer load patterns through the use of self-organizing maps. *IEEE Trans. on Power Systems*, 21(4):1672–1682, 2006.

[25] M. Zeifman and K. Roth. Nonintrusive appliance load monitoring: Review and outlook. *IEEE Trans. on Consumer Electronics*, 57(1):76–84, 2011.

[26] A. Zoha, A. Gluhak, M. A. Imran, and S. Rajasegarar. Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey. *Sensors*, 12(12):16838–16866, 2012.