



Exploring zero-training algorithms for occupancy detection based on smart meter measurements

Vincent Becker¹  · Wilhelm Kleiminger¹

Published online: 30 August 2017
© Springer-Verlag GmbH Germany 2017

Abstract Detecting the occupancy in households is becoming increasingly important for enabling context-aware applications in smart homes. For example, smart heating systems, which aim at optimising the heating energy, often use the occupancy to determine when to heat the home. The occupancy schedule of a household can be inferred from the electricity consumption, as its changes indicate the presence or absence of inhabitants. As smart meters become more widespread, the real-time electricity consumption of households is often available in digital form. For such data, supervised classifiers are typically employed as occupancy detection mechanisms. However, these have to be trained on data labelled with the occupancy ground truth. Labelling occupancy data requires a high effort, sometimes it even may be impossible, making it difficult to apply these methods in real-world settings. Alternatively, one could use unsupervised classifiers, which do not require any labelled data for training. In this work, we introduce and explain several unsupervised occupancy detection algorithms. We evaluate these algorithms by applying them to three publicly available datasets with ground truth occupancy data, and compare them to one existing unsupervised classifier and several supervised classifiers. Two unsupervised algorithms perform the best and we find that the unsupervised classifiers outperform the supervised ones we compared to. Interestingly, we achieve a similar classification performance on coarse-grained aggregated datasets and their fine-grained counterparts.

Keywords Occupancy detection · Smart meters · Unsupervised classification

✉ Vincent Becker
vincent.becker@inf.ethz.ch

¹ Department of Computer Science, ETH Zurich, Zurich, Switzerland

1 Introduction

Occupancy, i.e. whether the inhabitants of a dwelling are at home or not, is one of the major contextual features used in smart home applications. Determining occupancy patterns as shown by the examples in Fig. 1a, b could be used to control the heating, electronic devices, the burglar alarm, etc.

There are several ways to determine the presence or absence of inhabitants (cf. Sect. 4). A common approach is installing sensors in the dwelling, such as reed switches on the main doors or motion detectors indoors. However, these approaches are relatively obtrusive and require the installation of dedicated hardware. Another possibility is to use location-connected services on smartphones, such as the inhabitants' GPS location or the Wi-Fi networks their smartphones are connected to. However, the inhabitants would have to carry their smartphone with them at all times for this to work.

A different and promising possibility relies on monitoring the electricity consumption of the household. Previous research has shown that it is possible to detect occupancy from electrical load data using machine learning algorithms with sufficiently high accuracy [25, 27, 28]. Indeed, electrical load data is a good proxy for a household's occupancy since its magnitude and changes in the power consumption are indicators for human activity (i.e. interactions with appliances) in the household. At the same time, smart electricity meters, which continuously measure the electrical power demand of a household, are becoming more and more ubiquitous. In sixteen EU member countries, a smart meter penetration rate of 95% is expected by 2020 [18]. This large-scale deployment of smart meters makes it increasingly viable to use their measurements for purposes like occupancy detection.

The task of occupancy detection in our context can be defined as follows: Given a time series of electricity con-




Fig. 1 Average weekly schedules for two different households. The higher the value (as displayed by the *colour*) in a time slot, the likelier the home is occupied during that time slot (colour figure online)

sumption measurements, determine (i.e. estimate with a sufficiently high probability) whether the home is occupied or not for each time interval. We thus face a classification problem with two classes, *occupied* and *unoccupied*.

In machine learning, there are two main types of classifiers: supervised and unsupervised. Supervised classifiers have to be trained on data labelled with ground truth in order to learn the relevant patterns. In terms of our setting, electricity consumption samples labelled with the true occupancy class of the specific household would have to be provided. This labelling process entails a great effort, yielding supervised algorithms difficult to apply in a real-world scenario, where labelled data is not easily available. By contrast, unsupervised algorithms do not require any labelled samples, hence the term zero-training.

In this work we address two challenges: First, creating and exploring different unsupervised classifiers for occupancy detection, and second, coping with coarse-grained electricity consumption data with a sampling interval of half an hour. The classifiers presented are very lightweight and could easily be run locally, i.e. within the home without having to disclose information to the outside. In more detail, our contributions are:

- Developing unsupervised classifiers for occupancy detection from generally available electricity consumption data (only energy measurements, no voltages or currents).
- Being able to handle coarsely grained data at a sampling interval of 30 min.
- Validating and evaluating the algorithms on three publicly available datasets containing electrical energy consumption and ground truth occupancy values; further,

comparing them to previous algorithms including supervised classifiers.

The remainder of the paper is structured as follows: In Sect. 2, we show the design of our occupancy detection algorithms. We evaluate them in Sect. 3. In Sect. 4, we discuss related work done on occupancy detection. Finally, we draw conclusions in Sect. 5.

2 Occupancy detection from electricity consumption data

Our aim is to determine the occupancy state of a household by analysing its electricity consumption. We assume a coarse-grained sampling interval of 30 min. First, we pre-process the data. We take the logarithm of all power values (to the base 10) and use a moving average filter with a window size of 5 to smooth the data.

Then, we perform the classification, using one of the unsupervised classifiers mentioned below. The classifier assigns a label (either occupied or unoccupied) to each sample (i.e. 30 min interval). The sequence of the resulting labels is the occupancy schedule. After classification, we post-process the schedule by again performing moving average smoothing on the schedule. The whole process is depicted in Fig. 2.

In the following we will detail on the three unsupervised occupancy detection algorithms we developed: a Hidden Markov Model, a model using a geometric moving average, and one using a Page–Hinkley test. We face two challenges: firstly, in a real-world scenario, there are no labels available, i.e. the electricity data is not annotated with the occupancy ground truth. Hence, the use of supervised classifiers is not possible. Secondly, we want to be able to deal with



Fig. 2 The pipeline for occupancy detection

coarse-grained electricity consumption data. For a large sampling interval the raw data already gives a very aggregated view of the household's electricity consumptions. Therefore, the possibilities to calculate features over time windows are limited since one feature would aggregate a long period of time. To compare the algorithms and measure their quality, we apply each of them to three datasets for which occupancy ground truth is available, i.e. it is possible to determine how well they perform (cf. Sect. 3).

2.1 First algorithm: Hidden Markov Model (HMM)

The occupancy of a household can be modelled as a probabilistic model with two states—*occupied* and *unoccupied*. At any point in time, there is a certain probability of the household changing from one to the other. This kind of system can be represented by a Hidden Markov Model (HMM), a statistical model, through which a process is modelled by a Markov chain with hidden states. This means that the model randomly changes from one state to another with certain transition probabilities depending only on the current state. The states cannot be observed themselves, they are hidden and instead only the states' emissions are observable which emerge with certain probabilities depending on the emitting state. In our case the HMM models the binary occupancy of a home and hence has only two states, *unoccupied* and *occupied*. The emissions are power consumption values which are

drawn from emission distributions. The model is depicted in Fig. 3.

Figure 4 shows an example for a household from 4 a.m. to 12 a.m.

The only information we can observe are the power consumption values for each 30 min time slot. The goal is to find the most probable state sequence to a given sequence of observations (also known as decoding), which is solved by the Viterbi algorithm [48]. Usually, the parameters of the model would be learnt using a training algorithm, e.g. the Baum–Welch algorithm [38] following an expectation-maximisation approach. HMMs have been used for occupancy detection from electricity data in this supervised manner before [25,27,28]. For that however, training data would have to be available. Thus, we determine the emission distribution and transition probability on basic assumptions we make, which we detail on in the following.

Transition probabilities The transition probabilities define how probable it is in each time slot for the state to change from unoccupied to occupied or from occupied to unoccupied, respectively. In our method a single day has 48 half-hour time slots. The transition probabilities depend on the expectation how long the home is unoccupied or occupied, respectively and how many “leave” and “return” events there are. We calculate $\alpha = \frac{\#return}{\#unoccupied}$ and $\beta = \frac{\#leave}{\#occupied} = \frac{\#leave}{48 - \#unoccupied}$. Since it is the most common, we assume a typical working

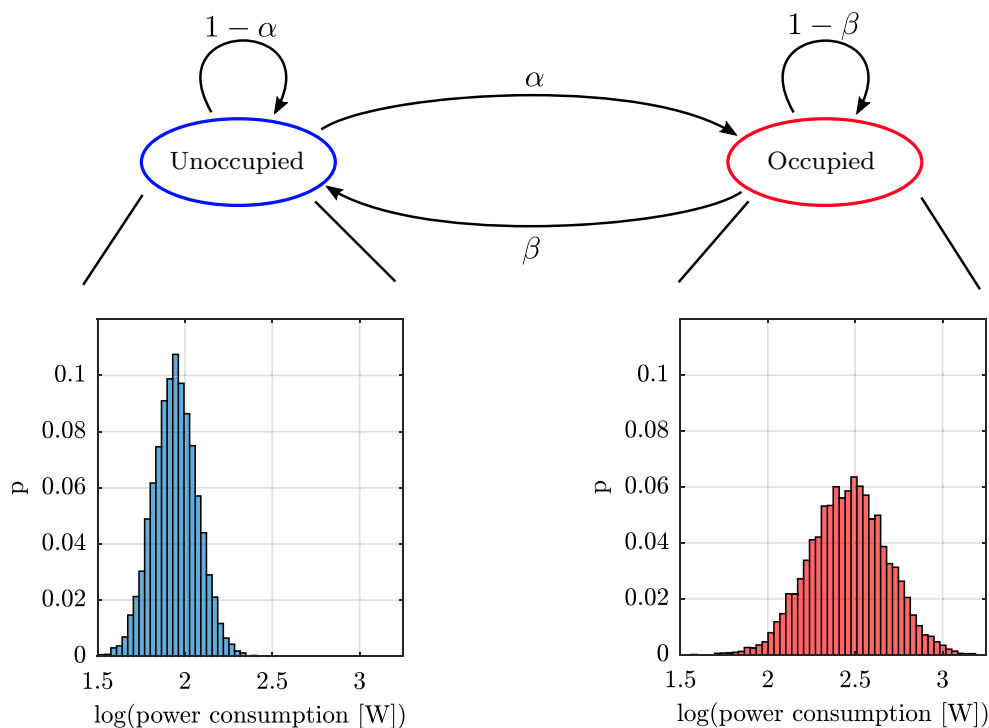


Fig. 3 The HMM for occupancy detection. Each state emits power values from a certain emission probability distribution and the transition from one state to the other takes place with a certain probability in each step




Fig. 4 An example of the 30min aggregated power consumption for a specific household from 4 a.m. to 12 a.m. The states are either occupied or unoccupied. The colours (*red* implying occu-

piated and *blue* unoccupied) are a guess for a schedule and depict the non-trivial task of estimating the state sequence (colour figure online)

day schedule in which the home is unoccupied nine continuous hours a day with a single “leave” event and a single “return” event. Hence there are 30 occupied and 18 unoccupied slots out of the 48 slots per day in total. Thus we calculate the transition probabilities by $\alpha = \frac{1}{18}$ and $\beta = \frac{1}{30}$. If further knowledge such as a rough estimation of the schedule of the household was available, this could be easily adjusted.

Emission probabilities For the emission probabilities we have to find a set of samples which we assume to belong to the unoccupied and the occupied state, respectively. To do so we use the mean over all power values from the household’s data as threshold and assume that all samples below the mean belong to the unoccupied state and all above or equal to the occupied state. In our experiments the mean has proven to be a good heuristic for a threshold to separate occupied from unoccupied emissions. For each sample set we fit a normal distribution and use this as the emission distribution.

2.2 Second algorithm: geometric moving average (GeoMA)

The motivation behind this strategy is that periods of absence will decrease the moving average of the electricity consumption. As soon as the inhabitants are home again the consumption will increase. The average will too, but naturally not as fast. Hence, the momentary electrical consumption will rise above the average and we will signal occupancy.

Implementing this idea, the algorithm using the geometric moving average follows a simple strategy. In each time step we calculate the geometric moving average. If the current sample is greater than the average, we set the schedule for the current time slot to *occupied*, otherwise to *unoccupied*. The procedure is shown in “Appendix” and Fig. 5 depicts an example for a single day.

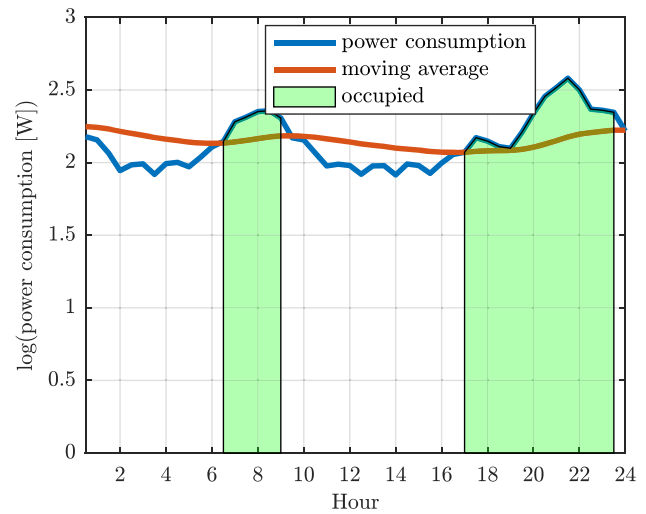


Fig. 5 An example for the geometric moving average, marked as the *red* line. Whenever the momentary electricity consumption (shown as the *blue* line) is higher than the geometric moving average, the household is considered to be occupied. The estimation for the given data is shown by the *green* areas (colour figure online)

2.3 Third algorithm: Page–Hinkley test (PHT)

The Page–Hinkley test [34] is an unsupervised concept change detection algorithm. In the area of data streams (a potentially unbounded sequence of data points, such as our power consumption values), the concept is considered as the probability distribution generating the stream data. In our case we can imagine two concepts, the unoccupied and occupied home, which incorporate two different distributions emitting the data. The aim is to find the changes, i.e. when the stream process moves from one concept (i.e. distribution) to the other, which corresponds to the change from unoccupied to occupied or vice versa. The Page–Hinkley test detects changes in signals by observing the difference of

cumulative variables from adapting averages. Basically, it is a more sophisticated version of the geometric moving average explained above. The procedure, fitted to our application, is shown in “Appendix”. The Page–Hinkley test can detect increasing and decreasing changes. If we find an increasing change, we set the schedule for that slot to 1, i.e. occupied, for a decreasing change to 0, i.e. unoccupied. In case we detect no change we set the slot to the state in the previous slot.

2.4 NIOM (non-intrusive occupancy monitoring)

NIOM [14] was presented by Chen et al. (cf. Sect. 4). We include it in our comparison of the three previously mentioned methods (and the supervised algorithms). NIOM calculates three features over a time window. These are the average, the standard deviation, and the maximum range of values in the window. The home is considered to be *occupied* in a specific time slot if one or more of the features are above a certain threshold. The thresholds are determined by the maxima of the features during the previous night. Thereby, the thresholds are dynamic. If any two time slots are detected to be *occupied* and are within a certain window $\tau_{cluster}$, then all slots in between are set to *occupied*. If the home was *occupied* in the evening, then it is also considered to be *occupied* during the night.

2.5 Adding a nightly schedule

Detecting nightly occupancy merely from electricity consumption data is not a simple task, since during sleep people do not interact with electrical devices and most of them are turned off or in standby mode. Similar to the authors of [14] we resort to an additional simple rule-based approach by adding a nightly schedule when applying any of our algorithms. If we detect occupancy with a duration of at least 1 h from 8 p.m. to midnight, we set the state of each time slot for the following night (until 9 a.m.) to occupied, beginning with the slot which is the last to be detected as occupied.

3 Validation and evaluation

We test our three unsupervised algorithms HMM, GeoMA, and PHT on three publicly available labelled datasets to be able to assess their performance. Due to the high effort and costs of annotating the data, such datasets are relatively rare. We downsampled each dataset by averaging to a sampling interval of half an hour to show we can indeed handle coarse-grained data. This downsampling does not impair the performance of the classification in most cases. Rather, it is often improved (cf. Sect. 3.5).

For comparison we also apply NIOM and three supervised algorithms, k-Nearest Neighbours ($k = 5$) (KNN), a

Support Vector Machine (SVM) with an RBF-kernel, and a random forest (RF). Additionally we show a baseline, which assumes that the home was occupied in every time slot. The baseline is a lower bound for the performance the other classifiers should achieve. For all supervised classifiers we use the standard implementations in MATLAB. For evaluation, we use 10-fold cross-validation (90% training and 10% testing). The features we use for the supervised classifiers are the mean, the standard deviation, the sum of absolute differences, and the maximum of the difference in a time window of two observations. Naturally, these supervised algorithms cannot be applied in real-world scenarios without any prior training. For GeoMA we set the adaptation rate $\lambda = 0.05$. For PHT we use 0.05 for the magnitude threshold and 0.3 for the detection threshold.

As metrics for the performance of a classifier we use the accuracy (ACC) and the Matthews correlation coefficient (MCC , [32]), which are defined in Eqs. 1 and 2. The bounds for the accuracy are 0 and 1, for the MCC -1 and 1 . In both cases a higher number indicates a better result. The ACC result may be misleading in case of an unbalanced class distribution (as it is in our case, since the homes are occupied more than they are unoccupied), whereas the MCC has the advantage that it compensates for skewed classes.

$$ACC = \frac{tp + tn}{tp + tn + fp + fn} \quad (1)$$

$$MCC = \frac{tp * tn - fp * fn}{\sqrt{(tp + fp)(tp + fn)(tn + fp)(tn + fn)}} \quad (2)$$

The arguments are the number of true positives (tp), false positives (fp), true negatives (tn), and false negatives (fn). If the classifier assigns all samples to one class, the MCC cannot be calculated and thus will not be shown in the results in such cases.

3.1 Dataset A: Tang et al.

The dataset is presented in [43]. It contains the electricity and occupancy data over 1 month for a single household in Victoria, BC, Canada. The consumption data was collected using off-the-shelf measuring devices and the occupancy information was derived from the GPS traces of the inhabitants’ mobile phones. The collection period was from the 23rd February 2015 to 23rd March 2015. The original sampling frequency was 0.1 Hz. Table 1 displays the results for the detection algorithms.

The HMM, the PHT, and the GeoMA perform the best, even better than the supervised algorithms. The baseline assumes the house to be occupied all the time. Thus, there are no true and false negatives and the MCC cannot be calculated. Since the occupancy is relatively high (65%), the SVM seems to have learnt to always classify a slot as occupied, i.e. it behaves just like the baseline.

Table 1 The average results of Tang's dataset

Algorithm	ACC	MCC
Baseline	0.65	–
HMM	0.89	0.76
GeoMA	0.90	0.78
PHT	0.89	0.75
NIOM	0.85	0.67
KNN	0.70	0.33
SVM	0.65	–
RF	0.70	0.33

Table 2 The average results over the three parts of Chen's dataset

Algorithm	ACC	MCC
Baseline	0.77	–
HMM	0.90	0.73
GeoMA	0.91	0.73
PHT	0.89	0.68
NIOM	0.88	0.55
KNN	0.88	0.63
SVM	0.85	–
RF	0.90	0.69

3.2 Dataset B: Chen et al.

Chen's dataset [14] is part of the Smart* dataset [6] augmented with occupancy information, which are again obtained from the GPS traces of the inhabitants' smartphones. It contains three parts, spring (1st April 2013 to 7th April 2013) and summer (8th July 2013 to 14th July 2013) measurements for one house, and only summer measurements for another. The original sampling interval is 1 min. The households both are in Western Massachusetts, US. Figure 6 show the results on this dataset and the averages are in Table 2.

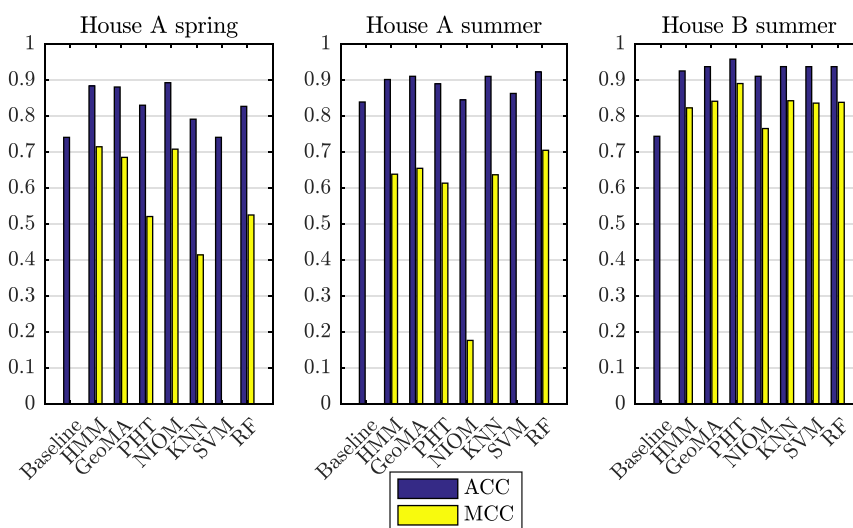
In terms of accuracy, all algorithms perform similarly well. For the MCC metric, however, the HMM and GeoMA are the best and NIOM, the algorithm which was evaluated on this dataset in [14], is outperformed.

3.3 Dataset C: ECO dataset

The ECO dataset presented in [7] contains the data of six households in Thun, Switzerland. It was collected over a

period of 8 months from June 2012 to January 2013. The data is split into two periods, summer and winter. The sixth household did not provide any occupancy information so we omit it here. To create occupancy ground truth the inhabitants manually registered presence and absence with a tablet. Additionally, a PIR sensor near the main door and several smart plugs were deployed in each household, connected to devices such as PCs, etc., to enhance the annotation. The original sampling rate is 1 Hz. Except household two, all households are occupied nearly all the time, hence it makes it difficult to perform better than the baseline. Figure 7 shows the results on this dataset and the averages are in Table 3.

The authors of [28] have already shown the difficulty of beating the baseline for this dataset, even with 1 Hz data and supervised learning algorithms. Here we work on 30 min sampling intervals and our methods do not compute features. The supervised algorithms perform the best, but often make predictions similar to the baseline, estimating occupancy for nearly all time slots. Among the unsupervised algorithms, the HMM and the GeoMA perform the best in terms of MCC.

**Fig. 6** The results on the dataset of Chen et al.

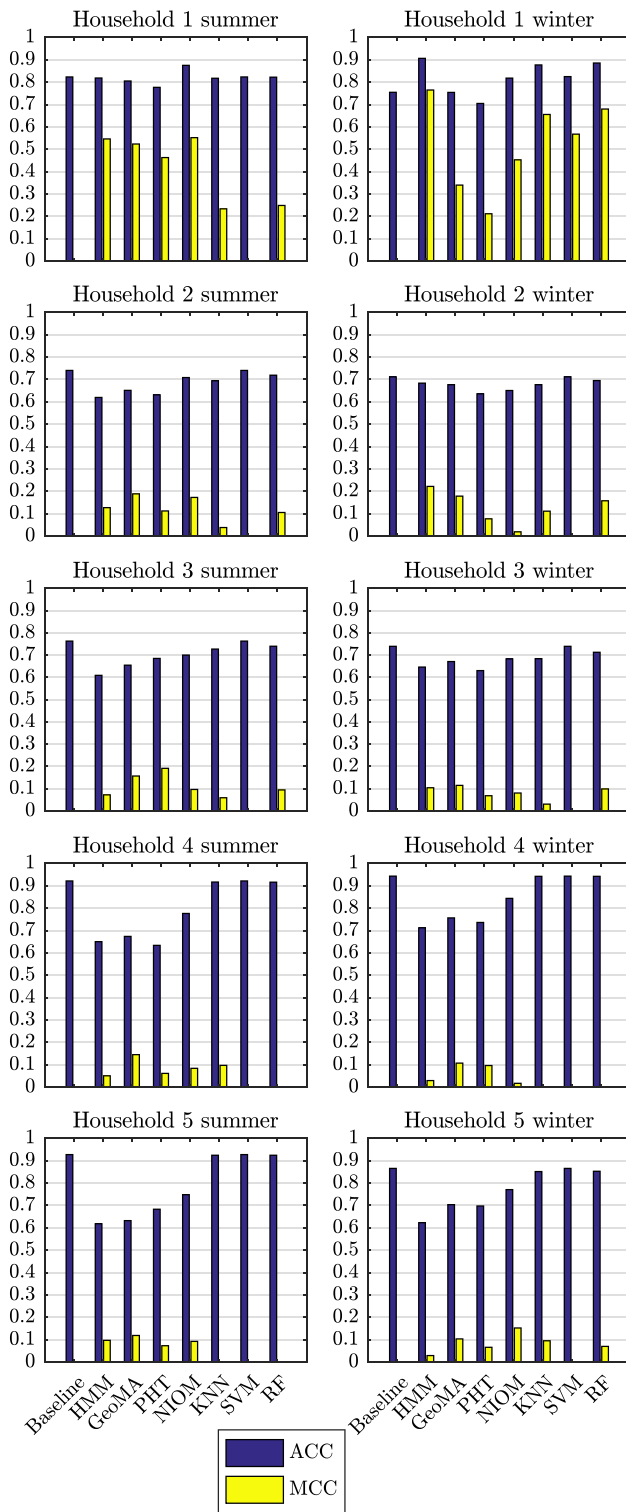


Fig. 7 The results on the ECO dataset

3.4 Overall results

In terms of MCC performance, which is a more appropriate performance measure than ACC, the HMM and the GeoMA

Table 3 The average results over the ECO dataset

Algorithm	ACC	MCC
Baseline	0.82	–
HMM	0.69	0.20
GeoMA	0.70	0.20
PHT	0.68	0.14
NIOM	0.76	0.17
KNN	0.81	–
SVM	0.83	–
RF	0.82	–

Table 4 The results of the HMM on the original datasets with higher sampling rates

Dataset	ACC	MCC
A (Tang)	0.74	0.78
B (Chen)	0.78	0.68
C (ECO)	0.75	0.38

The values show the average performance for each of the datasets

perform the best with the PHT as close runner-up. Note that the disadvantage of the HMM is that we need all the data prior to detection, while the GeoMA and the PHT work online. Whether the achieved classification performance is sufficient clearly depends on the application. We believe it is sufficient for many non-critical systems, such as an occupancy controlled heating or lighting system, which could be easily overruled by humans in case of false classification. In addition, since our algorithms achieve a significantly better performance than a random guess, they certainly would be beneficial when combined with other occupancy techniques to create an ensemble which achieves higher performance.

3.5 Performance with higher sampling rates

As mentioned before, we downsampled the datasets to a common sampling interval of 30 min. To show that this did not strongly decrease the performance we also evaluated the original datasets. The original sampling rates were 0.1 Hz for the dataset of Tang et al., once per minute for the dataset of Chen et al., and 1 Hz for the ECO dataset. We run the HMM on the original datasets to be able to compare to the downsampled version. The average classification performance for each dataset is shown in Table 4.

The results for the original and the downsampled version are similar for Tang’s dataset. For Chen’s dataset the results are better in the downsampled version. Only for the ECO dataset the results are significantly better using the original dataset.

4 Related work on occupancy detection

In this section we discuss the related work that has been done in the effort of detecting the occupancy in households. There are several approaches for occupancy detection, differing both conceptually and technologically.

4.1 Using the inhabitants' smartphones

One way to detect the occupancy of a household is to augment its members with devices which sense their location. This location information can then be used to check whether the inhabitants are at home or not. A device with localisation capabilities which many people already possess is a smartphone.

Gupta et al. [22] use GPS (Global Positioning System) information to calculate the inhabitants' distance to their home and employ a GPS-enabled thermostat to control the home's heating based on that distance. Thereby, the home can be reheated prior to the inhabitants' expected return. In that sense the systems actually applies occupancy *prediction* which naturally can be used for occupancy *detection*.

In the homeset algorithm, Kleiminger et al. [26] exploit the WLAN information sensed with an inhabitant's smartphone. In an initial stage a set of Wi-Fi networks, which are reachable from the home, is created. Whenever one of the networks in the set is reachable from the inhabitant's smartphone, the inhabitant is considered to be at home, otherwise not.

Smith et al. [40] use RFID tags on objects and a user-worn wristband to identify which activity a human is performing.

A disadvantage of these approaches is that the inhabitants have to carry their smartphones or other devices with them and that the location service has to be turned on at all times.

4.2 Using sensors inside the home

Often, a variety of sensors inside the home is used to directly detect the occupancy of a building or even single rooms.

In their review, Guo et al. [21] examine different occupancy detection sensors such as PIR sensors, ultra-sonic sensors, micro-wave sensors, light barriers, and video cameras in order to be able to control the lighting of a building. In their approach "Smart Thermostat", Lu et al. [31] use PIR sensors in rooms and magnetic reed switches on the main door of the home to create features. With these features, a Hidden Markov model is applied to infer the occupancy states, either being Active, Away, or Sleep. Using the occupancy information, the HVAC system of the house is controlled. To be able to preheat the home in time before the occupants arrive home again in order to avoid a loss in comfort, an occupancy schedule is set up incorporating historical occupancy data to predict arrival times. Hence the method is a mixture of occupancy detection and prediction.

Soltanaghaei et al. [41] present WalkSense, a system consisting of motion sensors distributed along the walkways of a home. Brown et al. [10] use an ultra-wideband radar module to sense occupancy. Woodstock et al. [50] employ time-of-flight sensors to detect occupancy. Wang et al. [49] measure the indoor CO₂ levels to detect if someone is present. Sensing changes in the CO₂ levels is a common approach in many other works [11, 15, 23, 42], also in combination with other environmental values, such as temperature [16, 17], light and humidity [12] and others [4, 37]. Zikos et al. [52] explore conditional random fields for fusing different combination of sensors including CO₂, motion, and acoustic sensors. Amayri et al. [3] calculate information gains to determine the most useful measurements. Teixeira et al. [45] present an approach to determine the number of people in a room using a camera sensor network. Gao and Whitehouse [20] present a self-programmable thermostat. The leaving and arrival time-points of the inhabitants are registered using simple sensors and with that data an automatic heating schedule is defined. Similarly, Barbato et al. [5] create user profiles from data gathered with a wireless sensor network in order to optimize energy consumption. Brackney et al. [9] present an image processing occupancy sensor which analyses video data aiming to detect humans in the pictures. It overcomes several disadvantages of PIR and ultra-sonic based motion sensors such as detecting humans who are not in motion. Furthermore it is possible to determine the number of people present. However, a video analysis system has severe privacy issues.

Moreover there are many other approaches which do not aim at occupancy detection directly but could be applied for that purpose. Patel et al. [35] sense differences in air pressure when doors are opened and closed using only a single sensor in a HVAC unit. The authors in [36] detect and classify electrical events by their pattern in the power line as occupants trigger switches using plug-in sensors. Froehlich et al. sense the water activity, e.g. the use of the kitchen sink, shower, or toilet with a single-point sensor [19]. Besides, there are more approaches for human activity recognition in homes using several types of sensors [44, 47]. There also are commercial providers for energy management systems using sensors for occupancy detection, such as motion or body heat sensors, in order to control the temperature of the home [46].

However, all these approaches entail the installation of sensors in the home.

4.3 Using electricity data from smart meters

As smart meters become more wide-spread, it is attractive to use their generated electricity data to infer the household occupancy. Basically, the smart meter is employed as a sensor in this approach, however, the advantage is that the sensor is

not installed solely for the purpose of occupancy detection and the inhabitants are not required to carry any devices with them.

When employing algorithms utilising energy consumption data, a common approach is to use supervised classification to determine the occupancy states. This means that labelled training data is necessary, i.e. the occupancy ground truth needs to be obtained. Usually, an individual classifier is trained for each household. Yang et al. [51] use conditional random fields to determine the number of people present in a household. As features the peak, mean, and variance for short time intervals were chosen. Moreover, they use supervised classifiers such as a random forest, a decision tree, KNN, NaiveBayes, and an MLP for the binary occupancy case and are able to reach accuracies up to 98% against a baseline of 88%. Akbar et al. [1] apply supervised machine learning methods such as KNN and SVM to detect if office desks are occupied. For this, they use energy meters at each desk. As features they use the real power, the root mean square of voltage and current, and the phase angle between them. They extend the binary setting and add a standby state to model short breaks of the people. Kleiminger et al. [25, 27, 28] also use supervised classifiers such as SVM, KNN, and a HMM to detect occupancy on the ECO dataset (cf. Sect. 3). They achieve accuracies of more than 80%. Chaney et al. [13] use an HMM as well, but combine electricity with sensed CO₂ levels and dew point temperature utilising the Dempster–Shafer theory for sensor fusion. Boait and Rylatt [8] exploit the electricity load and hot water usage data to infer occupancy. Additionally, they incorporate occupancy data from the previous week to set a prior probability of occupancy. This information is combined using the Bayes rule to infer the a-posteriori probability of the home being occupied and hence to control the heating system. Jin et al. [23] try to reduce the need for supervision by applying transfer learning, i.e. using supervised classifiers trained on similar households.

Another method of analysing electrical data is Non Intrusive Load Monitoring (NILM, [7]). The aim is to determine when single appliances are turned on or off from the aggregate electricity data. From the appliances' states, the occupancy of a home could be derived. However, NILM methods usually rely on high sampling rates and need additional information about the appliances. Alhamoud et al. [2] use appliance-level power consumption to derive human behavioural patterns.

The requirement of having to obtain ground truth data for supervised classification is a problem for the application of occupancy detection in practice. However, unsupervised classification is in principle more difficult, since the household's consumption patterns are unknown. Chen et al. [14] present their threshold-based algorithm NIOM (Non-Intrusive Occupancy Monitoring, cf. Sect. 2.4) which signals

occupancy as soon as one of the features exceeds its corresponding threshold. As features they use the mean, the standard deviation, and the maximum range of the electricity power data. The thresholds are calculated as the maximum of the features during the previous night.

Jin et al. [24] present PresenceSense, a zero-training algorithm for occupancy detection in office buildings based on plug loads. The algorithm uses a vague working schedule of the participant as initial estimate and iteratively refines the assessment through classifiers learning from the predictions in the previous iteration. As features they use the discrete power levels, the maximum absolute change, the mean of absolute difference, the mean of the length of changes, and the standard deviation. They state that a sampling interval of 1 min is sufficient, a relatively high rate compared to half an hour in our case.

Tang et al. present a framework named SHARK which requires no training. It models a household's appliances' states [43]. The mode states are decoded by solving an optimization problem. From the decoded state, the occupancy of the household is inferred. Although the algorithm needs no training, knowledge about the appliances within the house is necessary for the decoding step. Also the approach cannot be used online, since the optimisation process takes place over periods of time in the past.

A comparative study on occupancy prediction algorithms using electricity data can be found in [29]. Further occupancy prediction algorithms based on other principles among others are Preheat [39], Neurothermostat [33], or Presence Probabilities [30].

5 Conclusions

Our aim in this work was to advocate unsupervised classification algorithms for occupancy detection in private households which use the electricity consumption data measured by smart meters. One specific objective we approached was to examine the suitability of coarse-grained consumption data relative to fine-grained consumption data. We evaluated the performance of our algorithms on three datasets containing ground truth occupancy information. Among the algorithms we presented are also online algorithms, which are ready to be used in a real scenario. Besides, all algorithms we showed require little computational power and can easily be run inside the home. The best performing algorithms showed an accuracy of 69–90%, or an MCC of 0.20–0.78. In general we found that our unsupervised (i.e. zero-training) algorithms compare favourably to supervised algorithms and that the use of coarse-grained data is comparable to the use of fine-grained data to the performance in two out of three test cases.

Appendix: Algorithms

Algorithm GeoMA: t represents the timesteps, x is an array containing the electricity consumption values. Each entry represents a 30 min window of data. Accordingly, $schedule$ is an array containing the resulting occupancy estimations for each time slot. The parameter λ determines the importance of recent values over older values, i.e. $(1 - \lambda)$ determines the decay of the average.

```

function GEOMA( $x, \lambda$ )
  average  $\leftarrow x(1)$ 
  for all  $t$  do
    average  $\leftarrow \lambda * x(t) + (1 - \lambda) * average$ 
    if  $x(t) \geq average$  then
      schedule( $t$ )  $\leftarrow 1$ 
    else
      schedule( $t$ )  $\leftarrow 0$ 
    end if
  end for
  return schedule
end function

```

Algorithm PHT: t represents the timesteps, x is an array containing the electricity consumption values. Each entry represents a 30 min window of data. Accordingly, $schedule$ is an array containing the resulting occupancy estimations for each time slot.

```

function PHT( $x, magThreshold, detectThreshold$ )
  currentState  $\leftarrow 0$ 
  for all  $t$  do
    deviation  $\leftarrow x(t) - \bar{x} - magThreshold$ 
    mt  $\leftarrow mt + deviation$ 
    increasingMT  $\leftarrow \min(increasingMT, mt)$ 
    decreasingMT  $\leftarrow \max(decreasingMT, mt)$ 
    increasingPHT  $\leftarrow mt - increasingMT$ 
    decreasingPHT  $\leftarrow decreasingMT - mt$ 
    if increasingPHT > detectThreshold then
      schedule( $t$ )  $\leftarrow 1$ 
      currentState  $\leftarrow 1$ 
      mt  $\leftarrow 0$ 
    else if decreasingPHT > detectThreshold then
      schedule( $t$ )  $\leftarrow 0$ 
      currentState  $\leftarrow 0$ 
      mt  $\leftarrow 0$ 
    else
      schedule( $t$ )  $\leftarrow currentState$ 
    end if
  end for
  return schedule
end function

```

References

1. Akbar A, Nati M, Carrez F, Moessner K (2015) Contextual occupancy detection for smart office by pattern recognition of electricity consumption data. In: 2015 IEEE international conference on communications (ICC), pp 561–566. doi:[10.1109/ICC.2015.7248381](https://doi.org/10.1109/ICC.2015.7248381)
2. Alhamoud A, Xu P, Englert F, Reinhardt A, Scholl P, Boehnstedt D, Steinmetz R (2015) Extracting human behavior patterns from appliance-level power consumption data. In: Proceedings of the 12th European conference on wireless sensor networks, EWSN 2015, Porto, Portugal, pp 52–67. doi:[10.1007/978-3-319-15582-1_4](https://doi.org/10.1007/978-3-319-15582-1_4)
3. Amayri M, Arora A, Ploix S, Bandhyopadhyay S, Ngo QD, Badarla VR (2016) Estimating occupancy in heterogeneous sensor environment. Energy Build 129:46–58. doi:[10.1016/j.enbuild.2016.07.026](https://doi.org/10.1016/j.enbuild.2016.07.026)
4. Ardakanian O, Bhattacharya A, Culler D (2016) Non-intrusive techniques for establishing occupancy related energy savings in commercial buildings. In: Proceedings of the 3rd ACM international conference on systems for energy-efficient built environments, buildSys '16, pp 21–30. doi:[10.1145/2993422.2993574](https://doi.org/10.1145/2993422.2993574)
5. Barbato A, Borsani L, Capone A, Melzi S (2009) Home energy saving through a user profiling system based on wireless sensors. In: Proceedings of the 1st ACM workshop on embedded sensing systems for energy-efficiency in buildings, buildSys '09. ACM, New York, NY, USA, pp 49–54. doi:[10.1145/1810279.1810291](https://doi.org/10.1145/1810279.1810291)
6. Barker S, Mishra A, Irwin D, Cecchet E, Shenoy P, Albrecht J (2012) Smart*: an open data set and tools for enabling research in sustainable homes. In: Proceedings of the 2012 workshop on data mining applications in sustainability (SustKDD 2012), Beijing, China
7. Beckel C, Kleiminger W, Cicchetti R, Staake T, Santini S (2014) The ECO data set and the performance of non-intrusive load monitoring algorithms. In: Proceedings of the 1st ACM conference on embedded systems for energy-efficient buildings, buildsys '14. ACM, New York, NY, USA, pp 80–89. doi:[10.1145/2674061.2674064](https://doi.org/10.1145/2674061.2674064)
8. Boait PJ, Rylatt RM (2010) A method for fully automatic operation of domestic heating. Energy Build 42(1):11–16. doi:[10.1016/j.enbuild.2009.07.005](https://doi.org/10.1016/j.enbuild.2009.07.005)
9. Brackney LJ, Florita AR, Swindler AC, Polese LG, Brunemann GA (2012) Design and performance of an image processing occupancy sensor. In: The second international conference on building energy and environment 2012, Boulder, Colorado, USA, pp 987–994
10. Brown R, Ghavami N, Siddiqui HUR, Adjrard M, Ghavami M, Dudley S (2017) Occupancy based household energy disaggregation using ultra wideband radar and electrical signature profiles. Energy Build 141:134–141. doi:[10.1016/j.enbuild.2017.02.004](https://doi.org/10.1016/j.enbuild.2017.02.004)
11. Cali D, Matthes P, Huchtemann K, Streblov R, Müller D (2015) CO₂ based occupancy detection algorithm: experimental analysis and validation for office and residential buildings. Build Environ 86:39–49. doi:[10.1016/j.buildenv.2014.12.011](https://doi.org/10.1016/j.buildenv.2014.12.011)
12. Candanedo LM, Feldheim V (2016) Accurate occupancy detection of an office room from light, temperature, humidity and CO₂ measurements using statistical learning models. Energy Build 112:28–39. doi:[10.1016/j.enbuild.2015.11.071](https://doi.org/10.1016/j.enbuild.2015.11.071)
13. Chaney J, Owens EH, Peacock AD (2016) An evidence based approach to determining residential occupancy and its role in demand response management. Energy Build 125:254–266. doi:[10.1016/j.enbuild.2016.04.060](https://doi.org/10.1016/j.enbuild.2016.04.060)
14. Chen D, Barker S, Subbaswamy A, Irwin D, Shenoy P (2013) Non-intrusive occupancy monitoring using smart meters. In: Proceedings of the 5th ACM workshop on embedded systems for energy-efficient buildings, buildsys'13. ACM, New York, NY, USA, pp 9:1–9:8. doi:[10.1145/2528282.2528294](https://doi.org/10.1145/2528282.2528294)
15. Ebadat A, Bottegal G, Molinari M, Varagnolo D, Wahlberg B, Hjalmarsson H, Johansson KH (2015) Multi-room occupancy estimation through adaptive gray-box models. In: Proceedings of the 4th IEEE conference on decision and control (CDC), pp 3705–3711. doi:[10.1109/CDC.2015.7402794](https://doi.org/10.1109/CDC.2015.7402794)
16. Ebadat A, Bottegal G, Varagnolo D, Wahlberg B, Hjalmarsson H, Johansson KH (2015) Blind identification strategies for room occu-

- pancy estimation. In: 2015 European control conference (ECC), pp 1315–1320. doi:[10.1109/ECC.2015.7330720](https://doi.org/10.1109/ECC.2015.7330720)
17. Ebadat A, Bottegal G, Varagnolo D, Wahlberg B, Johansson KH (2015) Regularized deconvolution-based approaches for estimating room occupancies. *IEEE Trans Autom Sci Eng* 12(4):1157–1168. doi:[10.1109/TASE.2015.2471305](https://doi.org/10.1109/TASE.2015.2471305)
 18. European Commission (2014) Cost-benefit analyses & state of play of smart metering deployment in the EU-27. Technical Report 52014SC0189, EU Commission, Brussels. URL: <http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A52014SC0189>
 19. Froehlich J, Larson EC, Campbell T, Haggerty C, Fogarty J, Patel SN (2009) Hydrosense: infrastructure-mediated single-point sensing of whole-home water activity. In: Proceedings of the 11th international conference on Ubiquitous computing, UbiComp 2009, Orlando, Florida, USA, pp 235–244. doi:[10.1145/1620545.1620581](https://doi.org/10.1145/1620545.1620581)
 20. Gao G, Whitehouse K (2009) The self-programming thermostat: optimizing setback schedules based on home occupancy patterns. In: Proceedings of the 1st ACM workshop on embedded sensing systems for energy-efficiency in buildings, buildSys '09. ACM, New York, NY, USA, pp 67–72. doi:[10.1145/1810279.1810294](https://doi.org/10.1145/1810279.1810294)
 21. Guo X, Tiller D, Henze G, Waters C (2010) The performance of occupancy-based lighting control systems: a review. *Light Res Technol* 42(4):415–431. doi:[10.1177/1477153510376225](https://doi.org/10.1177/1477153510376225)
 22. Gupta M, Intille SS, Larson K (2009) Adding GPS-control to traditional thermostats: an exploration of potential energy savings and design challenges. In: Proceedings of the 7th international conference on pervasive computing, pervasive 2009, Nara, Japan, pp 95–114. doi:[10.1007/978-3-642-01516-8_8](https://doi.org/10.1007/978-3-642-01516-8_8)
 23. Jin M, Bekiaris-Liberis N, Weekly K, Spanos CJ, Bayen A M (2016) Occupancy detection via environmental sensing. *IEEE Trans Autom Sci Eng* 99:1–13. doi:[10.1109/TASE.2016.2619720](https://doi.org/10.1109/TASE.2016.2619720)
 24. Jin M, Jia R, Kang Z, Konstantakopoulos IC, Spanos CJ (2014) PresenceSense: zero-training algorithm for individual presence detection based on power monitoring. In: Proceedings of the 1st ACM conference on embedded systems for energy-efficient buildings, buildsys '14. ACM, New York, NY, USA, pp 1–10. doi:[10.1145/2674061.2674073](https://doi.org/10.1145/2674061.2674073)
 25. Kleiminger W (2015) Occupancy sensing and prediction for automated energy savings. Ph.D. thesis, ETH Zurich. doi:[10.3929/ethz-a-010450096](https://doi.org/10.3929/ethz-a-010450096)
 26. Kleiminger W, Beckel C, Dey AK, Santini S (2013) Using unlabeled Wi-Fi scan data to discover occupancy patterns of private households. In: Proceedings of the 11th ACM conference on embedded network sensor systems, sensys '13, Rome, Italy, pp 47:1–47:2. doi:[10.1145/2517351.2517421](https://doi.org/10.1145/2517351.2517421)
 27. Kleiminger W, Beckel C, Santini S (2015) Household occupancy monitoring using electricity meters. In: Proceedings of the 2015 ACM international joint conference on pervasive and Ubiquitous computing, UbiComp '15, pp 975–986. doi:[10.1145/2750858.2807538](https://doi.org/10.1145/2750858.2807538)
 28. Kleiminger W, Beckel C, Staake T, Santini S (2013) Occupancy detection from electricity consumption data. In: Proceedings of the 5th ACM workshop on embedded systems for energy-efficient buildings, buildsys'13. ACM, New York, NY, USA, pp 10:1–10:8. doi:[10.1145/2528282.2528295](https://doi.org/10.1145/2528282.2528295)
 29. Kleiminger W, Mattern F, Santini S (2014) Predicting household occupancy for smart heating control: a comparative performance analysis of state-of-the-art approaches. *Energy Build* 85:493–505. doi:[10.1016/j.enbuild.2014.09.046](https://doi.org/10.1016/j.enbuild.2014.09.046)
 30. Krumm J, Brush AJB (2011) Learning time-based presence probabilities. In: Proceedings of the 9th international conference on pervasive computing, pervasive 2011, San Francisco, CA, USA, pp 79–96. doi:[10.1007/978-3-642-21726-5_6](https://doi.org/10.1007/978-3-642-21726-5_6)
 31. Lu J, Sookoor T, Srinivasan V, Gao G, Holben B, Stankovic J, Field E, Whitehouse K (2010) The smart thermostat: using occupancy sensors to save energy in homes. In: Proceedings of the 8th ACM conference on embedded networked sensor systems, sensys '10. ACM, New York, NY, USA, pp 211–224. doi:[10.1145/1869983.1870005](https://doi.org/10.1145/1869983.1870005)
 32. Matthews BW (1975) Comparison predicted and observed secondary structure of T4 phage lysozyme. *Biochim Biophys Acta* 405(2):442–451. doi:[10.1016/0005-2795\(75\)90109-9](https://doi.org/10.1016/0005-2795(75)90109-9)
 33. Mozer M, Vidmar L, Dodier RH (1996) The neurothermostat: predictive optimal control of residential heating systems. In: Advances in neural information processing systems, vol 9, NIPS, Denver, CO, USA, pp 953–959
 34. Page E (1954) Continuous inspection schemes. *Biometrika* 41(1–2):100–115. doi:[10.1093/biomet/41.1-2.100](https://doi.org/10.1093/biomet/41.1-2.100)
 35. Patel SN, Reynolds MS, Abowd GD (2008) Detecting human movement by differential air pressure sensing in HVAC system ductwork: an exploration in infrastructure mediated sensing. In: Proceedings of the 6th international conference on pervasive computing, pervasive '08, pp 1–18. doi:[10.1007/978-3-540-79576-6_1](https://doi.org/10.1007/978-3-540-79576-6_1)
 36. Patel SN, Robertson T, Kientz JA, Reynolds MS, Abowd GD (2007) At the flick of a switch: detecting and classifying unique electrical events on the residential power line. In: Proceedings of the 9th international conference on Ubiquitous computing, UbiComp '07, Innsbruck, Austria, pp 271–288. doi:[10.1007/978-3-540-74853-3_16](https://doi.org/10.1007/978-3-540-74853-3_16)
 37. Pedersen TH, Nielsen KU, Petersen S (2017) Method for room occupancy detection based on trajectory of indoor climate sensor data. *Build Environ* 115:147–156. doi:[10.1016/j.buildenv.2017.01.023](https://doi.org/10.1016/j.buildenv.2017.01.023)
 38. Rabiner L, Juang B (1986) An introduction to hidden Markov models. *IEEE ASSP Mag*. doi:[10.1109/MASSP.1986.1165342](https://doi.org/10.1109/MASSP.1986.1165342)
 39. Scott J, Bernheim Brush A, Krumm J, Meyers B, Hazas M, Hodges S, Villar N (2011) Preheat: controlling home heating using occupancy prediction. In: Proceedings of the 13th international conference on Ubiquitous computing, UbiComp '11. ACM, New York, NY, USA, pp 281–290. doi:[10.1145/2030112.2030151](https://doi.org/10.1145/2030112.2030151)
 40. Smith JR, Fishkin KP, Jiang B, Mamishev A, Philipose M, Rea AD, Roy S, Sundara-Rajan K (2005) RFID-based techniques for human-activity detection. *Commun ACM* 48(9):39–44. doi:[10.1145/1081992.1082018](https://doi.org/10.1145/1081992.1082018)
 41. Soltanaghaei E, Whitehouse K (2016) Walksense: classifying home occupancy states using walkway sensing. In: Proceedings of the 3rd ACM international conference on systems for energy-efficient built environments, buildsys '16, pp 167–176. doi:[10.1145/2993422.2993576](https://doi.org/10.1145/2993422.2993576)
 42. Szczurek A, Maciejewska M (2016) Detection of occupancy events from indoor air monitoring data. In: 3rd international conference on mathematics and computers in sciences and in industry (MCSI), pp 229–234. doi:[10.1109/MCSI.2016.050](https://doi.org/10.1109/MCSI.2016.050)
 43. Tang G, Wu K, Lei J, Xiao W (2015) The meter tells you are at home! Non-intrusive occupancy detection via load curve data. In: 2015 IEEE international conference on smart grid communications (Smartgridcomm), Miami, FL, USA, pp 897–902. doi:[10.1109/SmartGridComm.2015.7436415](https://doi.org/10.1109/SmartGridComm.2015.7436415)
 44. Tapia EM, Intille SS, Larson K (2004) Activity recognition in the home using simple and ubiquitous sensors. In: Proceedings of the 2nd international conference on pervasive computing, pervasive '04, Vienna, Austria, pp 158–175. doi:[10.1007/978-3-540-24646-6_10](https://doi.org/10.1007/978-3-540-24646-6_10)
 45. Teixeira T, Savvides A (2008) Lightweight people counting and localizing for easily deployable indoors WSNs. *J Sel Top Signal Process* 2(4):493–502. doi:[10.1109/JSTSP.2008.2001426](https://doi.org/10.1109/JSTSP.2008.2001426)
 46. Telkonet. Telkonet products. URL: <http://www.telkonet.com> [cited 18.05.2017]
 47. van Kasteren T, Noulas AK, Englebienne G, Kröse BJA (2008) Accurate activity recognition in a home setting. In: Proceedings of

- the 10th international conference on Ubiquitous computing, UbiComp 2008, Seoul, Korea, pp 1–9. doi:[10.1145/1409635.1409637](https://doi.org/10.1145/1409635.1409637)
48. Viterbi AJ (1967) Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Trans Inf Theory* 13(2):260–269. doi:[10.1109/TIT.1967.1054010](https://doi.org/10.1109/TIT.1967.1054010)
 49. Wang S (1989) CO₂-based occupancy detection for on-line outdoor air flow. *Indoor Built Environ* 7(3):165–181. doi:[10.1177/1420326X9800700305](https://doi.org/10.1177/1420326X9800700305)
 50. Woodstock TK, Radke RJ, Sanderson AC (2016) Sensor fusion for occupancy detection and activity recognition using time-of-flight sensors. In: Proceedings of the 19th international conference on information fusion (FUSION), pp 1695–1701
 51. Yang L, Ting K, Srivastava MB (2014) Inferring occupancy from opportunistically available sensor data. In: Proceedings of the IEEE international conference on pervasive computing and communications, PerCom 2014, Los Alamitos, CA, USA, pp 60–68. doi:[10.1109/PerCom.2014.6813945](https://doi.org/10.1109/PerCom.2014.6813945)
 52. Zikos S, Tsolakis A, Meskos D, Tryferidis A, Tzovaras D (2016) Conditional random fields—based approach for real-time building occupancy estimation with multi-sensory networks. *Autom Constr* 68:128–145. doi:[10.1016/j.autcon.2016.05.005](https://doi.org/10.1016/j.autcon.2016.05.005)