

Privacy Implications of Room Climate Data*

Philipp Morgner¹, Christian Müller², Matthias Ring¹, Björn Eskofier¹,
Christian Riess¹, Frederik Armknecht², and Zinaida Benenson¹

¹ Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany
{philipp.morgner, matthias.ring, bjoern.eskofier,
christian.riess, zinaida.benenson}@fau.de

² University of Mannheim, Germany
{christian.mueller, armknecht}@uni-mannheim.de

Abstract. Smart heating applications promise to increase energy efficiency and comfort by collecting and processing room climate data. While it has been suspected that the sensed data may leak crucial personal information about the occupants, this belief has up until now not been supported by evidence.

In this work, we investigate privacy risks arising from the collection of room climate measurements. We assume that an attacker has access to the most basic measurements only: temperature and relative humidity. We train machine learning classifiers to predict the presence and actions of room occupants. On data that was collected at three different locations, we show that occupancy can be detected with up to 93.5% accuracy. Moreover, the four actions reading, working on a PC, standing, and walking, can be discriminated with up to 56.8% accuracy, which is also far better than guessing (25%). Constraining the set of actions allows to achieve even higher prediction rates. For example, we discriminate standing and walking occupants with 95.1% accuracy. Our results provide evidence that even the leakage of such ‘inconspicuous’ data as temperature and relative humidity can seriously violate privacy.

1 Introduction

The vision of the Internet of Things (IoT) is to enhance work processes, energy efficiency, and living comfort by interconnecting actuators, mobile devices and sensors. These networks of embedded technologies enable applications such as smart heating, home automation, and smart metering, among many others. Sensors are of crucial importance in these applications. Data gathered by sensors is used to represent the current state of the environment, for instance in smart heating, sensors measure the room climate. Using these information and a user-defined configuration of the targeted state of room climate, the application regulates heating, ventilation, and air conditioning.

While the collection of room climate data is obviously essential to enable smart heating, it may at the same time impose the risk of privacy violations. Consequently, it is commonly *believed* among security experts that leaking room climate data may result in privacy violations and hence that the data needs to be cryptographically protected [39,

* Authors’ version of the paper published in the Proceedings of the 22nd European Symposium on Research in Computer Security (ESORICS 2017). DOI: 10.1007/978-3-319-66399-9_18

12, 4]. However, these claims have not been supported by scientific *evidence* so far. Thus, one could question whether in practice additional effort for protecting the data would be justified.

The current situation with room climate data is comparable to the area of smart metering [18, 44, 24, 27]. In 1989, Hart [18] was the first to draw attention to the fact that smart metering appliances can be exploited as surveillance devices. Since then, research has shown far-reaching privacy violations through fine-granular power consumption monitoring, ranging from occupancy and everyday activities detection [31] up to recognizing which program a TV was displaying [14].

Various techniques have been proposed over the years to mitigate privacy risks of smart metering [3, 37, 25, 44, 36]. This issue has become such a grave concern that the German Federal Office for Information Security published a protection profile for smart meters in 2014 [2]. By considering privacy implications of smart heating, we hope to initiate consumer protection research and policy debate in this area, analogous to the developments in smart metering described above.

Research Questions. In this work, we are the first to investigate room climate data from the perspective of possible privacy violations. More precisely, we address the following research questions:

- *Occupancy detection:* Can an attacker determine the presence of a person in a room using only room climate data, i.e., temperature and relative humidity?
- *Activity recognition:* Can an attacker recognize activities of the occupant in the room using only the temperature and relative humidity data?

Our threat scenario targets buildings with multiple rooms that are similar in size, layout, furnishing, and positions of the sensors. These properties are typical for office buildings, dormitories, cruise ships, and hotels, among others. Assuming that an attacker is able to train a classifier that recognizes pre-defined activities, possible privacy violations are, e.g., tracking presence and working practices of employees in offices, or the disclosure of lifestyle and intimate activities in private spaces. All these situations present intrusions in the privacy of the occupants. In contrast to surveillance cameras and motion sensors, the occupant does not expect to be monitored. Also, legal restrictions regarding privacy might apply to surveillance cameras and motion sensors but not to room climate sensors.

Experiments. To evaluate these threats, we present experiments that consider occupancy detection and activity recognition based on the analysis of room climate data from a privacy perspective. We measured room climate data in three office-like rooms and distinguished between the activities reading, standing, walking, and working on a laptop. Although we assume that in smart heating applications, only one sensor per room is most likely to be installed, each room was equipped with several sensors in order to evaluate different positions of sensors in the room. These sensors measured temperature and relative humidity at a regular time interval of a few seconds. In our procedure, an occupant performed a pre-defined sequence of tasks in the experimental space. In sum, we collected almost 90 hours of room climate sensor data from a total

of 36 participants. The collected room climate data was analyzed using an off-the-shelf machine learning classification algorithm. To reflect realistic settings, we only evaluated data of a single sensor and did not apply sensor fusion.

Results. Evaluating our collected room climate data, the attacker detects presence of a person with detection rates up to 93.5% depending on location and the sensor position, which is significantly higher than guessing (50%). The attacker can distinguish between four activities (reading, standing, walking, and working on a laptop) with detection rates up to 56.8%, which is also significantly better than guessing (25%). We can also distinguish between three activities (sitting, standing and walking) with detection rates up to 81.0%, as opposed to 33.3% if guessing. Furthermore, we distinguish between standing and walking with detection rates up to 95.1%. Thus, we show that the fears of privacy violation by leaking room climate data are well justified. Furthermore, we analyze the influence of the room size, positions of the sensor, and amount of the measured sensor data on the accuracy. In summary, we provide the first steps in verifying the common belief that room climate data leaks privacy-sensitive information.

Outline. The remainder of this paper is organized as follows. In Section 2, we give an overview of related work. Section 3 presents the threat model considered in this work. In Section 4, we introduce the experimental design and methods. The results of our experiments are presented and discussed in Sections 5 and 6, respectively. We draw conclusions in Section 7. Additional information regarding the experimental procedure can be found in Appendix A.

2 Related Work

Over the last decade, several experiments have been conducted to detect occupancy in sensor-equipped spaces and to recognize people’s activities as summarized in Table 1. Activity recognition has been considered for basic activities, such as leaving or arriving at home, or sleeping [29], as well as for more detailed views, including toileting, showering and eating [41].

Most of the previous research uses types of sensors that are different from temperature and relative humidity. For example, CO₂ represents a useful source for occupancy detection and estimation [43]. Additionally, sensors detecting motion based on passive infrared (PIR) [1, 6, 15, 28, 17, 46], sound [11, 15], barometric pressure [30], and door switches [8, 9, 45] are utilized for occupancy estimation. For evaluation, different machine learning techniques are used, e. g., HMM [43], ARHMM [17], ANN [11], and decision trees [15, 45].

In contrast to previous work, our results rely exclusively on temperature and relative humidity. Previously published experimental results involved other or additional types of sensors, such as CO₂, acoustics, motion, or lighting (the latter three are referred to as AML in Table 1), door switches or states of appliances (also gathered with the help of switches), such as water taps or WC flushes. For this reason, our detection results are also not directly comparable to these works.

Table 1: Overview of previous experiments on occupancy detection (D), occupancy estimation (E), which aims at determining the number of people in a room, and activity recognition (A) with a focus on selected sensors; AML denotes acoustic, motion, and lighting sensors.

Work	Target	Rel. Humidity	Temperature	CO ₂	Ventilation	AML	Switches
van Kasteren et al., 2008 [41]	A	○	○	○	○	○	●
Lam et al., 2009 [28]	E	○	○	●	○	●	○
Dong et al., 2010 [6]	E	●	●	●	○	●	○
Lu et al., 2010 [29]	A	○	○	○	○	●	●
Hailemariam et al., '2011 [15]	D	○	○	○	○	●	○
Han et al., 2012 [17]	E	●	●	●	●	●	○
Zhang et al., 2012 [46]	E	●	●	●	○	●	○
Ekwevugbe et al., 2013 [11]	E	○	○	○	○	●	○
Ebadat et al., 2013 [8]	E	○	●	●	●	○	○
Ai et al., 2014 [1]	E	●	●	●	●	●	●
Wörner et al., 2014 [43]	D	○	○	●	○	○	○
Yang et al., 2014 [45]	D/E	●	●	●	○	●	●
Masood et al., 2015 [30]	E	○	●	●	○	○	○
Ebadat et al., 2015 [9]	E	○	●	●	●	○	●
This work	D/A	●	●	○	○	○	○

3 Threat Model

The overall goal of our work is to understand the potential privacy implications if room climate data is accessed by an attacker. The goal of the attacker is to gain information about the state of occupancy as well as the activity of the occupants without their consent.

Obviously, the more information an attacker can gather, the more likely she can deduce privacy-harming information from the measurements. Therefore, we base our analysis on the attacker model that considers a room climate system where only *one* sensor node is used to derive information. This is a realistic scenario since usually one sensor node per room is sufficient to monitor the room climate. Moreover, we assume that this sensor node takes only the two most basic measurements, temperature and relative humidity. These data are the fundamental properties to describe room climate. Note that our restricted data is in contrast to existing work (cf. Table 1 and Section 2) that based their experiments on more types of measurements or used data that is less common to characterize room climate.

We consider a sensor system that measures the climate of a room, denoted as *target location*. At the target location, a temperature and relative humidity sensor is installed that reports the measured values in regular intervals to a central database. We consider an attacker model where the attacker has access to this database and aims to derive information about the occupants at the target location. Furthermore, we assume that the attacker has access to either the target location itself, or a room similar in size, layout,

sensor positions, and furniture. Such situations are given, for example, at office buildings, hotels, cruise ships, and student dormitories. This location, denoted as *training location*, is used to train the classifier, which is a machine learning algorithm learning the input data labeled with the groundtruth. As the attacker has full control over the training location, she can freely choose what actions are taking place during the measurements. For example, she could do measurements while no persons are present at the training location, or one person is present and executes a predefined activity.

There are various scenarios, in which an attacker has incentives to collect and analyze room climate data. For example, the management of a company aims at observing the presence and working practices of employees in the offices. In another case, a provider of private spaces (hotels, dormitories, etc.) wants to disclose lifestyle and intimate activities in these spaces. This information may be utilized for targeted advertising or sold to insurance companies. In any case, the evaluation of room climate data provides the attacker with the possibility to undermine the privacy of the occupants.

The procedure of these attacks is as follows: First, the attacker collects training data at a training location, which might be the target location or another room similar in size, layout, sensor positions, and furniture. The attacker also records the groundtruth for all events that shall be distinguished. Examples of events are occupancy and non-occupancy, or different activities such as working, walking, and sleeping. The training data is recorded with a sample rate of a few seconds and split into *windows* (i.e., a temperature curve and a relative humidity curve) of same time lengths, usually one to three minutes. Using the collected training data, the attacker trains a machine learning classifier. After the classifier is trained, it can be used to classify windows of climate data from the target location to determine the events. The classifier works on previously collected data, thus reconstructing past events, and also on live-recorded data, thus determining current events “on-the-fly” at the target location.

4 Experimental Design and Methods

We conducted a study to investigate the feasibility of detecting occupancy and inferring activities in an office environment from temperature and relative humidity: From March to April 2016, we performed experiments at two locations simultaneously, *Location A* and *Location B*, with a distance of approximately 200 km between them. In addition, from January to February 2017, we conducted further experiments at a third location, denoted as *Location C*, which is located in the same building as Location B.

4.1 Experimental Setup and Tasks

The experimental spaces at the three locations are different in size, layout, and positions of the sensors. Thus, each target location is also the training location in our study. At Location A, the room has a floor area of 16.5 m² and was equipped with room climate sensors at four positions as shown in Figure 1ii. At Location B, the room has a floor area of 30.8 m², i. e., roughly twice as much as at Location A, and had room climate sensors installed at three positions as illustrated in Figure 1i. Location C has a floor area of 13.9 m² and was equipped with room climate sensors at five positions as shown

in Figure 1iii. In all locations, the room climate sensors measured temperature and relative humidity. The number of deployed sensors varied due to limitations of hardware availability.

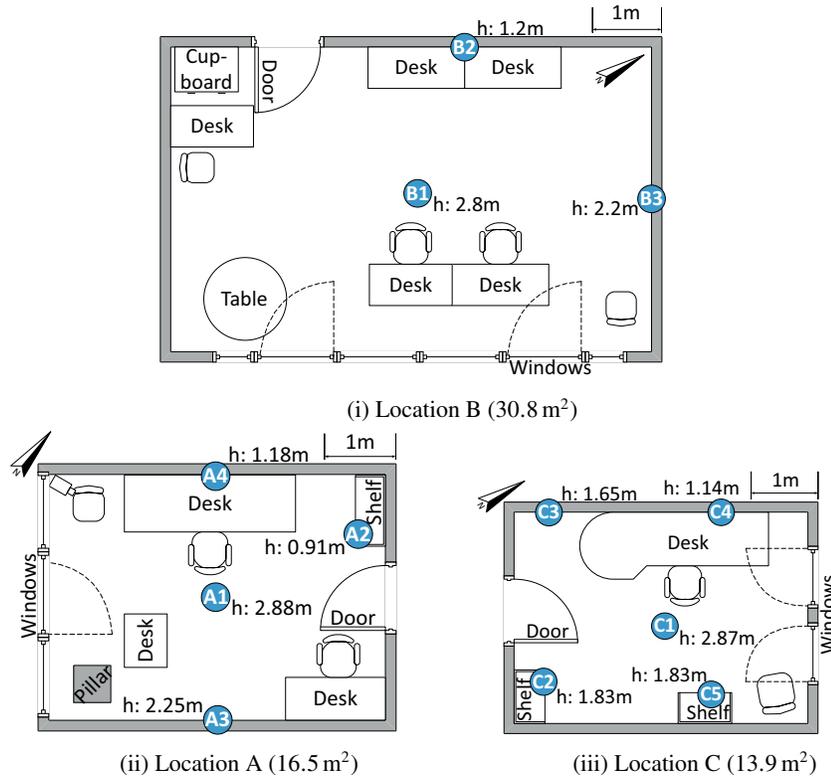


Fig. 1: Floor plans of the experiment spaces including sensor node locations, h indicates the node's height.

Our goal was to determine to which extent the presence and activities of an occupant influences the room climate data. Therefore, we measured temperature and relative humidity during phases of absence as well as phases of occupants' presence. If an occupant was present, this person had to perform one task or a sequence of tasks. We defined the following experimental tasks (see also Figure 2):

- Read** Sit on an office chair next to a desk and read.
- Stand** Stand in the middle of the room, try to avoid movements.
- Walk** Walk slowly and randomly through the room.
- Work** Sit on an office chair next to a desk and use a laptop, which is located on the desk.

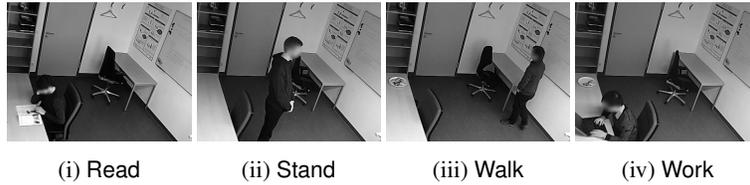


Fig. 2: The defined tasks performed by participants at Location A.

To eliminate confounding factors, we defined location default settings applying to all locations. Essentially, all windows were required to remain closed and no person was allowed in the room when not in use for the experiment. The rooms have radiators for heating, which were adjusted to a constant level. At Location A and B, we used shutters fixed in such positions that enough light was provided for reading and working.

4.2 Sensor Data Collection

We used a homogeneous hardware and software setup at all locations for data collection, which is described in the following.

Hardware. At each location, we set up a sensor network consisting of several Moteiv Tmote Sky sensor nodes with an integrated IEEE 802.15.4-compliant radio [32] as well as an integrated temperature and relative humidity sensor. The nodes have the Contiki operating system [7] version 2.7 installed. In addition, we deployed a webcam that took pictures in a 3-second interval at Location A. These were used for verification during the data collection phase only, and were not given to the classification algorithms.

Software. For sensor data collection, we customized the *Collect-View* application included in Contiki 2.7, which provides a graphical user interface to manage the sensor network. For our purposes, we implemented an additional control panel offering a customized logging system. The measurement settings of the *Collect-View* application were set to a report interval of 4 seconds with a variance of 1 second, i. e., each sensor node reported its current values in a time interval of 4 ± 1 seconds. The variance is a feature provided by *Collect-View* to decrease the risk of packet collisions during over-the-air transmissions.

Collected Data. We structured data collection in *units* and aimed for a good balance between presence and absence as well as the different tasks among all units, as this is needed for the later analysis using machine learning. Each unit has a fixed time duration, t , where exactly one person was present ($t \in \{10, 30, 60\}$, in minutes) who executed predefined activities. If the presence time was t minutes, then the absence time before and after it, respectively, was determined as $\frac{t}{2} + 5$ minutes, where 5 minutes served as buffer. This accounts for both, the equal distribution of presence time and absence time, respectively, and the fact that temperature and humidity settle within a 15-minute

period after the 60-minute presence of one person. For a detailed description of the experimental procedure, we refer to Appendix A.1.

Overall, we collected almost 90 hours of sensor data, 40 hours of which with a person being present. A more extensive overview of the amount of measured sensor data is shown in Table 2. To encourage replication and further investigations, all collected sensor data is available as open data sets on GitHub.³

Table 2: Measured sensor data of all locations (in hours)

Variable	Value	Recorded Time [h]		
		Location A	Location B	Location C
Occupancy	no	20:38:26	15:21:00	13:21:42
	yes	14:41:56	11:33:06	13:44:29
Task	Read	4:46:13	2:56:44	3:19:47
	Stand	2:45:27	2:34:20	3:28:27
	Walk	2:43:53	2:37:12	3:20:05
	Work	4:03:33	3:00:20	3:20:52

4.3 Participants and Ethical Principles

For participating in the experiment, 14 subjects volunteered at Location A, 12 subjects at Location B, and 10 subjects at Location C as shown in Table 3. Demographic data of participants was collected in order to facilitate replication and future experiments. All subjects provided written informed consent after the study protocol was approved by the data protection office.⁴ We assigned each participant to a random ID. All collected sensor data as well as the demographic data is only linked to this ID.

4.4 Classifier Design

We used classification to predict occupancy and activities in the rooms. We adopt an approach that has successfully been used in several applications of biosignal processing, namely extraction of a number of statistical descriptors with subsequent feature selection [26, 21].

The features use measurements from short time windows. We experimented with windows of different lengths, namely 60 s, 90 s, 120 s, 150 s, and 180 s. The offset between two consecutive windows was set to 30 s. We excluded all windows where only a part of the measurements belongs to the same activity.

The feature set was composed from a number of statistical descriptors that were computed on temperature and humidity measurements within these windows. These

³ <https://github.com/IoTsec/Room-Climate-Datasets>

⁴ Ethical review boards at both locations only consider medical experiments.

Characteristic	Location			
	A	B	C	
Gender	f:	3	2	5
	m:	11	10	5
Weight [kg]	μ :	74.9	81.7	63.1
	σ :	8.0	12.1	10.0
Height [cm]	μ :	175.9	178.4	170.7
	σ :	9.2	5.3	9.3
Age	μ :	33.7	30.3	25.6
	σ :	8.2	4.8	2.8

Table 3: Demographic data of participants, μ denotes the average, σ denotes the standard deviation.

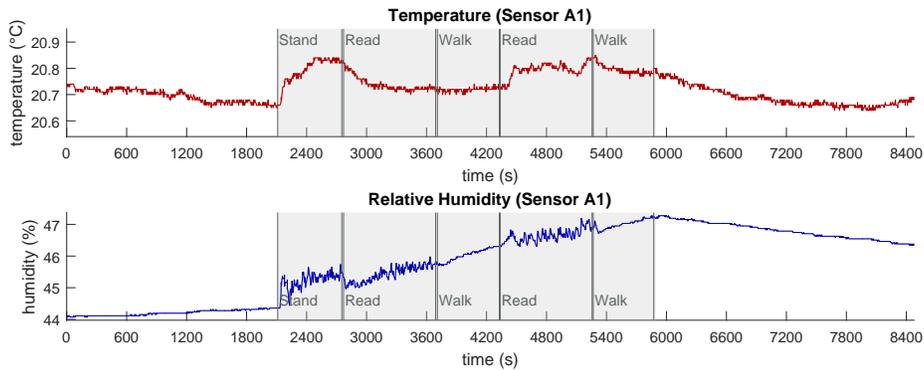
are mean value, variance, skewness, kurtosis, number of nonzero values, entropy, difference between maximum and minimum value of the window (i.e., value range), correlation between temperature and humidity, and mean and slope of the regression line for the measurement window before the current window. Additionally, we subtracted from the measurements their least-square linear regression line, and computed all of the listed statistics on the subtraction residuals. Feature selection was performed using a sequential forward search [42, Ch. 7.1 & 11.8], with an inner leave-one-subject-out cross-validation [19, Ch. 7] to determine the performance of each feature set. For classification, we used the Naïve Bayes classifier. To avoid a bias in the results, we randomly selected identical numbers of windows per class for training, validation and testing. For implementation, we used the ECST software [38], which wraps the WEKA library [16].

As performance measures, we use accuracy (i. e., the number of correctly classified windows divided by the number of all windows), and per-class sensitivity (i. e., the number of correctly classified windows for a specific class divided by the number of all windows of this class). Classification accuracy was deemed statistically significant if it was significantly higher than random guessing which is the best choice if the classifier could not learn any useful information during training. For each experiment, a binomial test with significance level $p < 0.01$ was carried out using the R software [34].

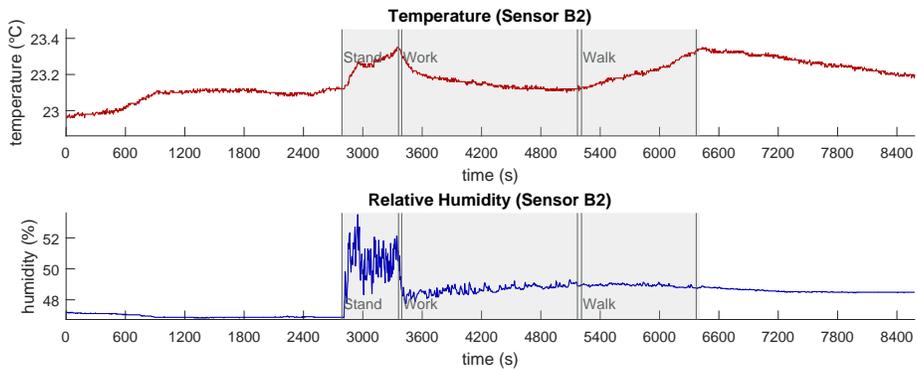
Note that neither the features nor the rather simple Naïve Bayes classifier are particularly tailored to predicting privacy leaks. However, we show that also such an unoptimized system is able to correctly predict occupancy and action types and hence produce privacy leaks. Higher detection rates results can be expected if more advanced classifiers are applied to this task.

5 Results

In this section, we present the experimental results. First, a visual inspection of the collected data is presented, followed by the machine learning-aided occupancy detection and activity recognition.



(i) Occupant is present for 60 minutes at Location A.



(ii) Occupant is present for 60 minutes at Location B.

Fig. 3: Visualization of two examples of room climate measurements. The grey background indicates the presence of the occupant in the experimental space.

5.1 Visual Inspection

We started our evaluation by analyzing the *raw* sensor data. Hence, we implemented a visualization script in MATLAB, which plots this data. The visualizations of two measurements are exemplarily depicted in Figure 3.

The visualizations show an immediate rise of the temperature and humidity as soon as an occupant enters the room. Furthermore, variations in temperature and humidity increase rapidly and can be clearly seen. Thus, one can visually distinguish between phases of occupancy and non-occupancy. One can also notice different patterns during the performance of the tasks. As Figure 3i shows, an occupant walking in the experimental space causes a constant increase of temperature and humidity with only small variations. In contrast, an occupant standing in the room causes the largest variations of humidity compared to the other defined tasks (cf. Figure 3ii). The effects of the tasks reading and working on temperature and humidity in the depicted figures are very sim-

ilar: both variables tend to increase showing medium variations. For further analysis of the data, we used machine learning as outlined in Section 4.4.

5.2 Occupancy Detection

Occupancy detection describes the binary detection of occupants in the experimental space based on features from windows with length of 180 seconds (cf. Section 4.4). This is a two-class task, namely to distinguish whether an occupant is present (true) or not (false). We only considered training and testing data within the same room (but separated training and testing both by the days and participants of the acquisition). We randomly selected the same number of positive and negative cases from the data. Thus, simply guessing the state has a success probability of 50%. However, our classification results are considerably higher than that. Table 4 shows that the highest accuracies per location were 93.5% (Location A), 88.5% (Location B), and 91.0% (Location C). Considering all sensors of all three locations, detection accuracy ranges between 66.8% (Sensor B3) and 93.5% (Sensor A1) as shown in Figure 4i. All classification accuracies were statistically significantly different from random guessing. This indicates that an attacker can reveal the presence of occupants in a target location with a high probability.

Scenario	Sensor	Sensitivity [%]		Guess [%]	Acc. [%]
		Occup.	No Occup.		
<i>Occupancy</i>	A1	94.1	93.0	50.0	93.5
	A2	94.5	85.0	50.0	89.7
	A3	92.0	76.4	50.0	84.2
	A4	77.8	79.1	50.0	78.4
	B1	91.9	85.1	50.0	88.5
	B2	85.3	77.2	50.0	81.3
	B3	69.7	63.9	50.0	66.8
	C1	92.9	89.2	50.0	91.0
	C2	89.9	87.4	50.0	88.6
	C3	90.0	82.0	50.0	86.0
	C4	89.8	87.6	50.0	88.7
	C5	92.5	88.8	50.0	90.7

Table 4: Classification accuracy for occupancy detection. Notations: ‘Occup.’, sensitivity for class *occupancy*. ‘No Occup.’, sensitivity for class *no occupancy*. ‘Guess’, probability of correct guessing. ‘Acc.’, classification accuracy.

5.3 Activity Recognition

Activity recognition reports the current activity of an occupant in the experimental space. The four activity tasks are described in Section 4.1. The recognition results for these tasks are shown in Figure 4.

Activity4 classifies between the activities Read, Stand, Walk, Work. As shown in Figure 4ii, the accuracy of recognizing activities achieved by the machine learning

pipeline ranged from 23.9% (Sensor C1) to 56.8% (Sensor A1). Overall, the accuracy of *Activity4* was statistically significantly better than the probability of guessing the correct task (25%) for 8 out of 12 sensors. Thus, the distinction between multiple activities is possible, but depends on the target location and the position of the sensor.

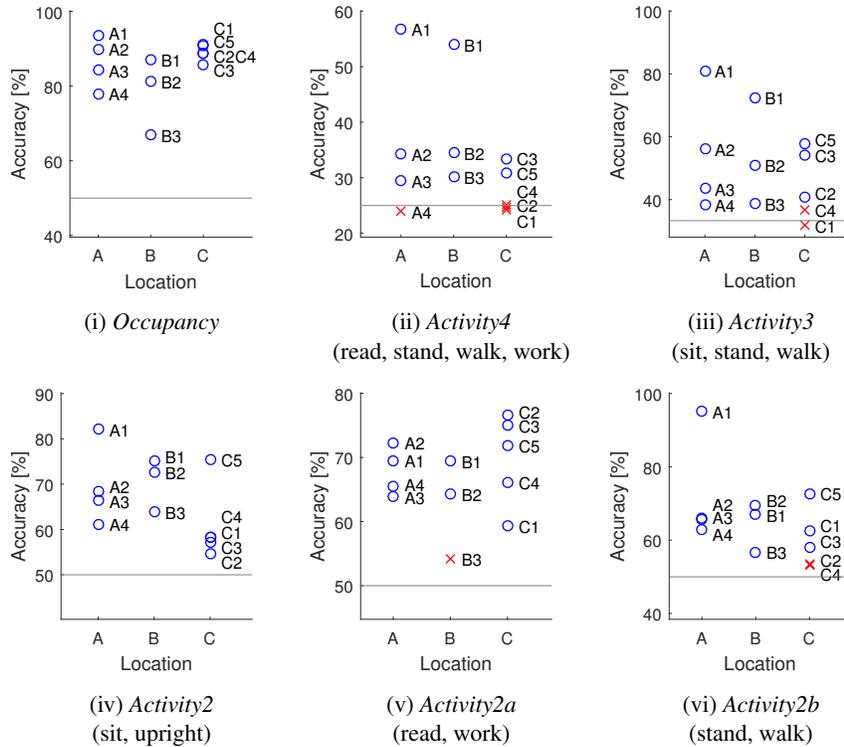


Fig. 4: Classification accuracy for occupancy detection and activity recognition. In each diagram, the guessing probability is plotted as a line. Each symbol represents the accuracy that we achieved with a single sensor. A blue dot marks a statistically significant result, while a red 'x' represents a statistically insignificant result.

In the next step, we investigated whether an attacker can increase the recognition accuracies by distinguishing between a smaller set of activities. To this end, we combined two tasks to a *meta task*, e.g., the tasks Read and Work became Sit. The model *Activity3* classifies between the tasks Sit, Stand, and Walk. The probability of correct guessing is thus 33.3%. This model is typical to represent activities of an occupant in a private space or an office room. For *Activity3*, the achieved accuracy ranged from 31.8% (Sensor C1) to 81.0% (Sensor A1). Our results were statistically significant for 10 out of the 12 sensors deployed in the three locations. Assuming a known layout of

the target location, the attacker might be able to determine the position of the occupant in the space and infer activities such as watching TV, exercising, cooking or eating.

The model *Activity2* classifies between the tasks Sit and Upright, whereby Sit is as previously Read or Work, and Upright combines Stand and Walk. In this classification, the attacker distinguishes whether an occupant is at a certain posture. The model *Activity2a* classifies between the tasks Read and Work, and the model *Activity2b* classifies between the tasks Stand and Walk. *Activity2a* indicates that an attacker can even distinguish between the sedentary activities, such as reading a book or working on the laptop. In contrast, *Activity2b* shows that an attacker can differentiate between standing and moving activities. Thus, an attacker can detect movements at the target location. For *Activity2*, *Activity2a*, and *Activity2b*, the probability to guess the correct class is 50%. Using these models, the attacker can infer various work and life habits.

For *Activity2*, our accuracy varies between 54.6% (Sensor C2) and 82.1% (Sensor A1), and all accuracies are statistically significant. For *Activity2a*, the lowest and highest accuracies were 54.2% (Sensor B3) and 76.6% (Sensor C2), respectively, which resulted in statistically significant results for 11 out of 12 sensors. For *Activity2b*, the achieved accuracy ranged from 53.3% (Sensor C4) to 95.1% (Sensor A1) and the results for 10 out of 12 sensors were statistically significant.

5.4 Further Observations

Length of Measurement Windows. The length of the measurement windows influences the accuracy of detection. We evaluated window sizes in the range between 60 and 180 seconds. Exemplarily, we analyzed the average accuracy of occupancy detection depending on the window size for all three locations. As shown in Figure 5, the accuracy increases with a longer window size. We achieved the best results with the longest window sizes of 180 seconds.

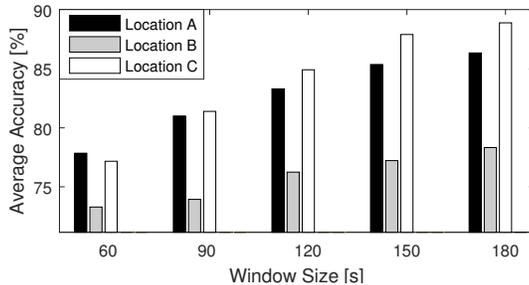


Fig. 5: Average accuracy over all sensors from each location for occupancy detection depending on the window size

This indicates that the highest accuracies are possible if longer time periods are considered. From a practical perspective, it is not advisable to extend the window size to a much larger duration than a few minutes since we assume that the performed activity is consistent for the whole duration of the window.

Selected Features. To assess the feasibility of an attacker that has only access to either temperature data or relative humidity data, we evaluated whether it might be enough to solely collect one type of room climate data. In the classification process, an attacker derives a set of features from temperature and relative humidity data and selects the best-performing features for each sensor and classification goal automatically (cf. Section 4.4). Analysis shows that features computed from temperature and relative humidity are of similar importance. In our evaluation, 57.9% of the selected features are derived from temperature measurements, and 52.3% from relative humidity measurements.⁵

We also compared the features in terms of differences between the three locations as well as differences between occupancy detection and activity recognition. In all these cases, there are no significant differences between the importance of temperature and relative humidity. An attacker restricted to either temperature or relative humidity data will perform worse than with both data.

Size and Layout of Rooms. All our locations are office-like rooms, which have a similar layout (rectangular) but differ in size and furnishing. In our evaluation, the accuracy correlates with the size of the target location. As shown in Figure 5, we had the highest average accuracy in occupancy detection with Location C, which has also the smallest ground area of 13.9m². Location A has a ground area of 16.5m², and has a slightly lower average accuracy. Location B is almost twice as large (30.8m²) and shows the worst average accuracy compared to the other locations. Thus, our experiment indicates that an increasing room size leads to decreasing accuracy on average. An attacker achieves higher accuracies by monitoring target locations of a small size compared to target locations of larger sizes.

Position of Sensors. According to our threat model in Section 3, the attacker controls layout of the target location. Thus, we assume an attacker that can decide at which position in the target location a room climate sensor is installed. We consider how the position of a room climate sensor influences the accuracy of derived information. For occupancy detection, we had the best accuracy with a sensor node that is located in the center point at the ceiling of the target location (Sensors A1, B1, C1). In this position, the sensor has the largest gathering area to measure the climate of the room. Sensors mounted to the walls or on shelves perform differently in our experiments. For activity recognition, the central sensor nodes performed best at Location A and B, but not at Location C.

From the attacker perspective, the best position to deploy a room climate sensor is at the ceiling in the center of the target location. In large rooms, multiple sensors at the ceiling could be installed, each covering a subsection of the room.

⁵ Note that some features are based on both, temperature and relative humidity, which is why the sum of both numbers exceeds 100%.

6 Discussion

As our experiments reveal, knowing the temperature and relative humidity of a room allows to detect the presence of people and to recognize certain activities with a significantly higher probability than guessing. By evaluating temperature and relative humidity curves of the length of 180 seconds, we were able to detect the presence of an occupant in one of our experimental spaces with an accuracy of 93.5% using a single sensor. In terms of activity recognition, we distinguished between four activities with an accuracy up to 56.8%, between three activities up to 81.0%, and between two activities up to 95.1%. Thus, an attacker focusing on the detection of a specific activity is more successful than an attacker that aims to classify a broader variety of activities. In the following, we discuss implications and limitations of our results.

Privacy Implications We show that an attacker might be able to infer life and work habits of the occupants from the room climate data. Thus, the attacker is able to distinguish between sitting, standing, and moving, which already might reveal the position and activities of the occupant in the room. Moreover, the attacker can distinguish between upright and sedentary activities, between moving and standing, and between working on the laptop or reading a book.

Given the limited amount of recorded sensor data, the achieved accuracies in occupancy detection and activity recognition give a clear indication that occupants are subject to privacy violations according to the threat model described in Section 3. However, activity recognition is not straightforward since the achieved accuracies differ between the different sensor positions and locations.

Further experiments are required for a better assessment of the privacy risks induced by the room climate data. Our work provides promising directions for these assessments. For example, we demonstrated the existence of the information leak with the Naïve Bayes classifier. Naïve Bayes is arguably one of the simplest machine learning classifiers. In future work, it would be interesting to explore upper boundaries for the detection of presence/absence and different activities by using more advanced classifiers such as the recently popular deep learning algorithms.

Location-Independent Classification An important question is whether it is possible to perform location-independent classification, i.e., to train the classifier with sensor data of one location and then use it to classify sensor data at the target location that is not similar to the training location in size, layout, and sensor positions. If this was possible, the service providers of smart heating applications would be able to detect occupancy and to recognize activities without having access to the target locations.

According to their privacy statements, popular smart thermostats from Nest [33], Ecobee [10], and Honeywell [20] send measured climate data to the service providers' databases. To evaluate these privacy threats, we used the room climate data of the best-performing sensor of a location as training data set for other locations. For example, to classify events of an arbitrary sensor of Location A, we trained the classifier with room climate data collected by Sensor B1 or Sensor C1. We gained statistically significant results for a few combinations in occupancy detection but the majority of our occupancy

detection results was not significant. For activity recognition, we were not able to gain statistically significant results.

However, the possibility of location-independent attackers cannot be excluded. Absence of significant results in our experiments may be merely due to the limited amount of data. Future studies should be conducted to gather data from various rooms up to a point where the combined results hold for arbitrary locations. Having more data from a multitude of rooms available would help the machine learning classifiers to recognize and ignore data characteristics that are specific to either of the experimental rooms. Consequently, the algorithms could better identify the distinct data characteristics of the different classes in occupancy detection and activity recognition. This would enable location-independent classification of room climate data, in which the training location is not similar to the target location regarding size, layout, furnishing, and positions of the sensors.

In a representative smart home survey of German consumers from 2015, 34% of the participants stated that they are interested in technologies for intelligent heating or are planning to acquire such a system [5]. Another survey with 1,000 US and 600 Canadian consumers found that for 72% of them, the most desired smart home device would be a self-adjusting thermostat, and 37% reported that they were likely to purchase one in the next 12 months [22]. Sharing smart home data with providers and third parties is a popular idea and a controversial issue for consumers. Thus, in a recent representative survey with 461 American adults by Pew Research [35], the participants were presented with a scenario of installing a smart thermostat *“in return for sharing data about some of the basic activities that take place in your house like when people are there and when they move from room to room”*. Of all respondents, 55% said that this scenario was not acceptable for them, 27% said that it was acceptable, with remaining 17% answering *“it depends”*. Furthermore, in a worldwide survey with 9,000 respondents from nine countries (Australia, Brazil, Canada, France, Germany, India, Mexico, the UK, and the US), 54% of respondents said that *“they might be willing to share their personal data collected from their smart home with companies in exchange for money”* [23].⁶

We think that the idea of sharing the smart home data for various benefits will continue to be intensively discussed in the future, and therefore, consumers and policy makers should be made aware of the level of detail inferable from smart home data. Which rewards are actually beneficial for consumers? Moreover, which kind of data sharing is ethically permissible? Only by answering these questions it would be possible to design fair policies and establish beneficial personal data markets [40]. In this work, we take the first step towards informing the policy for the smart heating scenario.

7 Conclusions

We investigated the common belief that the data collected by room climate sensors divulge private information about the occupants. To this end, we conducted experiments that reflect realistic conditions, i.e., considering an attacker who has access to typical room climate data (temperature and relative humidity) only. Our experiments revealed

⁶ Methodological details, such as representativeness, breakdown by country and the exact formulation of the questions, are not known about this survey.

that knowing a sequence of temperature and relative humidity measurements already allows to detect the presence of people and to recognize certain activities with high accuracy. Our results confirm that the assumptions that room climate data needs protection are justified: the leakage of such ‘inconspicuous’ sensor data as temperature and relative humidity can seriously violate privacy in smart spaces. Future work is required determine the level of privacy invasion in more depth and develop appropriate countermeasures.

Acknowledgement

The work is supported by the German Research Foundation (DFG) under Grant AR 671/3-1: WSNSec – Developing and Applying a Comprehensive Security Framework for Sensor Networks.

References

1. B. Ai, Z. Fan, and R. X. Gao. Occupancy estimation for smart buildings by an auto-regressive hidden Markov model. In *American Control Conference, ACC 2014, Portland, OR, USA, June 4-6, 2014*, pages 2234–2239. IEEE, 2014.
2. BSI. Protection Profile for the Gateway of a Smart Metering System (Smart Meter Gateway PP). <https://www.commoncriteriaportal.org/files/ppfiles/pp0073b.pdf>, Mar. 2014.
3. A. Cavoukian, J. Polonetsky, and C. Wolf. SmartPrivacy for the Smart Grid: embedding privacy into the design of electricity conservation. *Identity in the Information Society*, 3(2):275–294, 2010.
4. Chaos Computer Club: Guidelines for Smart Home Solutions. <https://www.ccc.de/en/updates/2016/smarthome>. Feb. 2016 (in German).
5. Deloitte. Ready for Takeoff? Consumer Survey, July 2015.
6. B. Dong, B. Andrews, K. P. Lam, M. Höynck, R. Zhang, Y.-S. Chiou, and D. Benitez. An information technology enabled sustainability test-bed (ITEST) for occupancy detection through an environmental sensing network. *Energy and Buildings*, 42(7):1038 – 1046, 2010.
7. A. Dunkels, B. Grönvall, and T. Voigt. Contiki – a lightweight and flexible operating system for tiny networked sensors. In *29th Annual IEEE International Conference on Local Computer Networks, 2004*, pages 455–462. IEEE, 2004.
8. A. Ebadat, G. Bottegal, D. Varagnolo, B. Wahlberg, and K. H. Johansson. Estimation of building occupancy levels through environmental signals deconvolution. In *BuildSys 2013, Proceedings of the 5th ACM Workshop On Embedded Systems For Energy-Efficient Buildings, Roma, Italy, November 13-14, 2013*, pages 8:1–8:8, 2013.
9. A. Ebadat, G. Bottegal, D. Varagnolo, B. Wahlberg, and K. H. Johansson. Regularized deconvolution-based approaches for estimating room occupancies. *IEEE Trans. Automation Science and Engineering*, 12(4):1157–1168, 2015.
10. Ecobee. Privacy policy & terms of use, April 2015.
11. T. Ekwevugbe, N. Brown, V. Pakka, and D. Fan. Real-time building occupancy sensing using neural-network based sensor network. In *7th IEEE International Conference on Digital Ecosystems and Technologies (DEST), 2013*, pages 114–119, July 2013.
12. European Union Agency For Network And Information Security. Security and Resilience of Smart Home Environments – Good practices and recommendations. <https://www.enisa.europa.eu>. Dec. 2015.

13. S. Fischer-Hübner and N. Hopper, editors. *Privacy Enhancing Technologies - 11th International Symposium, PETS 2011, Waterloo, ON, Canada, July 27-29, 2011. Proceedings*, volume 6794 of *Lecture Notes in Computer Science*. Springer, 2011.
14. U. Greveler, P. Glösekötterz, B. Justusy, and D. Loehr. Multimedia content identification through smart meter power usage profiles. In *Proceedings of the International Conference on Information and Knowledge Engineering (IKE)*, 2012.
15. E. Hailemariam, R. Goldstein, R. Attar, and A. Khan. Real-time occupancy detection using decision trees with multiple sensor types. In *2011 Spring Simulation Multi-conference, SpringSim '11, Boston, MA, USA, April 03-07, 2011.*, pages 141–148, 2011.
16. M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The WEKA Data Mining Software: An Update. *SIGKDD Explor. Newsl.*, 11(1):10–18, Nov. 2009.
17. Z. Han, R. X. Gao, and Z. Fan. Occupancy and indoor environment quality sensing for smart buildings. In *2012 IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*, pages 882–887, May 2012.
18. G. W. Hart. Residential energy monitoring and computerized surveillance via utility power flows. *Technology and Society Magazine, IEEE*, 8(2):12–16, 1989.
19. T. Hastie, R. Tibshirani, and J. H. Friedman. *The Elements of Statistical Learning*. Springer, New York, NY, USA, 2nd edition, 2009.
20. Honeywell. Honeywell connected home privacy statement, December 2015.
21. V. Huppert, J. Paulus, U. Paulsen, M. Burkart, B. Wullich, and B. Eskofier. Quantification of Nighttime Micturition With an Ambulatory Sensor-Based System. *IEEE Journal of Biomedical and Health Informatics*, 20(3):865–872, May 2016.
22. iconrol Networks: 2015 State of the Smart Home Report. <https://www.iconrol.com/blog/2015-state-of-the-smart-home-report>.
23. Intel Security: Intel Security’s International Internet of Things Smart Home Survey Shows Many Respondents Sharing Personal Data for Money. <https://newsroom.intel.com/news-releases/intel-securitys-international-internet-of-things-smart-home-survey>. Mar. 2016.
24. M. Jawurek, M. Johns, and F. Kerschbaum. Plug-in privacy for smart metering billing. In Fischer-Hübner and Hopper [13], pages 192–210.
25. M. Jawurek, F. Kerschbaum, and G. Danezis. SoK: Privacy technologies for smart grids – a survey of options. *Microsoft Res., Cambridge, UK*, 2012.
26. U. Jensen, P. Blank, P. Kugler, and B. Eskofier. Unobtrusive and Energy-Efficient Swimming Exercise Tracking Using On-Node Processing. *IEEE Sensors Journal*, 16(10):3972–3980, May 2016.
27. K. Kursawe, G. Danezis, and M. Kohlweiss. Privacy-friendly aggregation for the smart-grid. In Fischer-Hübner and Hopper [13], pages 175–191.
28. K. P. Lam, M. Höynck, B. Dong, B. Andrews, Y. shang Chiou, D. Benitez, and J. Choi. Occupancy detection through an extensive environmental sensor network in an open-plan office building. In *Proc. of Building Simulation 09, an IBPSA Conference*, 2009.
29. 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 *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, pages 211–224. ACM, 2010.
30. M. K. Masood, Y. C. Soh, and V. W. Chang. Real-time occupancy estimation using environmental parameters. In *2015 International Joint Conference on Neural Networks, IJCNN 2015, Killarney, Ireland, July 12-17, 2015*, pages 1–8. IEEE, 2015.
31. A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin. Private memoirs of a smart meter. In *Proceedings of the 2Nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building, BuildSys '10*, pages 61–66, New York, NY, USA, 2010. ACM.
32. Moteiv Corporation. Tmote Sky Datasheet, 2006.

33. Nest. Privacy statement for nest products and services, March 2016.
34. R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2014.
35. L. Rainie and M. Duggan. Pew Research: Privacy and Information Sharing. <http://www.pewinternet.org/2016/01/14/privacy-and-information-sharing>. Jan. 2016.
36. A. Reinhardt, F. Englert, and D. Christin. Averting the privacy risks of smart metering by local data preprocessing. *Pervasive and Mobile Computing*, 16:171–183, 2015.
37. A. Rial and G. Danezis. Privacy-preserving smart metering. In *Proceedings of the 10th Annual ACM Workshop on Privacy in the Electronic Society, WPES '11*, pages 49–60, New York, NY, USA, 2011. ACM.
38. M. Ring, U. Jensen, P. Kugler, and B. Eskofier. Software-based performance and complexity analysis for the design of embedded classification systems. In *Proceedings of the 21st International Conference on Pattern Recognition, ICPR 2012, Tsukuba, Japan, November 11-15, 2012*, pages 2266–2269. IEEE Computer Society, 2012.
39. M. Selinger. Test: Smart Home Kits Leave the Door Wide Open – for Everyone. <https://www.av-test.org/en/news/news-single-view/test-smart-home-kits-leave-the-door-wide-open-for-everyone/>. Apr. 2014.
40. S. Spiekermann, A. Acquisti, R. Böhme, and K.-L. Hui. The challenges of personal data markets and privacy. *Electronic Markets*, 25(2):161–167, 2015.
41. T. van Kasteren, A. Noulas, G. Englebienne, and B. Kröse. Accurate activity recognition in a home setting. In *Proceedings of the 10th International Conference on Ubiquitous Computing*. ACM, 2008.
42. I. H. Witten, E. Frank, and M. A. Hall. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann, Burlington, MA, USA, 3rd edition, 2011.
43. D. Wörner, T. von Bomhard, M. Roeschlin, and F. Wortmann. Look twice: Uncover hidden information in room climate sensor data. In *4th International Conference on the Internet of Things, IoT 2014, Cambridge, MA, USA, October 6-8, 2014*, pages 25–30. IEEE, 2014.
44. W. Yang, N. Li, Y. Qi, W. Qardaji, S. McLaughlin, and P. McDaniel. Minimizing private data disclosures in the smart grid. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, pages 415–427. ACM, 2012.
45. Z. Yang, N. Li, B. Becerik-Gerber, and M. D. Orosz. A systematic approach to occupancy modeling in ambient sensor-rich buildings. *Simulation*, 90(8):960–977, 2014.
46. R. Zhang, K. P. Lam, Y.-S. Chiou, and B. Dong. Information-theoretic environment features selection for occupancy detection in open office spaces. *Building Simulation*, 5(2):179–188, 2012.

A Additional Material

A.1 Experimental Procedure

The participants were assigned to at least one experimental unit with fixed presence times and tasks, and provided with a script for their actions (that is, for how long and in which order the tasks should be performed). Every participant performed each unit twice, with the same tasks, but possibly on different days and in a permuted chronological order. Tasks were performed in blocks of 10, 20, or 30 minutes. Thus, 10-minute units contained only one task of 10 minutes; 30-minute units consisted of either three tasks of 10 or one task of 10 plus one of 20 minutes; 60-minute units were composed of either two tasks of 20 plus two of 10, or one task of 10, 20, and 30 minutes each.

At the beginning of the presence time for each unit, i.e., the time period where a person had to be present, the experimental supervisor unlocked the room door to let the participant in. The participant started with the first task and was instructed by phone (at Locations A and C) or through the glass pane (at Location B) when it was time to change activities or to leave the room.

Overall, we defined 22 units per location, consisting of six 60-minute plus eight 30-minute and eight 10-minute units. Furthermore, the distribution of units and tasks was identical for all locations. Both, **Read** and **Work** account for 180 minutes each, whereas **Stand** and **Walk** provide 160 minutes each. A comprehensive overview of the distribution of tasks, number of tasks (per unit and block), and aggregated values is provided by Table 5.

Table 5: Overview of the number and distribution of tasks and units at one location. $n \times t$ denotes the number of n recorded t -units (i. e., the time of presence in minutes, $t \in \{10, 30, 60\}$), t_{task} denotes the defined task block lengths per unit. For instance, in a total of six 60-minute units, **Read** and **Work** account for two 30-minute blocks, whereas in a total of eight 10-minute units, all tasks account for two blocks of 10 minutes each.

Units		Tasks			
$n \times t$	t_{task}	Read	Stand	Walk	Work
	30	2	–	–	2
6 x 60	20	2	–	4	2
	10	–	8	–	–
8 x 30	20	2	–	2	2
	10	2	6	2	2
8 x 10	10	2	2	2	2
Total time [min]		180	160	160	180