How to Log Sleeping Trends? A Case Study on the Long-Term Capturing of User Data

Holger Becker, Marko Borazio and Kristof Van Laerhoven

Embedded Sensing Systems, TU Darmstadt,
Hochschulstr. 10, 64289, Darmstadt Germany
{h.becker,borazio,kristof}@ess.tu-darmstadt.de
http://www.ess.tu-darmstadt.de

Abstract. Designing and installing long-term monitoring equipment in the users' home sphere often presents challenges in terms of reliability, privacy, and deployment. Taking the logging of sleeping postures as an example, this study examines data from two very different modalities, high-fidelity video footage and logged wrist acceleration, that were chosen for their ease of setting up and deployability for a sustained period. An analysis shows the deployment challenges of both, as well as what can be achieved in terms of detection accuracy and privacy. Finally, we evaluate the benefits that a combination of both modalities would bring.

Keywords: actigraphy, sleep postures, long-term monitoring

1 Introduction

Long-term monitoring of users at home has been the focus of many recent studies. The applications that are argued for in research often tend to be medical in nature, such as the support for elderly care with systems that observe activity, fitness, and distinct events like falls. Other work targets a wider audience with context-aware home environments that display and activate services in appropriate situations, such as ubiquitous displays [8] or smart appliances [10].

Sleep posture logging and analysis is part of an emergent field in sleep monitoring research, with patients in professional sleeping laboratories being monitored by infrared cameras and a chest-worn tilt measurement unit. Sleep postures has been identified as a key factor in a variety of scenarios, including personality evaluation [2] and obstructive sleep apnea [9]. This scenario will be used here as a case study for long-term domestic monitoring systems, where the sleeping postures of the user will be detected by a system that remains running for a long period, from weeks to years of continuous operation.

The emphasis in this paper is put on how the modality of the sensed data affects the performance of a *long-running* system that is supposed to monitor the user in her *domestic* environment. This set of broad requirements tends to bring in several challenges that researchers are facing when deploying such systems.

1.1 Challenges

We argue that three challenges in particular are important to observe in long-term deployments in domestic environments. These three will be used to frame our evaluation and discussion in the remainder of the paper:

Reliability is in general a concern for most monitoring systems; for those that are supposed to run over longer time spans (from months to even years), this is even more of a challenge. The system not only has to keep running, it also has to perform and keep an accurate detection of the phenomena to be monitored. The presence of noise and irrelevant data, as well as limitations on the sensor's behalf, mean that a perfect capturing of the phenomena is often impossible to attain out in real-world deployments. This for instance will become clear in our sleep case study, where direct video footage is often covered and body-worn sensors cannot be worn in the most optimal places (chest or back).

Privacy is another factor that becomes pressing as larger sections of our lives are recorded. As the sensors can pick up more detail and cover more in both time and space, their data become more sensitive and should be safeguarded, and designing the system should very much take this into account. For the case study of sleep posture logging, the sensor could automatically record and reveal any activities related to the user's bedroom environment, which is together with the bathroom a most sensitive place in the home [5].

Deployment of the monitoring system needs to be easy and modular, so little time is spent on installation and moving the system from one environment to another does not bring along a costly installation procedure. Dependability and usability of the system also belong in this category (e.g., not having to frequently reboot or maintain the system, or avoiding failures in critical medical applications). In our case study, the recording system needs to remain active for at least several hours per day, and this without intervening much in the user's daily life (as frequent battery changes or system interaction would cause).

1.2 High-fidelity versus low-fidelity sensors

The two types of modalities that are studied in this paper also emphasize a trade-off between good sensor modalities from the point of views related to a detection and monitoring, privacy, and reliability.

The stronger, high-fidelity sensors tend to produce rich information at a high rate so that monitored phenomena are captured in detail. As a result, high accuracy, precision, and recall figures are feasible. Often, maintaining this system is harder and the data is, especially in home environments, sensitive in nature. Video footage of the user's bedroom is by far the best sensor data to observe sleeping postures, but as it is also placed in one of the most sensitive environments, few users would actually permit installation. Keeping a video system running also tends to require more maintenance and bulkier hardware.

The weaker, low-fidelity sensors on the other hand often are associated with less privacy concerns. This comes at the price of less detailed data being captured however, with a significant amount of uncertainty being present for the analysis algorithms to handle. These sensors produce less data that are harder to mine.

1.3 Paper overview

The contributions of this paper are threefold. An evaluation of all feasible scenarios is carried out to find the most workable solution to capture, process, and log images from sleeping postures. Two studies then investigate how well both types of sensors on their own would work in terms of accuracy when deployed in a domestic setting, and assuming they are trained by the user. Finally, a combination of both modalities is proposed and an experiment is described to indicate under which parameters this combination would result in a better system.

This paper is structured as follows: First, the case study of sleep posture detection and logging is presented as an illustration for the type of long-term domestic monitoring task we pursue in this paper. Then, video footage and wrist acceleration are presented and analyzed on their individual merit for the purpose of sleep posture monitoring. Section 5 will then offer the combination of these two as a better alternative and studies associated trade-offs. The last section will finally present our conclusions and point to future and ongoing work.

2 Case study: sustained logging of sleep postures

The posture in which we sleep can have a large effect on the quality of sleep, with research [4] looking at sleep postures as one of the main pieces of information to record in studies. Patients with obstructive sleep apnea, a sleep disorder characterized by pauses in breathing during sleep, should avoid sleeping on the back, and are strongly encouraged by sleeping laterally (on one's side). Other research [7] correlates usual sleep postures with a person's personality profile.

There are several types of postures that are of interest in the aforementioned studies: Medical articles mostly investigates the *lateral* (lying on the left or right side), *supine* (lying on back), and *prone* (lying on chest) sleep postures. Other studies take a more detailed look at the full body and follow a different naming convention for common postures such as *foetus*, *soldier*, or *starfish*. Some examples for both are depicted in figure 1.

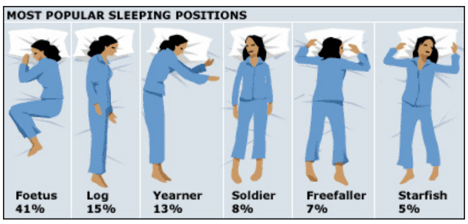

Fig. 1: Some common sleeping postures (reprinted from [2] with permission from Prof. Idzikowski from the Edinburgh Sleep Centre.)

Sleep postures are therefore not only important indicators to watch for a typical one-to-two-night study, having a longer record would mean that trends could be observed and correlated to daily activities and habits. In the following sections we will discuss two methods to capture these sleep posture trends.

3 High-fidelity data: IR video footage

Videometry is a fixed part of every cardiorespiratory polysomnography, which is the central measuring method of stationary diagnostics in sleep laboratory [11]. However, in most cases the posture is measured by tilt sensors, which are worn on the chest. The data from the night vision camera is used redundantly to monitor the patient. Although this is currently not done, this video material could be used to automatically detect the posture of the patients.

The reliability of performing (semi-) automatic posture analysis can be expected to be quite high. The resolution and camera angle can be optimally chosen for covering the user's full body, and state-of-the-art cameras are widely available with built-in IR lighting that is strong enough to reach an entire bed. The straightforward approach taken in this paper is based on having enough personal training data, but other approaches in pedestrian pose recognition [1] could be employed for immediate deployments.

The privacy aspects of video footage are not as good as for other polysomnographic data sources. The pictures, respectively videos, of the patient's sleep are highly personal and tend to be a major concern of patients, even more so than the other data recorded by polysomnography since that requires additional analysis and interpretation. In the sleeping laboratory, the privacy of the video data together with the recorded sensor data is determined by local policies and regulations. For a deployment at home such regulations in terms of a privacy policy are mandatory.

Videometry is in general easy to deploy and maintain, since video cameras can be wall-powered in the sleeping environment. Another fact that makes videometry relatively easy to deploy is that the video camera is independent from the patient, unlike other sensors that are being worn by the patient, and that due to the patient's motion during sleep tend to be affected. In our research a mounted camera and a recording medium are needed, and a wide selection of both is available as standard components.

3.1 Implementation

Physical setup Our prototype for recording videographic data consists of a night vision camera and a stand. The camera is a commercial product with built in IR-LEDs and a pan-tilt camera head. The lightweight stand that supports the camera is capable of positioning the camera above the subject, so that at least the subject's head and chest area (see figure 2) can be recorded in detail. This particular camera is capable of different communication channels. Its interface can be accessed either by LAN or WLAN with WEP, WPA, or WPA2 encryption. The video data can also be stored via USB mass media, or via video streaming over LAN, WLAN, or BNC. We will discuss some possible communication strategies in section 3.2.

(a) A stand allowing full body capture

(b) Example of a commercial IR camera

(c) Camera's good resolution and lighting

Fig. 2: Components for the video modality

Posture detection To detect the patients' postures, we are using a semi-automatic procedure. As we are operating on images recorded by the night vision camera, the challenge of the procedure is to extract those images, which mark a new sleeping posture by pruning images automatically in three steps followed by a manual analysis of the user. Thus the procedure is subdivided into four steps:

1. Automatically reducing the number of images by limiting the camera's frame rate to 1 fps. This can be done without loosing important information, because we want only to detect sleep phases, that are lasting longer than a few seconds.
2. Automatically reducing the number of images by using the camera's built-in motion detection algorithm. To ensure 100% recall the threshold of the motion detection algorithm is lowered as much as possible. This causes a comparatively high false positive rate, but reduces the false negative rate.
3. Automatically reducing the number of images by applying an averaged image differences algorithm (ADIA) on the remaining images. The algorithm is described below.
4. Manually pruning false positives from the remaining images by the user on the automatically reduced image set.

The averaged image differences algorithm (ADIA) works on the differences between successive images. It can be subdivided into three steps:

1. Calculate the differences of the successive images. We only store the average number of differing pixels, which is calculated as follows:

$$0 \leq ad = \frac{sum}{width * height * depth * num\_channels} \leq 1 \qquad (1)$$

   where *sum* is the sum of all pixels for every channel. The averaged difference for every pair of successive images is stored in a vector *ads*.

2. The vector of averaged differences is then mean filtered with a dynamic kernel size. A size of one percent of the total vector size has shown to be working well. While changing a posture we typically detect heavy variation of the averaged differences. By mean filtering *ads* we can reduce this unwanted behavior.
3. To detect posture transitions, we threshold the filtered vector *ads* and extract the falling edges to ensure that we select images that differ enough from each other. The threshold is calculated as the arithmetic mean of the vector of averaged differences.
   Next, the image is determined were the motion has settled down and the postures are "stable". It turned out that the images, which correspond to the next local minima after falling edges in *ads* are appropriate for this procedure. The new posture starts in the vicinity of the last local maxima before falling edges in *ads*.

The images which correspond to the local minima are shown to the user who then annotates these images. If the posture has changed compared to the last identified posture, we store the last local maximum as a starting point.

3.2 Analysis / Experiments

Privacy and deployment Security is a fundamental concern, especially due to the personal nature of the recorded video data. Therefore, we introduce five possible scenarios for collecting the data and discuss their advantages and disadvantages. The analysis is subdivided into three parts: configuration, recording and transfer. The configuration comprises the user's setup of the camera which consists of the configuration of the recording scope by panning and tilting the camera's head. This can be done only by accessing the camera via the provided web interface and therefore requires a connection via LAN or WLAN. The recording part consists of the storage for the pictures or video material on an external device (via the built-in USB port or by accessing directly the video stream of the camera via LAN or WLAN). The transfer part comprises the transmission of the recorded visual data from the storage medium to a computer, where it can be post-processed. The different scenarios are depicted in figure 3.

The data has to be securely transferred due to the personal nature of the recorded data. In scenario 4 data is recorded via an ad-hoc connection. The camera's software is only capable of transferring data with WEP encryption in ad-hoc mode. As WEP encryption is highly insecure [12], scenario 4 is dismissed. Another important factor is that a person's sleep should not be influenced by the used equipment to ensure an usual sleeping environment. Scenario 2 and 4 contradict this assumption since a wireless connection would disturb a person's sleep due to radiation from a wireless connection [6]. Besides this fact the reliability of the overall system is increased by using a wired connection in contrast to the more failure-prone wireless connections. Scenarios 1, 3 and 5 differ only in the configuration part. In scenario 1 a secure WPA2 encrypted connection is used. However this depends on an additional WLAN router. This router is not

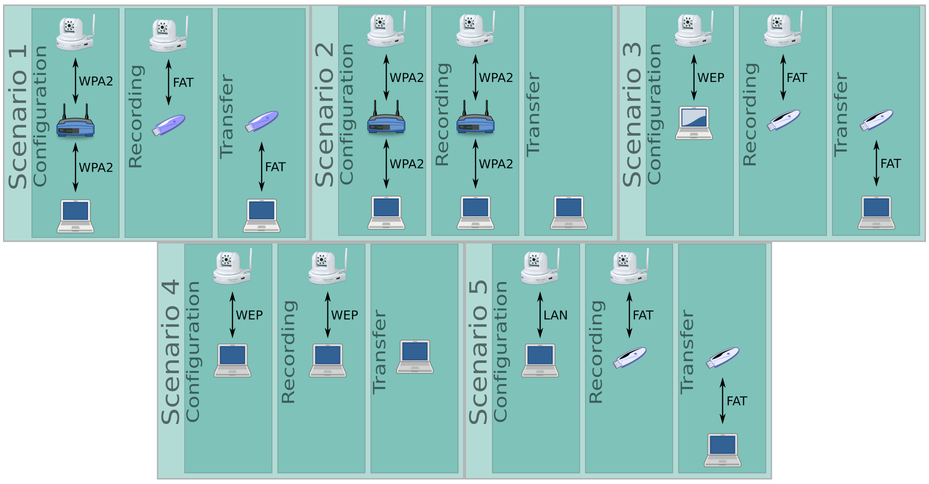

Fig. 3: Video capture scenarios

needed in scenario 3, where the camera is configured via an insecure WEP encryption. In scenario 5 no WLAN router is needed at all and the security factor can be dismissed here. However, this scenario depends on a cabled connection, which requires a nearby PC or laptop.

As already mentioned, security is crucial and an additional WLAN router would be far more complex than using a simple cabled connection, reducing the required hardware to a minimum. With scenario 5 the subjects will consequently be confronted with the fewest inconveniences. Trust in the setup is increased by the fact that the camera is physically disconnected from any network connection during recording.

Prototypical implementation To test the described algorithms we have implemented a prototype. Figure 4a shows a screenshot of the prototype with an already tagged dataset. The postures are described by the colored areas in the plot. We have tested the prototype on three datasets which where recorded on different nights. The number of false positives were determined immediately after recording (phase 1), the built-in motion detection algorithm (phase 2), ADIA (phase 3) and the manual pruning of the user (phase 4). The reduction of the false positives to zero is depicted in figure 4b.

Many of the false positives after running ADIA on the images result from heavy motion without changing posture. To automatically eliminate those false positives a highly sophisticated algorithm would be needed. Due to the low number of false positives, it is not necessary to implement another algorithm which can detect the postures automatically. Because the recorded scenes vary heavily for different users, it is doubtful that an algorithm can automatically detect the postures reliably for every user. There are some promising research results in sleep posture detection [14].

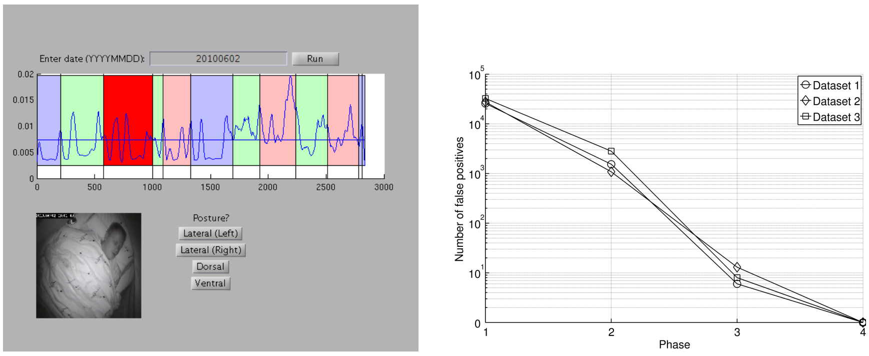

(a) GUI of the prototypical implementation.

(b) Reduction of false positives.

Fig. 4: GUI and reduction of false positives.

The implemented algorithm ADIA deteriorates on changing light conditions. The number of false positives raises when the illumination of the recorded scene is changing fast. This behavior can be observed e.g. on cloudy days after dawn. Another problem arises from sleep postures which cannot be visually detected because the blanket is covering the complete body.

4 Low-fidelity data: Wrist-worn accelerometer

Actigraphy devices are traditionally used for monitoring sleep disorders outside sleeping labs, as they detect movements which are used to infer sleep/wake patterns and quality of sleep. Here, we use an actigraph-like device with a 3D accelerometer instead which can also record postures.

Unlike the video camera of the previous section, this sensor modality is a lot weaker and therefore the detection reliability can be expected to be lower. The 3D accelerometer allows to detect various fine-grained orientations and can be read a relatively high frequencies (typically 100Hz or more). However, as this device is worn at the wrist, a correlation between wrist- and body posture is assumed. An indication for this can be found in [13], where a wrist-worn tilt sensor was successfully used to extract sleeping data and sleeping postures.

The data is stored directly on the sensor unit, and therefore this approach is not as sensitive as wireless monitoring. Also, the data is not immediately human-readable. However, privacy issues still remain when the device is handed over or the data is uploaded into a central system.

The sensor unit is designed to run over long periods of time and continuously without user interaction and therefore is relatively easy to deploy. Although current implementation of the prototype requires that users take off the sensor before showering or swimming, the device needs minimum maintenance compared to the video modality.

4.1 Implementation

The sensor unit is worn by the user at the dominant wrist. Before the sensor unit is handed to the user it is fully charged and switched on to record accelerometer data. The user is not concerned about recharging it or restarting the sensing mode. Should nevertheless the sensing unit not perform appropriate would this be observable after the sensing unit was given back by the user. As already mentioned, the current design of the sensor does not allow to be used with water.

The sensor unit is equipped with a 3D accelerometer and captures data at 100Hz, storing it on a 2 GB SD card. The sensor board is powered by a lithium battery and lasts for almost a month when used only in accelerometer mode. The board can be equipped with even more and especially different sensors (e.g. temperature and light), but currently only the accelerometer is used. Such a set-up is very cost-effective and can be deployed easily. The program running on the sensor board is very power-efficient for the hardware since data is only stored when necessary and equal sensor readings are enumerated and stored with the counter to the memory.

After the sensor board was given back by the user the data is analyzed. Night segments are being extracted by calculating the variance in the accelerometer data over a window of one minute. For that a threshold t is set that detects movement against non-movement by marking a window sleeping whenever the variance is below t.

The residual data consists of values that describe a persons postures. These values remain constant over a longer time span and characterize a certain posture. We want to detect these postures and describe our approach in the following section.

4.2 Analysis / Experiments

Experiments were conducted with one test subject who wore the sensor at the dominant wrist for five nights. For initial examination the recording started one hour prior and stopped one hour after sleep. To display how well postures can be detected, three different cases are described.

Case 1: artificial data set The first case uses an artificial dataset to detect different postures. The artificial dataset is obtained by lying for ten seconds in the postures left-, right lateral, supine and prone each time before sleep.

The posture estimation is shown in figure 5. The ground truth with the estimation of the postures is shown in the top plot. For case 1 we will focus on the artificial set, the other two sets will be explained in the following parts. By visual inspection it is evident that posture changes are detected accurately, whereas the posture itself has been detected poorly. A few overlaps in ground truth and estimation of the posture are visible.

Precision and recall for the postures are in the range of 62% and 35% respectively (see figure 6, left plot). Remarkably, the detection of the right lateral posture exhibits a recall of 95% and a precision of 53%. The false positive rate for the other postures is very high and can be explained by the fact that each posture is not represented by the value obtained from the artificial dataset. Due to the spontaneity of a person's positioning during sleep it is almost impossible to simulate an accurate artificial dataset. The position of the wrist during a posture while sleeping varies frequently, resulting in different accelerometer values.

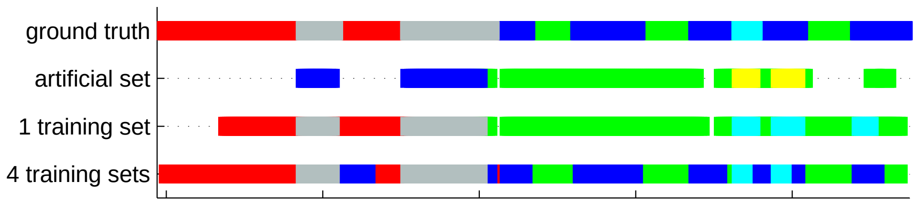

Fig. 5: Detected postures with increasing amount of training sets.

Case 2: 1 training set To compare the first case to a different approach, in the second case a k-nearest-neighbor (KNN) classifier is trained on one night only. The classifier is then tested on the four other nights and compared to the ground truth. This procedure is repeated for each of the 5 nights.

The visual results are shown in figure 5, described by the one day training set. In contrast to case 1 the results of detecting posture changes and the posture itself improved. Overall a precision of 48% and recall of 83% are obtained. Precision dropped in contrast to case 1 but note that recall increased (see figure 6, middle plot). A classifier was obtained by training on one night only. In theory a person takes in only characteristic postures, limiting the amount of postures of a person [4]. Therefore the number of training datasets (or nights) has to be sufficient to gain a classifier that covers all possible postures.

Case 3: 4 training sets Finally in the last case a KNN classifier is obtained by training with posture labels and accelerometer data using 5-fold cross-validation, where one fold corresponds to one night. The difference to case 2 is that the classifier is trained on four datasets.

The results are summarized in figure 5, showing again an improvement of detection of posture changes and the estimated postures, almost overlapping with the ground truth. Notice that in contrast to the previous plots of case 1 and 2 no unclassified posture is detected, leaving no whitespaces in the plot. By training a classifier on more than one night the precision is increased to 88%, whereas recall stays steady in the range of 80% (see figure 6, right plot). Using only an artificial dataset is not feasible due to user-specific postures and is therefore not further followed.

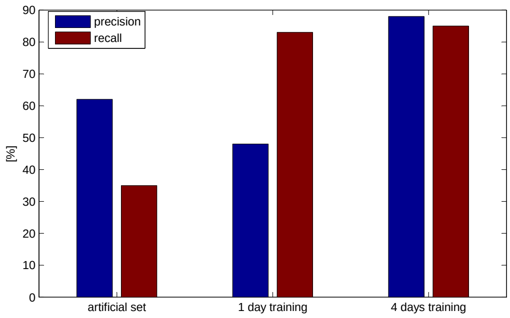

Fig. 6: Precision and recall for all three cases.

4.3 Combining both modalities

We have shown that both systems can be operated individually to detect sleeping postures. We have found that, as expected, the video footage is superior to a wrist-worn sensor in terms of detection reliability, since its data contains more details on the sleeping posture, assuming that the user prunes the system by sorting out false positives manually after each night. This is first and foremost feasible since the camera's output can be directly interpreted by the user. The wrist-worn sensor on the other hand operates on its own without the need of installing environmental devices and without capturing data as sensitive as video. Although a full study still remains to be performed on the privacy issues of this paper's particular wrist-worn sensor, we can deduce from similar devices, actigraphs, that they bring along far less privacy concerns compared to video footage.

By combining both modalities into one, we can overcome the limitations of the individual parts: In the combined system we propose to use the visual data to obtain ground truth, which is then used to train the wrist-worn sensor unit for multiple nights. By following this approach, we improved for the earlier mentioned dataset the overall detection rate to over 80% in both precision and recall by only training the wrist-worn sensor units on the data from four nights (see figures 5, 6). This approach maintains the privacy of the wrist-worn sensor unit by using the inertial posture data for detecting the subjects sleeping postures, and storing this information alone. The data of the camera can be immediately discarded after usage in the training phase, which can be done at the user's home without involvement of third parties. The video data is thus accessible only by the user, thereby reducing the privacy concerns of the overall system to the wrist-worn sensor unit's privacy issues. The price to pay for these two advantages are (1) a slightly degraded applicability, due to the deployment of the video system and the fact that the inertial data has to be downloaded after each night *for the duration of the training phase*, and (2) the video-based privacy issues remaining during the training phase as well.

As the main weakness of the combined system is the training phase, it is important that this lasts as short as possible. To investigate how many days are needed to obtain an acceptable precision and recall for each posture, we increased the number of training nights in the cross-validation.

Figure 7 shows the given training sets' results for the right lateral posture: recall and precision are increased from 77% to 84% and 72% to 82% respectively by using two training sets instead of one only. Both are improved by four training sets, leading to a recall of 97% and a precision of 84%. We conclude that one training set is not sufficient to detect the postures appropriately, but a training set of four is leading to an overall precision and recall of over 85% respectively.

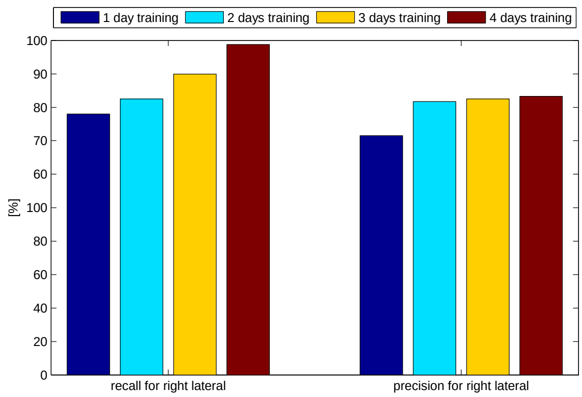

Fig. 7: Precision and recall for right lateral: more training nights improve the recognition.

5 Conclusions

This paper reported on our efforts to design a domestic sleep posture monitoring system for long-term deployment, which allows sleep researchers to track their patients outside the laboratory. We argued that the long runtime and the home environment make for a difficult design: We identified three challenges that such a system must meet, and used these as a framework to describe and analyze two sensor modalities: video monitoring of the sleeping user, and a wrist-worn accelerometer logger. Through several studies, we have explored how well these fare on their own, and proposed finally a multi-modal solution that uses both in a system that reduces all challenges.

By combining video information and sensor data we obtained precision and recall of over 80% respectively which are promising results for long-term studies. The setup is easy to deploy and requires only little user interference, which increases the acceptance on the persons' side. Such a system provides a potentiality for the use in sleep studies where patients suffer from for example sleep apnea and have to be monitored in their usual environment.

Further experiments over more nights have to be conducted to confirm our findings. The sensor modality is continuously improved to make it more energy efficient and thereby enhancing recording time for long-term experiments. We are also researching on a full automatic posture detection by using only the video modality. In our approach we follow the procedure, which is used to detect people in [3].

6 Acknowledgements

This work was sponsored by the project *Long-Term Activity Recognition with Wearable Sensors* (LA 2758/1-1) from the German Research Foundation (DFG).

References

1. Andriluka, M., Roth, S., Schiele, B.: Pictorial structures revisited: People detection and articulated pose estimation. In: Proc. CVPR. vol. 1, p. 4 (2009)
2. BBCNews: Sleep position gives personality clue. http://news.bbc.co.uk/2/hi/health/3112170.stm (2003)
3. Dalal, N., Triggs, B., Rhone-Alps, I., Montbonnot, F.: Histograms of oriented gradients for human detection. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2005. CVPR 2005. pp. 886–893 (2005)
4. Gordon, S., Grimmer, K., Trott, P.: Self reported versus recorded sleep position: an observational study. In: Internet Journal of Allied Health Sciences and Practice Vol2No1 (2004)
5. Hong, J., Ng, J., Lederer, S., Landay, J.: Privacy risk models for designing privacy-sensitive ubiquitous computing systems. In: Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques. pp. 91–100. ACM (2004)
6. Hyland, G.: Physics and biology of mobile telephony. The Lancet 356(9244), 1833–1836 (2000)
7. Idzikowski, C.J.: Learn to Sleep Well. Duncan Baird (2000)
8. Ojala, T., Kukka, H., Linden, T., Heikkinen, T., Jurmu, M., Hosio, S., Kruger, F.: Ubi-hotspot 1.0: Large-scale long-term deployment of interactive public displays in a city center. Internet and Web Applications and Services, International Conference on 0, 285–294 (2010)
9. Oksenberg, A., Arons, E., Greenberg-Dotan, S., Nasser, K., Radwan, H.: The significance of body posture on breathing abnormalities during sleep: data analysis of 2077 obstructive sleep apnea patients. Harefuah 148(5), 304 (2009)
10. Peng, J., He, H., Zhu, P., Liu, Y., Zhang, X., Jin, Y.: Zigbee-based new approach to smart home. Journal of Shanghai University (English Edition) 14(1), 12–16 (2010)
11. Peter, H., Penzel, T., Peter, J.: Enzyklopädie der Schlafmedizin. Springer, Berlin, 1 edn. (03 2007)
12. Tews, E., Weinmann, R., Pyshkin, A.: Breaking 104 bit WEP in less than 60 seconds. Cryptology ePrint Archive 2007/120, Cryptography and Computeralgebra group, Darmstadt University of Technology (2007), `http://eprint.iacr.org/2007/120`
13. Van Laerhoven, K., Borazio, M., Kilian, D., Schiele, B.: Sustained logging and discrimination of sleep postures with low-level, wrist-worn sensors. In: ISWC (2008)
14. Wang, C., Hunter, A.: Robust Pose Recognition of the Obscured Human Body. International Journal of Computer Vision pp. 1–18 (2010)