Zero-Effort In-Home Sleep and Insomnia Monitoring using Radio Signals

CHEN-YU HSU, AAYUSH AHUJA, SHICHAO YUE, RUMEN HRISTOV, ZACHARY KABELAC, and DINA KATABI, Massachusetts Institute of Technology

Insomnia is the most prevalent sleep disorder in the US. In-home insomnia monitoring is important for both diagnosis and treatment. Existing solutions, however, require the user to either maintain a sleep diary or wear a sensor while sleeping. Both can be quite cumbersome. This paper introduces EZ-Sleep, a new approach for monitoring insomnia and sleep. EZ-Sleep has three properties. First, it is zero effort, i.e., it neither requires the user to wear a sensor nor to record any data. It monitors the user remotely by analyzing the radio signals that bounce off her body. Second, it delivers new features unavailable with other devices such as automatically detecting where the user sleeps and her exact bed schedule, while simultaneously monitoring multiple users in different beds. Third, it is highly accurate. Its average error in measuring sleep latency and total sleep time is 4.9 min and 10.3 min, respectively.

Additional Key Words and Phrases: Sleep and Insomnia Monitoring, Wireless Sensing, Device-Free Localization

ACM Reference format:
Chen-Yu Hsu, Aayush Ahuja, Shichao Yue, Rumen Hristov, Zachary Kabelac, and Dina Katabi. 2017. Zero-Effort In-Home Sleep and Insomnia Monitoring using Radio Signals. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3, Article 59 (September 2017), 18 pages.
https://doi.org/10.1145/3130924

1 INTRODUCTION

Insomnia and sleep deprivation are common health problems in the US. One in every three Americans do not get enough sleep, and about 10% of the population suffers from chronic insomnia [37]. Chronic insomnia increases the risk of heart disease, kidney failure, high blood pressure, diabetes, and stroke [33]. The statistics are even worse among the elderly, where 50% of seniors experience periods of insomnia that last for weeks, months or even years [37, 48]. In-home sleep monitoring is important for both detecting insomnia and treating it. People's impression of their sleep can be wildly wrong, off by hours in some cases [16, 42]. On the other hand, older patients may not report their insomnia symptoms in the first place [37]. Further, there is an increasing interest in replacing drug-based treatment with Cognitive Behavior Therapy for Insomnia (CBT-I). CBT-I is based on gradual adjustments of sleep schedule and the time spent in bed, and hence in-home monitoring of those parameters is beneficial for treatment tracking [24].

Monitoring insomnia is difficult. The gold standard in sleep monitoring is overnight Polysomnography (PSG), conducted in a hospital or sleep lab, where the subject sleeps with EEG, ECG, EMG, respiration, and pulse monitors. PSG does not work well for insomnia because sleeping away from one's bed with obtrusive sensors causes difficulties in falling asleep that are unrelated to insomnia (the first night effect [14, 40]). In fact, most

2474-9567/2017/9-ART59 $
https://doi.org/10.1145/3130924

insomnia studies use patient diaries, requiring people to keep daily records of when they go to bed, how long it takes them to fall asleep, how often they wake up at night, etc. This approach creates significant overhead and is hard to sustain over long periods. More recently, actigraphy-based solutions have been used to track motion and infer sleep patterns, but they must instrument the user with accelerometers at the wrist or hip [36]. Many people do not feel comfortable sleeping with wearable devices [17], and older adults are encumbered by wearable technologies and may simply take the device off creating adherence issues.

Insomnia and long term sleep monitoring should be zero effort without sacrificing accuracy. It should provide in-home continuous monitoring without requiring the user to wear a device or write a diary. It should also measure the key sleep parameters used for insomnia assessment. This means it needs to measure the time between going to bed and falling asleep, or sleep latency (SL), the percentage of sleep time to the time in bed, or sleep efficiency (SE), the total sleep time (TST), and the amount of wakefulness after falling asleep (WASO).

We introduce EZ-Sleep, a sleep sensor that achieves these goals. EZ-Sleep is zero effort – all that the user has to do is to put EZ-Sleep in her bedroom and plug it to the power outlet. EZ-Sleep works by transmitting radio frequency (RF) signals and listening to their reflections from the environment. By analyzing these RF reflections, EZ-Sleep automatically detects the location of the user's bed, identifies when she goes to bed and when she leaves the bed, and monitors the key insomnia assessment parameters SL, SE, TST, and WASO. Further, because it directly measures the user's bed routine, i.e., when she enters and exits the bed, EZ-Sleep can be used in CBT-I to monitor patient compliance with prescribed changes in her bed schedule.

The design of EZ-Sleep builds on recent advances in wireless systems, which show that by transmitting a wireless signal and analyzing its reflections, one can localize a person and track her vital signs without any wearables [11, 12]. However, past solutions that leveraged these advances in the context of sleep have limited themselves to analyzing the user's vital signs (mainly breathing) as extracted from the RF signals [32, 39, 44, 49]. In contrast, EZ-Sleep uses the RF signal to extract both the user's breathing and location, and combines both to infer the user's bed schedule and sleep quality. In the absence of information on when the user goes to bed and leaves the bed, past work cannot compute key insomnia parameters like sleep latency, sleep efficiency and WASO. Furthermore, most past solutions rely on the Doppler effect which is highly sensitive to interference from other sources of motion in the environment such as potential neighbors or flatmates (see 9.1 for empirical results). In contrast, by leveraging RF-based localization, EZ-Sleep is not only robust to motion in the environment but can also monitor the sleep of multiple subjects at the same time.

Fig. 1. EZ-Sleep setup in one of our subjects' bedroom.

We designed EZ-Sleep as a standalone sensor as shown in Figure 1. Our design involves four components that work together to deliver the application:

- *Monitoring the user via RF signals:* EZ-Sleep uses an FMCW radio and an antenna array to separate RF signals based on the location of the reflecting body. This allows it to track both the location of the user and her breathing, and use both to extract the user's sleep patterns.
- *Identifying bed location:* EZ-Sleep introduces a novel algorithm for automatically detecting where the user goes to sleep. The algorithm first identifies locations where the user stays stationary like chairs, desks, couches and beds. It then classifies these locations as bed vs. non-bed. To do so, EZ-Sleep leverages RF-based localization to analyze how the user uses the space. It builds a map of where the user spends her time. Using tools from image processing and machine learning, EZ-Sleep segments the map into meaningful areas.
- *Detecting bed entries and exits:* Knowing when the person goes to bed is essential for both assessing insomnia and monitoring compliance with CBT-I. But, since RF-localization error can be as high as one meter, we cannot simply rely on the location estimate to detect when the person goes to bed. Instead, EZ-Sleep employs a Hidden Markov Model (HMM), where location measurements act as observations of the hidden state which takes one of the two values: in-bed or out-of-bed.
- *Classifying sleep and awake periods, and computing sleep parameters:* Once it knows the user is in bed, EZ-Sleep zooms in on the RF signal reflected from the bed region. EZ-Sleep feeds this signal to a deep neural network model, which operates in two phases. In the first phase it learns to classify the time in bed into sleep and awake epochs. In the second phase, it zooms in on the first transition from awake to sleep, and trains a second model customized for learning sleep onset. Once it knows entry and exit from the bed, the sleep onset, and any later awakenings, EZ-Sleep can compute all of the insomnia assessment parameters described above.

We have implemented EZ-Sleep and evaluated it in 8 homes. We collected a total of 100 nights of sleep, including 30 nights with an EEG-based FDA-approved sleep monitor [7]. Our results show correct bed identification in all places. The results also show that EZ-Sleep is highly accurate. Specifically, its average errors in computing SE, SL, TIB, TST, and WASO are 2.8%, 4.9 min, 3.2 min, 10.3 min, 8.2 min respectively. These results show that EZ-Sleep's accuracy is comparable or better than medical grade actigraphy devices [19, 20, 30].[1] Thus, EZ-Sleep can deliver clinically meaningful sleep parameters without asking the user to wear any sensor or record her sleep data. Furthermore, the results also show that EZ-Sleep can monitor the sleep of two subjects simultaneously, making it the first RF-based sleep monitor that works with multiple people.

Contributions: To our knowledge, EZ-Sleep is the first RF-based sleep sensor that automatically detects bed locations, identifies bed entries and exits, and monitors key insomnia parameters like sleep latency, sleep efficiency, total sleep time, and wake after sleep onset. It is also the first RF-based sleep sensor that can monitor multiple users simultaneously. These properties are enabled by the observation that sleep monitoring can benefit from combining location tracking with temporal analysis of the user's breathing signal, both extracted from RF reflections. We introduce new algorithms that use tools from image processing and machine learning to analyze the user's location and breathing signals and infer her bed schedule and sleep quality. We further implement our design as a standalone sensor which we deploy and evaluate in eight homes, demonstrating its high accuracy.

2 BACKGROUND AND RELATED WORK

We first provide background on insomnia monitoring and then describe past solutions, both in industry and academia.

[1]Medical grade actigraphy errors are as follows: SE: 3~5%, SL: 0.2~3 min, TST: 8~19 min, and WASO: 12 min [19, 20, 30]. Since actigraphy cannot measure TIB, these estimates require the user to push a button to indicate going-to-bed and leaving-bed.

2.1 Key Insomnia Parameters

Insomnia and sleep deprivation are typically assessed using the following sleep parameters [30, 41]:

- *Sleep latency (SL)* is the time between going to bed and falling asleep. Sleep latency is perhaps the most common metric for assessing insomnia. A sleep latency that is longer than 30 minutes for more than two nights per week is a sign of insomnia.
- *Total sleep time (TST)* is the total time in bed actually spent in sleep. It captures whether the person gets enough sleep.
- *Time in Bed (TIB)* which is also called time from light out to light on. It puts an upper bound on the sleep time and can show that the person does not allow herself enough time to sleep.
- *Sleep efficiency (SE)* is the ratio of the total sleep time to the total time in bed, (i.e., TST/TIB).
- *Wake after sleep onset (WASO)* is the total duration of wakefulness occurring after sleep onset. It captures sleep fragmentation, or the inability to sustain sleep. For example, some older people wake up after a few hours of sleep and cannot go back to sleep.
- *Number of awakenings:* This metric counts the number of awakenings that last for more than 5 minutes (NA > 5). In this paper, we experiment with healthy individuals, and have not had any reports of night awakenings. Hence, we do not report results for this metric. However, our design can compute this metric because it classifies all 30-second epochs as awake or asleep, as explained in section 7.

Figure 2 shows the relation between these sleep parameters.

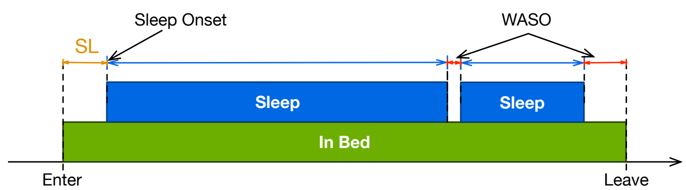

Fig. 2. The definition of all sleep parameters.

2.2 Sleep Monitoring Solutions

Laboratory-based Polysomnography (PSG) is the medical gold standard for sleep monitoring. It requires the person to spend the night in a sleep lab connected to a dozen sensors, including EEG scalp electrodes, an ECG monitor, a respiratory chest band, a nasal probe, etc. PSG monitors many aspects of sleep including sleep stages, sleep parameters, and sleep apnea. Yet, PSG is both highly obtrusive and impractical for long term studies [28]. The discomfort with the sensors changes the user's sleep behavior, making it hard to measure insomnia. Further, since night-to-night sleep variability is a important factor for diagnosing insomnia [15], a single night of study is often insufficient.

Patient diaries are commonly used for monitoring insomnia. The person is asked to keep daily records of when they go to bed, how long it takes them to fall asleep, how often they wake up at night, etc. Writing sleep diaries requires a significant effort and it is hard to sustain for long periods of time. Medical grade actigraphy has been used to track user movement to infer sleep. Actigraphy measures user motion using accelerometers tied to a person's wrist or hip. However, some people feel encumbered to sleep with wearable devices [17] , and older people and kids may simply take the device off during the night.

The consumer industry has developed wellness devices that track sleep. They encompass activity trackers such as FitBit and Jawbone, smartphone-based solutions like Sleepbot and Sleep Cycle, bed side sensors or bed pads inserted under the sheets such as Beddit [2–4, 6, 8]. These solutions have lower accuracy than the medical grade devices [22]. Some of these solutions require the user to input when they are physically in bed [5] or may have lower accuracy if this information is missing [2].

The research community has shown a great interest in monitoring sleep using smartphones or systems that sense environmental factors [17, 18, 23, 25, 27, 35]. They infer sleep quality using data from microphones, accelerometers, cameras, phone usage, etc. However, recording audio and video information throughout the night may still be considered obtrusive for some people due to privacy and comfort reasons [39]. Also, these systems may not be easy to configure for certain sectors of the populations, like the elderly or children.

In terms of technology, EZ-Sleep is closest to past work on using radio signals for monitoring sleep stages or in-bed movements [31, 32, 39, 44, 49]. All of these systems work by extracting a person's breathing using radio signals that bounce off the person's body. They then classify the night into periods that correspond to different sleep stages or body movements. EZ-Sleep differs from these systems along four axes. First, EZ-Sleep analyses sleep behavior both in time and space. It does so by combining RF capabilities to extract breathing and location. Second, EZ-Sleep computes sleep parameters unavailable to past systems including TIB, SL, and WASO. Third, EZ-Sleep introduces new algorithms that automatically identify where a person sleeps and track her bed entries and exists. Fourth, EZ-Sleep can simultaneously monitor multiple users with one device.

3 EZ-SLEEP OVERVIEW

EZ-Sleep is an in-home sleep monitoring system that requires zero user effort. It operates by transmitting a low power wireless signal and capturing its reflections off users in the vicinity of the device. By analyzing the signals, it learns the bed location, when the user goes to bed, and when she falls asleep. Based on this knowledge, it generates sleep parameters including Time in bed (TIB), Sleep latency (SL), Total sleep time (TST), Sleep efficiency (SE), and Wake after sleep onset (WASO).

The EZ-Sleep sensor is a software-hardware system whose operation involves the following key steps:

(1) Capturing the user's location and breathing using RF signals.
(2) Inferring bed areas by observing the user's movements through space.
(3) Detecting bed entry and exit events by interpreting noisy location measurements.
(4) Classifying sleep versus awake when the person is in bed and estimating sleep parameters.

In the following sections, we explain each step in detail.

4 EXTRACTING LOCATION AND BREATHING FROM RF SIGNALS

The relationship between motion and sleep is intuitive and constitutes the core of all actigraphy-based sleep sensors. But sleep sensors, so far, have looked only at the amount of motion over time – i.e., actigraphy. We observe however that some sleep parameters, such as going-to-bed, can be better inferred by analyzing motion as a function of space. Wireless signals are a powerful tool to capture a person's motion both in time and space. When the person is stationary, her motion over time is mostly her breathing motion [13]. Thus, in designing EZ-Sleep, we leverage these RF capabilities. The literature already shows the capacity of RF signals to capture both location [11, 12, 29, 45] and breathing [13, 39, 46]. We introduce these techniques briefly and refer the reader to [10, 11, 13] for more information.

EZ-Sleep uses a combination of an FMCW radio and an antenna array. The FMCW technique measures the distance of the reflecting body from the device, whereas the antenna array measures the spatial direction of the reflector with respect to the device. Thus, together, they allow us to divide the x-y plane into pixels and separate RF reflections from different pixels in space.

Once we have separated the signal from each x-y pixel, we need to first identify human reflections from reflections of other objects in the environment such as walls and furniture. To do so, we apply two different filters to the signals to capture two types of human motions. The first is a high pass filter that captures fast and non-periodic human motion, including walking and any hand or limb movements. The second is a band pass filter around the breathing frequency that captures a user's chest motion when she is breathing and not actively moving other body parts. Note that since furniture and walls neither move at a periodicity similar to human breathing, nor have non-periodic motion, their reflections disappear at the output of both filters.[2] Applying these filters to the signal allows us to detect whether the x-y pixel has a human at that time, and whether the human is only breathing or he is performing a bigger non-periodic motion.

In our implementation, we generate one measurement per pixel every 50 milliseconds. Each measurement contains a time sample of the RF signal reflected from that pixel at that time, and a tag that states whether the pixel has a moving person, a stationary person, or is empty (i.e., has no person). We project the area covered by EZ-Sleep into a 500-by-500 pixel image, where each square pixel occupies a 0.024-by-0.024 square meter space.

Finally, we note that RF-based localization suffers from environment noise and multi-path problems, and its 90 percentile error can be more than one meter [11, 29, 45]. Similarly, breathing signals extracted from RF reflections are highly sensitive to body motion, and can be erroneous when the user moves her limbs [13]. Our system is designed to be robust to these phenomena and can measure sleep parameters accurately in the presence of noise.

5 BED IDENTIFICATION

In this section, we describe EZ-Sleep's algorithm for detecting beds in the home. The algorithm first identifies areas in the home where the user is mostly stationary. Such areas span the bed and various seating places such as chairs and couches. Next, EZ-Sleep identifies which of these stationary areas are beds.

Potential bed areas are places where people stay stationary for long periods of time. Thus, we would like to identify areas in the home where the radio sees the person stationary for significant periods. Recall that at every 50 millisecond, we have a tag that identifies every pixel as being stationary, moving, or empty. Thus, the first step of our algorithm computes a 2D histogram of the stationary pixels in the x-y plane. The histogram is computed every day. Regions where the person sits or lies down will accumulate many stationary tags and hence will show up as peaks in the histogram.

Figure 3 shows example histograms from two of the homes where we deployed EZ-Sleep. The locations of EZ-Sleep are shown as the blue rectangles and the black lines indicate walls or boundaries of beds or desks. The darker the pixels are the more stationary time is spent in that location. The figure reveals three observations: First, the distribution of stationary pixels does indeed reveal the location of actual seating or sleeping areas. For example, the dark regions in the figure refer to real beds, couches and desks. Second, EZ-Sleep can detect more than one bed and can identify beds and seating areas outside the room where the device is installed. Third, the stationary pixels do not span the whole bed region. This is because most people sleep on a particular side of the bed. Besides, stationary pixels correspond to locations where there is some breathing signal –i.e., the chest, face, and abdomen. This description of the sleeping area is more relevant to our task than the whole physical space of the bed. Fourth and most importantly, the stationary regions can be connected in the histogram. For example, one of the beds in the figure appears connected to the desk. Simply by looking at the histogram, EZ-Sleep cannot tell whether this is one big bed or two distinct stationary areas. Thus, our next step is to separate the distinct seating or sleeping areas observed by the device.

In the second step, our algorithm takes the 2D histogram as input and extracts distinct seating or sleeping areas. We frame this problem as an image segmentation problem – i.e., we would like to assign each pixel a label such that pixels with the same label share the same area identity (i.e., desk, couch, or bed). Common approaches to image segmentation take initial labels of a few pixels, and finds boundaries dividing regions with different initial labels. Such initial labels are called markers. But, how can we assign initial labels to the images in Figure 3a

[2]Objects with mechanical vibration such as fans and air conditioners are detected and eliminated. They behave as outliers because their periodicity is much higher than breathing and they do not move from one pixel to another for multiple days.

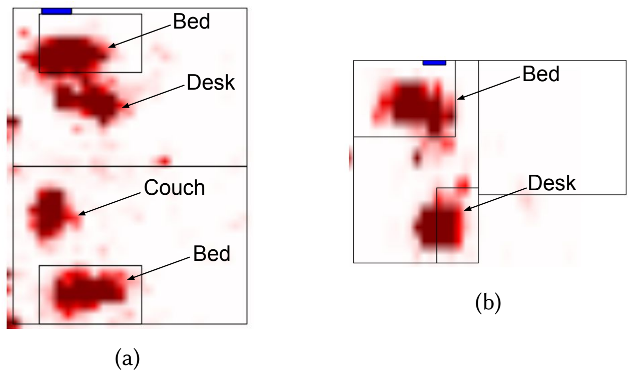

Fig. 3. Images of the distribution of stationary location measurements from two homes. The locations of EZ-Sleep are shown as the blue rectangles and the black lines indicate walls or boundaries of beds or desks. Dark regions indicate areas where people stay stationary for long periods of time, such as beds, desks or couches.

such that the bed and the desk have different labels automatically? Simply using local maxima as initial distinct labels will lead to over segmentation because each area can have multiple local maxima. Smoothing the image first does not help because it blurs region boundaries.

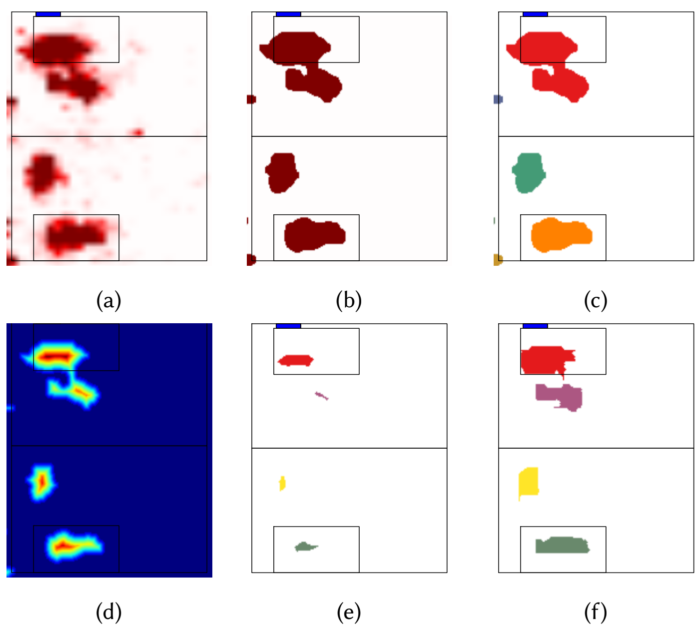

Fig. 4. Estimating potential bed areas. The image of the stationary location distribution is shown in (a). This image is binarized using Otsu thresholding, resulting in (b) where the red pixels form the foreground. To obtain the markers for watershed, we first compute the connected components as shown in (c). Computing a distance transform results in (d) which is thresholded to obtain the markers in (e). The watershed segmentation outputs the potential bed areas as shown in (f).

To assign initial labels for the segmentation algorithm, we use a processing pipeline explained in Figure 4. To estimate the initial labels, we binarize the input image (Figure 4a) into foreground and background (red and white regions in Figure 4b) using Otsu thresholding [38]. In the binarized image (Figure 4b), the desk and the bed are still connected. We can make the distinction between two areas bigger by computing the distance transform, i.e., the nearest distance to the background, for each pixel of this image. Given the transformed image (Figure 4d), we can apply a threshold to get the central regions of both the desk and the bed (Figure 4e). Note that to set a threshold that generalizes to different areas in different homes, we normalize the transformed image with respect to the size of each connected component (Figure 4c) before thresholding. After thresholding, the connected components of the thresholded image (Figure 4e) can be used as the initial labels for the segmentation algorithm. Figure 4f shows the final segmentation result using Watershed algorithm [34]. We reject regions whose area is too small to be a bed. The remaining regions are the potential bed areas.

The third and last step in our algorithm is to identify beds from other areas where the user may spend some time while stationary, like a couch or a desk. People may spend long time in areas other than beds. Using only the amount of time spent in an area as the classifying feature is not enough. Also, people's schedule changes from day to day. A user can spend more time on a couch watching TV than sleeping in the bed during the weekend. Similarly, before a deadline a user can spend more time working at her desk than sleeping in her bed. The task is further complicated by the fact that we want to discover multiple beds in the same house and different people may have very different schedules.

To address these problems, we analyze the temporal activities in each potential bed area to extract additional features and look at the consistency of features across multiple days. Specifically, in addition to the amount of time the user stays in an area, we also use the ratio of time being stationary versus moving. To compute this feature, we take the sum of all stationary tags for all pixels in the area and divide it by the sum of all moving tags for all pixels in the area. Figure 5a visualizes around two weeks of data for 10 potential bed areas identified at the final output of the segmentation pipeline in Figure 4. The two axes refer to the two features: total amount of time, and ratio of stationary time to moving time spent in that area. Each day contributes one data point to each detected area. The circles refer to real beds, whereas the triangles refer to couch and desk areas. We remove data points where people spent zero amount of time from this figure. We can see that the desk and couch areas form a cluster in the lower left corner. There are outliers as a person may sleep less on some days or a couch may look like a bed occasionally. To better understand this, we highlight data from one bed region and one desk region in Figure 5b. If we look across multiple days, data from the bed and the desk stay in their respective clusters.

Given the analysis, we use the above two features to train a bed classifier using SVM with a linear kernel. To determine if a region is a bed, we look at the its predicted labels over the past D days. We identify the region as bed if the percentage of days labeled as bed is greater than γ:

$$\hat{y_i} = \begin{cases} Bed, & if \ \frac{\sum_{d=0}^{D-1} \|f(A_i;d)\|}{D} > \gamma \\ Not\ Bed, & otherwise \end{cases}, \tag{1}$$

where $\hat{y_i}$ is the final prediction for area A_i, $f(.)$ is the classifier, $f(A_i;d)$ is the binary predicted results for area A_i on the dth day. Our default is $D = 7$ and $\gamma = 5/7$. Figure 5 shows the classification boundary of our classifier.

6 DETECTING BED ENTRY AND EXIT

After identifying the bed areas, we consider the problem of detecting when a user goes to bed and when she leaves the bed. RF-based localization returns the location of the people in the environment. However, as mentioned in Section 4 their errors can be as large as one meter. An error of plus or minus one meter is pretty much the size of the bed. It can easily move a person from the bed to outside the bed. Similarly, a person who is changing his clothes next to the bed may look already in bed. These localization errors have little impact on detecting the

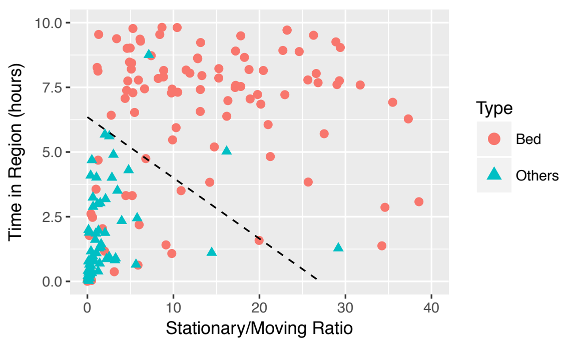

(a) Figure shows the two features for 10 potential bed areas. Using the two features, we can separate the bed and non-bed areas.

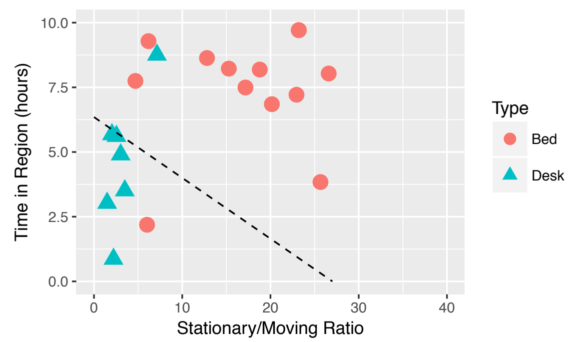

(b) Figure shows the two features for a specific bed and desk region. We use data across multiple days to classify a region as bed or non-bed.

Fig. 5. Identifying actual beds.

bed region because in that case we can use the histogram of location data over a whole day, which allows us to average out the errors. In contrast, when detecting entries and exits from the bed, we would like to detect them as quickly as possible because any delay in detecting such event will appear as an error in our estimate of sleep latency.

So, how do we accurately track when the user enters or leaves her bed? Instead of directly mapping the measured locations to in or out of the bed, EZ-Sleep considers the location measurements as noisy observations of the true state. The true state –i.e., whether the user is in or out of bed – is hidden. EZ-Sleep uses a Hidden Markov Model (HMM) to infer the true state from the observations. Below we give a short primer on HMM followed by our particular HMM design.

6.1 HMM Background

A hidden Markov model is a statistical model that tries to explain a sequence of observations with a sequence of hidden states. Time progresses in steps. In each time step, the model is presented with a new observation which it uses to decide whether to stay in the current state or transition to some other state. The key assumption underlying an HMM is that the probability of an observation given the current state is independent of any other observation or state. Mathematically, let us model the observation at time, t, as o_t and the hidden state as s_t. If $P(X)$ denotes the probability of event X, then the assumption can be represented as:

$$P(o_t|s_0, s_1, s_2, ..., s_{t-1}, s_t) = P(o_t|s_t) \tag{2}$$

Thus, the joint probability of having a sequence of states, $\{s_t\}_{t=1}^{N_t}$ and observations, $\{o_t\}_{t=1}^{N_t}$ is given by:

$$P(\{o_t, s_t\}_{t=1}^{N_t}) = P(s_0)P(o_0|s_0)\prod_{t=1}^{N_t} P(o_t|s_t)P(s_t|s_{t-1}) \tag{3}$$

To model a problem as an HMM one has to define a set of possible states $\mathcal{S} = \{\mathcal{S}_1, ..., \mathcal{S}_M\}$, a set of observations $\mathcal{O} = \{\mathcal{O}_1, ..., \mathcal{O}_N\}$, a $M \times M$ transition probability matrix $\mathbf{T}$, and a $M \times N$ emission probability matrix $\mathbf{E}$. Transition probability is the probability of transitioning from state, $\mathcal{S}_i$ to $\mathcal{S}_j$ in consecutive time steps, i.e., $\mathbf{T}_{ij} = P(s_t = \mathcal{S}_i|s_{t-1} = \mathcal{S}_j)$. The emission probability is the probability of an observation given a state, i.e, $\mathbf{E}_{ij} = P(o_t = \mathcal{O}_j|s_t =$

$\mathcal{S}_i$). The set of observations and states is typically picked by the designer, and the transmission probabilities and emission probabilities are learned from the data.

6.2 Design of Our HMM

We use the HMM hidden states to represent whether the user is in the bed or outside the bed at each time step t. We denote the state corresponding to being in the bed as $\mathcal{S}_0$ and outside the bed as $\mathcal{S}_1$.

We define our observations by first dividing the space around the bed region into three areas: center area (R_0), buffer area (R_1), and outer area (R_2). The center area is the bed region we found in section 5, and the buffer area is a 50-centimeter-wide area that encircles the center area. The outer area is the rest of the space. We define 9 possible observations which correspond to the user transitioning from one of the three areas to another or staying in the same area. Thus, the observations for the HMM are the set of tuples

$$O = \{(R_s, R_e)\ \forall R_s, R_e \in R_0, R_1, R_2\}\,.$$

To detect which observation has occurred in each time step, we divide time into consecutive windows of 5 seconds each. For each window we detect the location of the user from the moving and stationary pixels. We average the location in each second and map it to one of the three area: R_0, R_1, or R_2. Then R_s denotes the area that corresponds to the first second of the window while R_e is the area that correspond to the last second of the window.

We specify the buffer area around the bed because we are often unsure of the person's exact location. To decide a person's state, we need to look a bit into the past and the future. Having the buffer area helps us encode different beliefs in these observations as opposed to simply dividing the space into inside versus outside the bed. Given the learned bed regions from the last section, this choice of observations does not require any information specific to the home. Thus, we can use the same HMM for all homes. In addition, we transform raw location data into discrete observations for better generalization of our model. If we use a continuous input instead (e.g., the distance to the center of the bed), because people stay at different locations away from beds in different homes, the model trained at few places cannot generalize well to different homes.

If multiple bed regions are detected, we generate different observations for each bed. EZ-Sleep uses mechanisms similar to [11] to track raw locations of users to form observations for the HMM.

Learning and Inference: The goal of the learning algorithm is to learn the transition probabilities matrix, T, and the emission probabilities matrix, E. In the absence of priors, the typical way to learn T and E is to compute their statistics using labeled data. Thus, we manually label the training data with the correct states and use them to learn T and E.

After the learning phase, the model can be used independent of the environment given the bed region. In the inference phase, we give the HMM the sequence of observations for one day and find a sequence of states that best explain those observations using the Viterbi algorithm [21]. This means that we can give the HMM the observations extracted from the areas around each detected bed, and have the HMM predict for us when a person enters the bed (i.e., the state transition from $\mathcal{S}_1$ to $\mathcal{S}_0$) and when she exits the bed (a transition from $\mathcal{S}_0$ to $\mathcal{S}_1$).

7 ESTIMATING SLEEP PARAMETERS

So far, we have learned the bed area and when the user enters and leaves the bed. To estimate the sleep parameters in Section 2, we still need to detect when the user falls asleep, and the sleep-awake intervals throughout a night. To do so, we divides the time in bed into 30-second epochs, as typical in sleep studies. We would like to classify each epoch into sleep vs. awake and identify the first epoch in which the person falls asleep, which is the sleep onset.

The information of whether a person is asleep or awake is encoded in her breathing pattern and body movements [23, 39, 44]. To extract this information, we need to ensure that the RF signal that we provide to the classifier, in each epoch, reflects the user's breathing and movements and is filtered from any source of interference. This is challenging because there are many other sources of motion in the environment, e.g., rotating fans, neighbors whose movements may impact the signal, and even curtains moving due to air conditioning. These extraneous movements can impact the signal and completely overshadow the small movements caused by human breathing. As we show in Section 9.1, having the neighbor walking in another room from behind the wall can easily interfere with the breathing signal.

To separate breathing from extraneous interference, EZ-Sleep uses the identified bed areas as spatial filters to filter out motion coming from outside the bed. Specifically, we zoom in on the pixels in the bed areas and analyze only the signals reflected from those pixels. For each epoch, we consider the signals reflected from pixels tagged as stationary or moving (i.e., non empty). We use these signals as estimates of the user's breathing pattern and body movements.

For our classifier, we build on recent success of convolutional neural network (CNN) in classifying time signals [9]. The benefit of this classifier is that we do not need to pick features manually. We can directly pass the signal in each epoch to the classifier. We build a CNN classifier with a 14-layer residual network model [26]. Our network architecture is similar to the 18-layer network described in [26] but with no repetition for the conv4 and conv5 layers. The classifier takes the spectrogram of the signal in each epoch as input, and outputs the probability of the person being asleep. Hence for each night, we obtain a series of probabilities $\{p_i\}_n$, where n is the total number of epochs for that night, and i is the i_{th} epoch. We use this CNN classifier as a building block for estimating sleep onset time, Total sleep time (TST), and Wake after sleep onset (WASO).

Sleep Onset Time: Predicting the onset of sleep, using the CNN classifier alone would not give us enough accuracy. The reason is that in training the CNN, the model tries to maximize the overall accuracy of predicting asleep-awake for the entire night, but not the accuracy of detecting the exact sleep onset time. Thus, we use this first CNN classifier as the first phase in detecting sleep onset. Specifically, for each night, we consider the first epoch in which the probability of sleep p_i is larger than 0.5. We then take a window of 15 minutes before and 15 minutes after that epoch. This gives us the sequence of epochs around the sleep onset time.

We build a Gradient Boosting Regressor on top of the above CNN with special focus on the sleep onset transition. The regressor takes as input the signal in the epochs in the above window to learn the exact sleep onset. In this case, we would like to learn which epoch is the sleep onset. Since this function is an impulse function, it cannot be directly learned by the regressor. It has to be smoothed first. Thus, we smooth it by convolving it with a Gaussian kernel. For each epoch, the regressor predicts the probability of being the sleep onset epoch. The epoch with the largest probability score is considered the sleep onset epoch.

TST and WASO: With the predicted sleep probabilities $\{p_i\}_n$ and sleep onset time κ, we can estimate TST and WASO. Specifically, TST is the total duration of all epochs for which $p_i > 0.5$, starting after epoch κ. Similarly, WASO is the total time durations of all epochs for which $p_i \leq 0.5$ starting after epoch κ.

SE and SL: Finally, sleep efficiency (SE) is computed directly as TST/TIB, and sleep latency (SL) is the time difference between the last entry to bed before epoch κ and the beginning of epoch κ.

8 EVALUATION SETUP

We evaluated EZ-Sleep through actual deployments in 8 homes. We collected more than 100 nights of sleep data from 10 healthy subjects whose age spans 23 to 45 years. Out of these, all subjects slept with SleepProfiler [1, 43], which is an FDA-approved medical grade sleep monitoring device, for a total of 30 nights to obtain the ground truth for our sleep parameters. We install the EZ-Sleep device in the bedroom of the subject. Figure 7 shows the location of the device in all 8 homes. In our deployments, two of the homes have two beds in the device's coverage area. All subjects have been consented in accordance with our IRB.

SleepProfiler (used to obtain the ground-truth) has three frontal EEG electrodes to measure brain activity, accelerometers for detecting motion, a chest band for breathing, and a pulse rate sensor for monitoring heart rate. The device has a push button for the user to indicate when she goes to bed and when she leaves the bed. The subjects are instructed to push the button when they enter and leave the bed. Other parameters can be extracted directly from the sleep report provided by the device. The subject may make mistakes in setting up the sleep profiler or wearing the sensors. For example, they may forget to attach the adhesives that prevent the EEG electrodes from moving, or remove the sensors and go back to sleep. To ensure that all nights considered in the study do not have such errors, we ask the subjects to keep a sleep diary, in which they record when they go to bed and when they wake up, whether they experienced any awakenings during the night, and other comments that regarding removal of the device. The diary is used to check for consistency with the sleep profiler data and exclude inconsistent nights.

The ground truth data is also used to train the classifiers. Training and testing are done on different people and different homes −i.e., for each home, we test a classifier trained on the other homes.

9 RESULTS

Below we start by showing that EZ-Sleep is more robust to interference than past solutions that use RF-based signals but rely on the Doppler effect. We then evaluate EZ-Sleep's ability to detect the bed region and extract the sleep parameters.

9.1 Sensor Robustness

There are two benefits for using the spatial information extracted from RF-signals. The first is the ability to detect where and when the user goes to bed. The second, is the fact that we can separate the RF reflections from different pixels in space. This latter property allows EZ-Sleep to eliminate interference from other sources of motion, such as fans, neighbors, etc. In this section, we demonstrate empirically the added robustness due to this spatial filtering. We also provide mathematical reasons why Doppler-based solutions common in past work [5, 31, 32, 39], can be easily confused in the presence of extraneous motion.

We did an experiment where we ask a person to lie on a bed in the same room as the radio device. We ask the person to be stationary, so that we can focus only on his breathing motion. The RF signal is measured both using the EZ-Sleep radio and a similar radio, where the only difference is that we replace the FMCW and antenna array setup with measuring the Doppler effect as described in [39]. Figure 6a shows that both the Doppler radar and EZ-Sleep capture the person's breathing when he is the only source of motion in the environment. We repeat the same experiment, but this time with another person present in an adjacent room. Figure 6b shows the results. The second person is stationary in the first half of the experiment, and starts walking around the 12th second. For the Doppler-based approach, even before the second person starts walking, his presence disturbs the signal making it hard to track the breathing of the person in the bed. This gets worse when the person starts walking, as shown in the second half of Figure 6b. In contrast, EZ-Sleep's signal stays clean and clearly reflects the breathing of the subject in the bed.

The above results can be explained as follows. A Doppler-based sensor transmits a single frequency and detects the chest displacement $x_0(t)$ by looking at the baseband signal $B(t)$ at the receiver [39]:

$$B(t) = A_0 cos(\frac{4\pi x_0(t)}{\lambda} + C_0), \tag{4}$$

where A_0 is proportional to the amplitude of the reflection, λ is the wavelength, and C_0 is a constant that depend on the person's distance from the device. The chest motion $x_0(t)$ can be inferred from the phase of the signal $B(t)$.

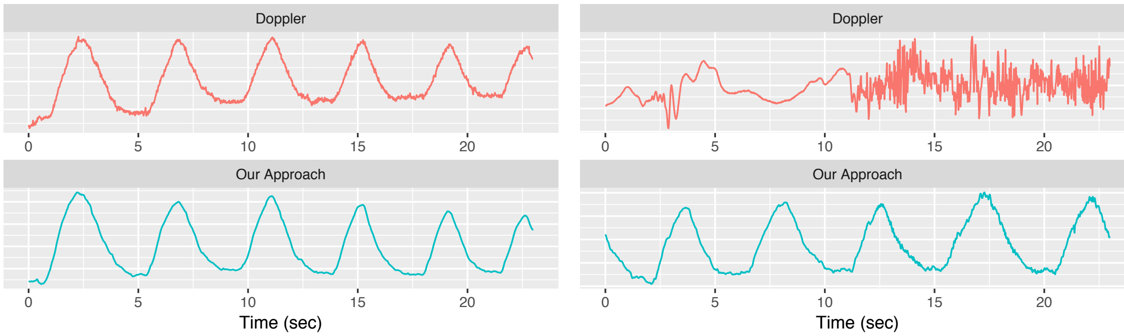

(a) Single Person Scenario

(b) Two Person Scenario

Fig. 6. EZ-Sleep vs. Doppler based approaches. The signal using the Doppler effect gets disturbed in the presence of other motion in the environment.

However, in the presence of another source of motion with displacement $x_1(t)$, the baseband signal becomes:

$$\widetilde{B}(t) = A_0 cos(\frac{4\pi x_0(t)}{\lambda} + C_0) + A_1(t) cos(\frac{4\pi x_1(t)}{\lambda} + C_1(t)) \tag{5}$$

where $A_1(t)$ is proportional to the amplitude of the second reflection and $C_1(t)$ is related to the second person's distance. Without knowing how the second person moves, both $A_1(t)$ and $C_1(t)$ are unknown, and one can no longer infer $x_0(t)$. Moreover, since both terms could change over time, estimating the periodicity of $x_0(t)$ becomes challenging. Hence, Doppler-based approaches cannot estimate the breathing rate accurately in this scenario. In contrast, combining FMCW with an antenna array (Section 4) allows EZ-Sleep to separate the two reflections using the fact that they come from different spatial pixels. Therefore, we can extract cleaner breathing signal, as shown in Figure 6b, even in the presence of other people.

9.2 Identifying Bed Areas

We compare the bed areas identified by EZ-Sleep with the ground truth location of the bed. The ground truth is obtained using careful measurements of the location of the bed and the layout of each home using laser distance meters with an accuracy of 0.06". To train our bed model we divide the data into two sets, each cover 4 homes. We train the model on one set and test it on the other. We then swap the training and test sets. Thus, the model is asked to predict the bed area for new homes that it did not see in the training phase.

Figure 7 shows the results of our bed identification. The figure shows the floor plans from all 8 homes. The location of the EZ-Sleep device is illustrated using a blue rectangle. We label the ground truth bed areas using the red boxes and show the bed areas identified by EZ-Sleep in dark green. Note that in the homes in (c) and (g), EZ-Sleep monitor two beds simultaneously.

As shown in Figure 7, EZ-Sleep correctly identified the bed areas for all 8 homes, even at places where there are two beds in the coverage area (7c and 7g). Note that the identified bed area usually covers a subregion of the physical bed. As explained in Section 5, these areas are obtained using the location data when the person is stationary and breathing is the only motion. These areas represent the chest, abdomen and face of the person. Having this ability to accurately estimate the spacial pixels that reflect the user's breathing allows us to zoom in on the subject during their sleep to estimate the sleep parameters avoiding a scan of a larger 3D space.

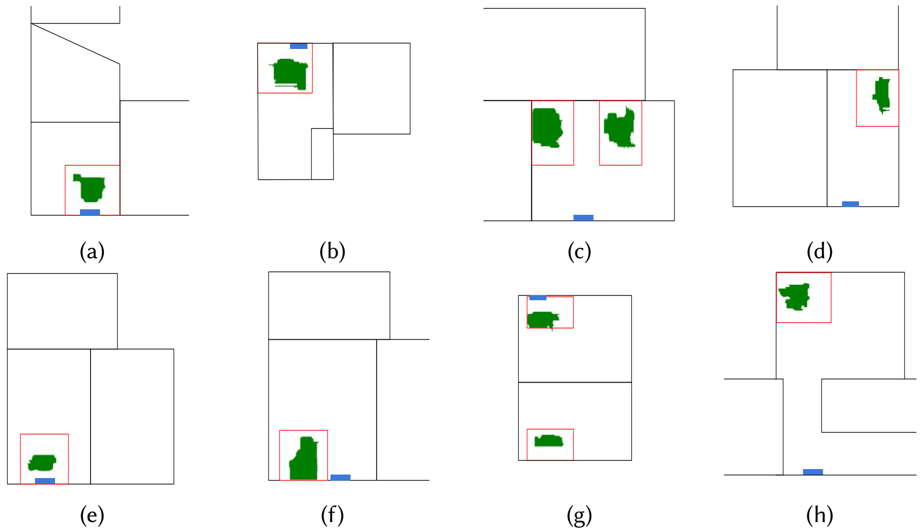

Fig. 7. Bed areas identified by EZ-Sleep. Locations of EZ-Sleep are illustrated with blue rectangles. The red boxes represent the ground truth bed areas and the green areas show the identified bed areas.

9.3 Accuracy of Sleep Parameters

In this section, we evaluate EZ-Sleep's accuracy of detecting bed entry and exit events, and estimating the various sleep parameters. We distinguish scenarios in which the sensor monitors one subject and one bed from scenarios in which the sensor monitors two subjects in two beds.

9.3.1 Single User Scenarios. Let us start by reporting the accuracy of detecting bed entry and exit events and the resulting time in bed (TIB). As explained earlier, EZ-Sleep uses an HMM based approach to identify bed entries and exits. The average error of the HMM in detecting bed entries and exits are 1.8 minutes and 1.3 minutes, respectively. The entries and exits are used to compute the TIB, which has an average error of 3.15 minutes.

To show the individual errors for all subjects and all nights, we show in Figure 8 a scatter plot of the predicted TIB values versus the actual TIB values. The TIB values in our dataset range from 273 minutes to 564 minutes. Most of the points in Figure 8 lie on the diagonal line. This shows that EZ-Sleep achieves high accuracy of TIB estimation across a wide range of sleep schedules.

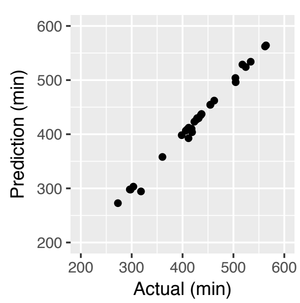

Fig. 8. Predicted versus actual Time in Bed.

Next, we compare the other sleep parameters predicted by EZ-Sleep to their ground truth values. Figure 9 shows scatter plots of the predicted sleep parameters vs. their ground truth values taken over all subjects and all nights. Table 1 summarizes the statistics in the scatter plots and presents the average, mean and standard deviation of the prediction error, for the various sleep parameters. The table shows that EZ-Sleep has high accuracy. Specifically, its average error in predicting TST, SL, SE, and WASO is 10.3 min, 4.9 min, 2.8%, and 8.2 min, respectively. These results are comparable to medical grade actigraphy-based insomnia monitors and within the clinically meaningful ranges [30].

Sleep Parameter	Average Error	Median Error	Standard Deviation of Error
Time in Bed (TIB)	3.15 (min)	0.14 (min)	6.11 (min)
Total Sleep Time (TST)	10.3 (min)	8.5 (min)	7.7 (min)
Sleep Latency (SL)	4.9 (min)	4.3 (min)	3.1 (min)
Sleep Efficiency (SE)	2.8 (%)	2.6 (%)	2.1 (%)
Wake After Sleep Onset (WASO)	8.2 (min)	6.2 (min)	8.5 (min)

Table 1. Accuracy of Sleep Parameters for Single User Scenarios

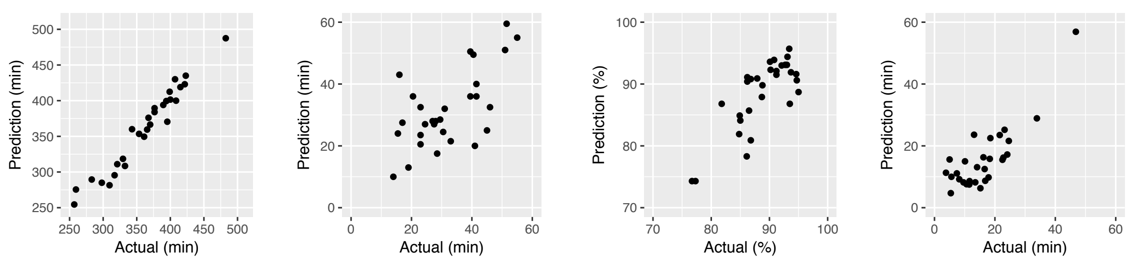

(a) Total Sleep Time (b) Wake After Sleep Onset (c) Sleep Efficiency (d) Sleep Latency

Fig. 9. Scatter plots of the predicted sleep parameters versus their ground truth values.

9.3.2 Monitoring Multiple Users. Finally, we show that EZ-Sleep can simultaneously monitor multiple subjects sleeping in their corresponding beds. We have two homes where the RF signal covers two beds, as shown in Figure 7c and 7g. Table 2 presents the error statistics for simultaneously monitoring two people sleeping in the same home. The table shows that the errors in TST, SL, and WASO are a few minutes higher than in the case of a single person. This is expected because while mechanisms like antenna array separate signals from different directions, the separation is not perfect (due to the array's side lobes). Thus, the presence of a second person can add a small disturbance to the signal from the other person, particularly if the two people are close to each other. However, the errors are still small and comparable to medical grade actigraphy [30]. This result is particularly interesting since it shows that EZ-Sleep is the first RF-based sleep sensor that is capable of monitoring multiple people simultaneously.

10 CONCLUSION

We have presented EZ-Sleep, a new approach to monitoring insomnia and sleep deficiencies. EZ-Sleep is both accurate and easy to use. All that the user has to do is to put the device in her bedroom and plug it to the power

Sleep Parameter	Average Error	Median Error	Standard Deviation of Error
Time in Bed (TIB)	0.2 (min)	0.1 (min)	0.2 (min)
Total Sleep Time (TST)	15.8 (min)	15.5 (min)	11.8 (min)
Sleep Latency (SL)	7.6 (min)	7.0 (min)	6.4 (min)
Sleep Efficiency (SE)	1.8 (%)	2.2 (%)	1.2 (%)
Wake After Sleep Onset (WASO)	13.1 (min)	13.5(min)	12.1 (min)

Table 2. Accuracy of Sleep Parameters for Multi User Scenarios.

outlet. EZ-Sleep automatically figures out where the user sleeps, and continuously monitors her sleep parameters. It has high accuracy in monitoring sleep latency, sleep efficiency, total sleep time, time in bed, and wake after sleep onset. Furthermore, one device can monitor multiple users sleeping in different beds.

We believe EZ-Sleep provides an important improvement over the state-of-the-art insomnia monitoring. We expect future research to continue to expand the capabilities of RF-based monitoring devices. In particular, in this paper we associate the identities of users with the beds they sleep on. One could incorporate past work that identifies users from radio reflections [10, 47] with EZ-Sleep to deal with cases where users do not sleep in the same place across multiple nights. Also, we only looked at scenarios where each user sleep alone in the bed. More advances can lead to accurate sleep monitoring of multiple users in the same bed using radio signals. Finally, our experiments were conducted on healthy subjects. Future work should experiment with more diverse populations and might integrate monitoring with treatment.

ACKNOWLEDGMENTS

The authors thank the members of the NETMIT group for their feedbacks, the reviewers for their insightful comments, and the volunteers for participating in our study.

REFERENCES

[1] 2017. Advanced Brain Monitoring - Sleep Profiler. (2017). Retrieved May 12, 2017 from http://www.advancedbrainmonitoring.com/sleep-profiler/
[2] 2017. Beddit. (2017). Retrieved April 23, 2017 from http://www.beddit.com/
[3] 2017. Fitbit. (2017). Retrieved April 23, 2017 from https://www.fitbit.com/
[4] 2017. Jawbone. (2017). Retrieved April 23, 2017 from https://jawbone.com/
[5] 2017. S-Plus. (2017). Retrieved April 23, 2017 from http://splus.resmed.com/
[6] 2017. Sleep Cycle. (2017). Retrieved April 23, 2017 from https://www.sleepcycle.com/
[7] 2017. Sleep Profiler. (2017). Retrieved April 23, 2017 from http://www.advancedbrainmonitoring.com/sleep-profiler/
[8] 2017. Sleepbot. (2017). Retrieved April 23, 2017 from https://mysleepbot.com/
[9] Ossama Abdel-Hamid, Abdel-rahman Mohamed, Hui Jiang, Li Deng, Gerald Penn, and Dong Yu. 2014. Convolutional neural networks for speech recognition. *IEEE/ACM Transactions on audio, speech, and language processing* 22, 10 (2014), 1533–1545.
[10] Fadel Adib, Chen-Yu Hsu, Hongzi Mao, Dina Katabi, and Frédo Durand. 2015. Capturing the human figure through a wall. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 219.
[11] Fadel Adib, Zachary Kabelac, and Dina Katabi. 2015. Multi-person localization via RF body reflections. In *Proceedings of the 12th USENIX Conference on Networked Systems Design and Implementation*. USENIX Association, 279–292.
[12] Fadel Adib, Zachary Kabelac, Dina Katabi, and Robert C Miller. 2014. 3D Tracking via Body Radio Reflections.. In *NSDI*, Vol. 14. 317–329.
[13] Fadel Adib, Hongzi Mao, Zachary Kabelac, Dina Katabi, and Robert C Miller. 2015. Smart homes that monitor breathing and heart rate. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 837–846.
[14] S Ancoli-Israel, R Cole, C Alessi, M Chambers, W Moorcroft, and C Pollak. 2003. The role of actigraphy in the study of sleep and circadian rhythms. American Academy of Sleep Medicine Review Paper. *Sleep* 26, 3 (2003), 342–392.
[15] Daniel J Buysse, Yu Cheng, Anne Germain, Douglas E Moul, Peter L Franzen, Mary Fletcher, and Timothy H Monk. 2010. Night-to-night sleep variability in older adults with and without chronic insomnia. *Sleep medicine* 11, 1 (2010), 56–64.

[16] Mary A Carskadon, William C Dement, MM Mitler, Christian Guilleminault, Vincent P Zarcone, and Rene Spiegel. 1976. Self-reports versus sleep laboratory findings in 122 drug-free subjects with complaints of chronic insomnia. *Am J Psychiatry* 133, 12 (1976), 1382–1388.

[17] Zhenyu Chen, Mu Lin, Fanglin Chen, Nicholas D Lane, Giuseppe Cardone, Rui Wang, Tianxing Li, Yiqiang Chen, Tanzeem Choudhury, and Andrew T Campbell. 2013. Unobtrusive sleep monitoring using smartphones. In *Pervasive Computing Technologies for Healthcare (PervasiveHealth), 2013 7th International Conference on*. IEEE, 145–152.

[18] Eun Kyoung Choe, Bongshin Lee, Matthew Kay, Wanda Pratt, and Julie A Kientz. 2015. SleepTight: low-burden, self-monitoring technology for capturing and reflecting on sleep behaviors. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 121–132.

[19] Roger J Cole, Daniel F Kripke, William Gruen, Daniel J Mullaney, and J Christian Gillin. 1992. Automatic sleep/wake identification from wrist activity. *Sleep* 15, 5 (1992), 461–469.

[20] Luciane de Souza, Ana Amélia Benedito-Silva, Maria Laura Nogueira Pires, Dalva Poyares, Sergio Tufik, and Helena Maria Calil. 2003. Further validation of actigraphy for sleep studies. *Sleep* 26, 1 (2003), 81–85.

[21] G David Forney. 1973. The viterbi algorithm. *Proc. IEEE* 61, 3 (1973), 268–278.

[22] Kristina Grifantini. 2014. How's My Sleep?: Personal sleep trackers are gaining in popularity, but their accuracy is still open to debate. *IEEE pulse* 5, 5 (2014), 14–18.

[23] Weixi Gu, Zheng Yang, Longfei Shangguan, Wei Sun, Kun Jin, and Yunhao Liu. 2014. Intelligent sleep stage mining service with smartphones. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 649–660.

[24] Reinder Haakma and R Beun. 2012. Unobtrusive Sleep Monitoring. In *Measuring Behavior*, Vol. 2012. Citeseer, 122.

[25] Tian Hao, Guoliang Xing, and Gang Zhou. 2013. iSleep: unobtrusive sleep quality monitoring using smartphones. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*. ACM, 4.

[26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 770–778.

[27] Matthew Kay, Eun Kyoung Choe, Jesse Shepherd, Benjamin Greenstein, Nathaniel Watson, Sunny Consolvo, and Julie A. Kientz. 2012. Lullaby: A Capture & Access System for Understanding the Sleep Environment. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12)*. ACM, New York, NY, USA, 226–234. https://doi.org/10.1145/2370216.2370253

[28] Anastasi Kosmadopoulos, Charli Sargent, David Darwent, Xuan Zhou, and Gregory D Roach. 2014. Alternatives to polysomnography (PSG): a validation of wrist actigraphy and a partial-PSG system. *Behavior research methods* 46, 4 (2014), 1032–1041.

[29] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. 2015. Spotfi: Decimeter level localization using wifi. In *ACM SIGCOMM Computer Communication Review*, Vol. 45. ACM, 269–282.

[30] Kenneth L Lichstein, Kristen C Stone, James Donaldson, Sidney D Nau, James P Soeffing, David Murray, Kristin W Lester, and R Neal Aguillard. 2006. Actigraphy validation with insomnia. *SLEEP-NEW YORK THEN WESTCHESTER-* 29, 2 (2006), 232.

[31] Feng Lin, Yan Zhuang, Chen Song, Aosen Wang, Yiran Li, Changzhan Gu, Changzhi Li, and Wenyao Xu. 2016. SleepSense: A Noncontact and Cost-Effective Sleep Monitoring System. *IEEE Transactions on Biomedical Circuits and Systems* (2016).

[32] Xuefeng Liu, Jiannong Cao, Shaojie Tang, and Jiaqi Wen. 2014. Wi-Sleep: Contactless sleep monitoring via WiFi signals. In *Real-Time Systems Symposium (RTSS), 2014 IEEE*. IEEE, 346–355.

[33] National Heart Lung and Blood Institute (NHLBI). 2012. Why Is Sleep Important? (Feb. 2012). Retrieved April 23, 2017 from https://www.nhlbi.nih.gov/health/health-topics/topics/sdd/why

[34] Fernand Meyer. 1992. Color image segmentation. In *Image Processing and its Applications, 1992., International Conference on*. IET, 303–306.

[35] Jun-Ki Min, Afsaneh Doryab, Jason Wiese, Shahriyar Amini, John Zimmerman, and Jason I Hong. 2014. Toss'n'turn: smartphone as sleep and sleep quality detector. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 477–486.

[36] Standards of Practice Committee of the American Academy of Sleep Medicine and others. 2003. Practice parameters for the role of actigraphy in the study of sleep and circadian rhythms: an update for 2002. *Sleep* 26, 3 (2003), 337–341.

[37] Maurice M Ohayon. 2002. Epidemiology of insomnia: what we know and what we still need to learn. *Sleep medicine reviews* 6, 2 (2002), 97–111.

[38] Nobuyuki Otsu. 1975. A threshold selection method from gray-level histograms. *Automatica* 11, 285-296 (1975), 23–27.

[39] Tauhidur Rahman, Alexander T Adams, Ruth Vinisha Ravichandran, Mi Zhang, Shwetak N Patel, Julie A Kientz, and Tanzeem Choudhury. 2015. Dopplesleep: A contactless unobtrusive sleep sensing system using short-range doppler radar. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 39–50.

[40] Brant W Riedel, Carolyn F Winfield, and Kenneth L Lichstein. 2001. First night effect and reverse first night effect in older adults with primary insomnia: does anxiety play a role? *Sleep medicine* 2, 2 (2001), 125–133.

[41] Deepak Shrivastava, Syung Jung, Mohsen Saadat, Roopa Sirohi, and Keri Crewson. 2014. How to interpret the results of a sleep study. *Journal of community hospital internal medicine perspectives* 4, 5 (2014).

[42] Simon Smith and John Trinder. 2001. Detecting insomnia: comparison of four self-report measures of sleep in a young adult population. *Journal of sleep research* 10, 3 (2001), 229–235.

[43] C Stepnowsky, D Levendowski, D Popovic, I Ayappa, and DM Rapoport. 2013. Scoring accuracy of automated sleep staging from a bipolar electroocular recording compared to manual scoring by multiple raters. *Sleep medicine* 14, 11 (2013), 1199.

[44] Alexander Tataraidze, Lyudmila Korostovtseva, Lesya Anishchenko, Mikhail Bochkarev, Yurii Sviryaev, and Sergey Ivashov. 2016. Bioradiolocation-based sleep stage classification. In *Engineering in Medicine and Biology Society (EMBC), 2016 IEEE 38th Annual International Conference of the*. IEEE, 2839–2842.

[45] Deepak Vasisht, Swarun Kumar, and Dina Katabi. 2016. Decimeter-level localization with a single wifi access point. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*. USENIX Association, 165–178.

[46] Hao Wang, Daqing Zhang, Junyi Ma, Yasha Wang, Yuxiang Wang, Dan Wu, Tao Gu, and Bing Xie. 2016. Human respiration detection with commodity wifi devices: do user location and body orientation matter?. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 25–36.

[47] Wei Wang, Alex X Liu, and Muhammad Shahzad. 2016. Gait recognition using wifi signals. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 363–373.

[48] L Welstein, WC Dement, D Redington, C Guilleminault, and MM Mitler. 1983. Insomnia in the San Francisco Bay area: a telephone survey. In *Sleep/wake disorders: natural history, epidemiology, and long-term evolution*. Raven Press New York, 73–85.

[49] Mingmin Zhao, Shichao Yue, Dina Katabi, and Tommi S Jaakkola. 2017. Learning Sleep Stages from Radio Signals: A Deep Adversarial Architecture. In *The 34th International Conference on Machine Learning (ICML)*.