SleepMonitor: Monitoring Respiratory Rate and Body Position During Sleep Using Smartwatch

XIAO SUN, LI QIU, YIBO WU, YEMING TANG, and GUOHONG CAO, The Pennsylvania State University, University Park.

Respiratory rate and body position are two major physiological parameters in sleep study, and monitoring them during sleep can provide helpful information for health care. In this paper, we present SleepMonitor, a smartwatch based system which leverages the built-in accelerometer to monitor the respiratory rate and body position. To calculate respiratory rate, we design a filter to extract the weak respiratory signal from the noisy accelerometer data collected on the wrist, and use frequency analysis to estimate the respiratory rate from the data along each axis. Further, we design a multi-axis fusion approach which can adaptively adjust the estimates from the three axes and then significantly improve the estimation accuracy. To detect the body position, we apply machine learning techniques based on the features extracted from the accelerometer data. We have implemented our system on Android Wear based smartwatches and evaluated its performance in real experiments. The results show that our system can monitor respiratory rate and body position during sleep with high accuracy under various conditions.

CCS Concepts: •**Human-centered computing** →*Ubiquitous and mobile computing systems and tools; Mobile devices;*

Additional Key Words and Phrases: Smartwatch, respiratory rate, body position, filtering, fusion

ACM Reference format:
Xiao Sun, Li Qiu, Yibo Wu, Yeming Tang, and Guohong Cao. 2017. SleepMonitor: Monitoring Respiratory Rate and Body Position During Sleep Using Smartwatch. *PACM Interact. Mob. Wearable Ubiquitous Technol.* 0, 0, Article 0 (2017), 22 pages.
DOI: 10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Sleep plays a significant role in maintaining good physiological and emotional health. In sleep study, respiratory rate and body position are two important and useful physiological parameters. Respiratory rate can be used to assess the sleep quality [6] and to predict some sleep-related diseases [12, 21]. Body position is also a significant metric for sleep quality assessment [7], and it is related to many medical conditions such as apnea [22] and back pain [8]. Thus, monitoring the respiratory rate and body position during a user's sleep is very important and can provide helpful information for health care.

Traditionally, sleep parameters such as respiratory rate and body position are monitored in the hospital by using polysomnography (PSG) [29]. By collecting data from a lot of sensors and electrodes worn on various parts of a user's body, PSG is able to provide fine-grained sleep data. However, using PSG

DOI: 10.1145/nnnnnnn.nnnnnnn

involves hospital visit, specialized wearable sensors and professional installation, which is impractical for long-term sleep monitoring and limited only to clinical usage.

With the rapid development of wearable devices and wireless technologies, a lot of smart health applications have been proposed [19, 20, 27, 28], including sleeping respiratory rate and body position monitoring [1, 2, 14, 15, 17, 18, 23]. Recently, commercial wrist-worn devices such as Fitbit or Jawbone have been released, which are able to track users' sleep. However, they can only record some coarse-grained sleep data such as sleep duration and body movement, and none of them monitors respiratory rate and body position. Radio frequency (RF) based respiratory rate monitoring approaches have been proposed in [1, 2, 15, 17, 23]. In these approaches, respiratory rate is estimated by capturing the variation in the wireless signal's CSI (Channel State Information) or RSSI (Received Signal Strength Indicator), which is caused by the chest movement during respiration. Although RF based approaches can monitor respiratory rate non-invasively, they require wireless transmitting and receiving devices and their performance is sensitive to the environmental changes. Using wireless sensors to monitor body position during sleep has been studied in [14, 18]. Sensor data collected from accelerometers placed at different parts of the mattress [14] and sensor data collected from pressure sensors built in the specialized bedsheet [18] are respectively utilized to monitor body position during sleep. These systems are unobtrusive. However, they require sensor installation on the bed and their detection accuracies are affected by the user's height and weight.

Different from the aforementioned works, in this paper, we propose SleepMonitor, a smartwatch based system which leverages the built-in accelerometer to monitor a user's respiratory rate and body position during sleep. Our basic idea for respiratory rate monitoring is to capture the slight periodic wrist movement caused by the repeated inhalation and exhalation process during respiration, and the body position monitoring is based on the observation that in different sleeping positions, the user's wrist is likely to be put at different positions with respect to the body and with different poses.

As shown in Figure 1, the accelerometer data along three axes is continuously collected from the user's wrist-worn smartwatch, and segmented into sampling windows for processing. To estimate the user's respiratory rate, the window of raw data is first fed into a filter to remove noise. Since the signal variation in the filtered accelerometer data is mainly caused by the periodic patterns of inhalation and exhalation while the user breathes during sleep, we apply Fast Fourier Transform (FFT) to the filtered data and use the frequency with the largest magnitude to estimate the user's respiratory rate. The estimates from three axes are then fused together as the respiratory rate at the corresponding time. To monitor the user's body positions during the night, different features are extracted from the segmented windows. With the extracted features and the ground truth knowledge of the user's body positions at each time, different machine learning techniques can be applied to train classifiers for recognizing different body positions.

The main contributions of the paper are as follows:

- We propose a smartwatch based system to monitor respiratory rate and body position during sleep.
- We design a multi-axis fusion approach which can adaptively adjust the estimates from the three axes and then significantly improve the accuracy of respiratory rate monitoring.
- We have implemented our system on Android Wear based smartwatches and evaluated its performance in real experiments. The results show that our system can monitor respiratory rate and body position during sleep with high accuracy under various conditions.

The rest of this paper is organized as follows. Section 2 reviews related work. Section 3 and Section 4 illustrate the design details for respiratory rate monitoring and body position monitoring, respectively.

In Section 5, we evaluate the performance of our system under various conditions. Section 6 concludes the paper.

2 RELATED WORKS

Respiratory rate estimation is useful for health monitoring, and it has been studied in previous works. Traditionally, polysomnography(PSG) is used to estimate the patient's respiratory rate [29]. However, it involves wearable sensors with professional installation, and thus is limited to clinical usage. Some radio frequency (RF) based systems have been proposed in [1, 2, 15, 17, 23] to monitor respiratory rate in a non-invasive way. The intuition behind these schemes is to capture the variation in the wireless signal during its propagation, which is caused by the user's chest movement when he/she breathes, to estimate the respiratory rate. However, to implement these systems, wireless transmitting and receiving devices have to be installed. Besides, the performance of these RF-based systems may degrade due to the environment related changes or wireless signal interference. By leveraging the built-in sensors, researchers have also proposed smartphone based respiratory rate monitoring systems. In [24], the microphone is exploited to estimate the respiratory rate by capturing and analyzing the breathing sound. In [4], a smartphone is bound on the user's chest and the accelerometer and gyroscope is utilized to extract the chest movement pattern for estimating respiratory rate. However, both of them are inconvenient and not applicable for monitoring respiratory rate during sleep. In [24], to achieve accurate estimation, earphones or earbuds have to be worn during sleep, and in [4], the smartphone has to be bound on the chest for the whole night. Smartwatch's accelerometer data is exploited in *SeismoTracker* [10] and *BioWatch* [13] for respiratory rate estimation. However, *SeismoTracker* does not provide details on how to process the accelerometer data. *BioWatch* only uses data collected from one axis (the one with the largest magnitude in frequency domain), without fully taking advantage of information obtained from other axes. In addition, *BioWatch* also uses gyroscope, which may dramatically shorten the smartwatch's battery life since gyroscope consumes much more power than accelerometer. In [21], respiratory rate is monitored by using the smartphone as an active sonar system. However, to achieve accurate estimation, the smartphone has to be put close to the user. Different from all these works, we propose a smartwatch-based system to estimate respiratory rate during sleep and design a multi-axis fusion approach to get high accuracy estimation by exploiting information from different axes.

Body position is an important physiological parameter for sleep quality assessment, and many works have been proposed to monitor body position during sleep. A camera is utilized in [30] to capture the sequence of images during a user's sleep, and then image processing techniques are applied to the recorded images for detecting body positions. However, to implement this monitoring system, a camera has to be installed. Also, the use of camera may raise privacy concerns, which makes this approach impractical. By placing 3 WISP (Wireless Identification and Sensing Platform) tags at different positions of the mattress, Hoque *et al.* [14] monitored body positions through the variation of the accelerometer data collected from the WISP tags. In [18], a dense pressure sensitive bedsheet containing 64×128 pressure sensors was employed and the pressure images were collected and analyzed to recognize body positions. However, these systems are not practical since specialized hardware devices have to be installed on the bed.

3 RESPIRATORY RATE MONITORING

In SleepMonitor, the accelerometer data collected from the smartwatch worn on a user's wrist during his/her sleep is used to estimate the respiratory rate. This is motivated by the observation that a user's breaths can cause the periodic slight movement on his/her wrists while he/she is sleeping. As shown in Figure 2, respiration is the process of exchanging oxygen and carbon dioxide between the cells of the body and external environment, which consists of repeated cycles of inhalation and exhalation. During the

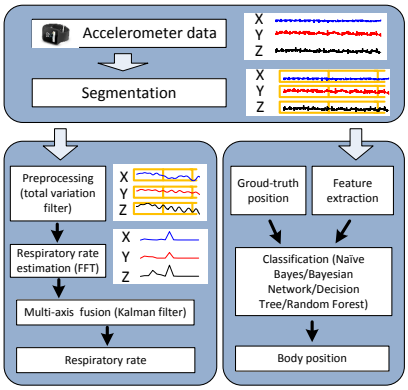

Fig. 1. System overview.

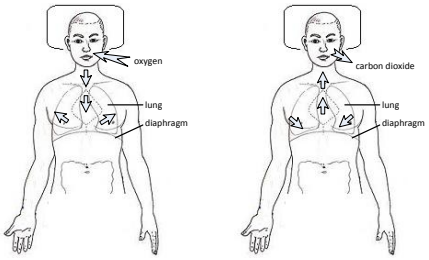

Fig. 2. The process of respiration.

inhalation cycle, air is taken into the lungs, and the diaphragm and intercostal muscle contract, causing the thoracic cavity to increase in size. During the exhalation cycle, air is moved out from the lungs, and the diaphragm and intercostal muscle relax, causing the thoracic cavity to decrease in size. The periodic increase and decrease in the volume of the thoracic cavity will lead to the periodic contractions and relaxations of the corresponding muscles, which eventually lead to the periodic movement of the chest, abdomen, arms and wrists. SleepMonitor leverages the accelerometer data collected on the user's wrist to detect his/her respiration cycles during sleep. The collected accelerometer data along three axes is first preprocessed to remove body movements and filter out noise. Then, frequency analysis is applied to estimate respiration rate from data along each axis. After that, a Kalman filter is designed to improve the estimation accuracy by fusing estimates from different axes together. In what follows, we discuss the design details for monitoring respiratory rate.

3.1 Preprocessing

In our system, the accelerometer data along three axes (denoted as a_x, a_y and a_z respectively) is continuously sampled with a sampling rate of 16 Hz, and then segmented into windows of 30 seconds for processing. SleepMonitor is designed to work when a user is in motionless sleep (i.e., quiet sleep without body movement), which occupies the majority of time of sleep. However, during sleep, the user may toss and turn occasionally and the accelerometer data collected at these moments should be discarded. To detect whether a window is sampled in motionless sleep or not, we calculate the total acceleration $a = \sqrt{a_x^2 + a_y^2 + a_z^2}$ for each accelerometer sample and compare it with a predefined threshold γ. Since the 3-axis accelerometer equipped on the smartwatch actually measures all the accelerations that affect the device, including the gravity, the total acceleration when the smartwatch is motionless will be equal to the gravity in magnitude. Thus, we can compare the total acceleration with the gravity in magnitude to detect if an accelerometer sample is collected in motionless sleep or not. In our system, γ is set to 10 m/s^2. If there are more than 5% of the accelerometer samples (i.e., 1.5 seconds) in a window with total acceleration larger than γ, this window is detected as containing body movement and discarded; otherwise, it is preserved for respiratory rate estimation.

During motionless sleep, the fluctuations caused by respiration in the accelerometer data collected from the wrist are weak and easy to be dominated by noises (shown in Figure 3(a)). In order to extract

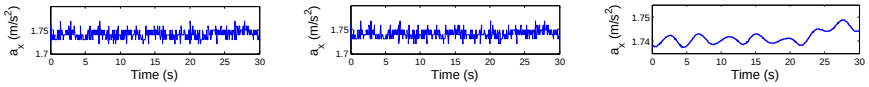

(a) Raw data (b) Low-pass filter (c) TV filter

Fig. 3. Denoising raw accelerometer data (along X axis) using low-pass filter and TV filter respectively.

respiratory signal from the weak and noisy accelerometer data, we first design a filter to remove the noise from the raw accelerometer data.

In SleepMonitor, the total variation filter (TV filter) [25] is used for denoising. By reducing the total variation of the raw signal, TV filter returns a filtered signal which is a close match to the original one but is much smoother. Different from low-pass filter, which only removes the high-frequency noise, TV filter removes both high and low-frequency noise and preserves the peaks and troughs which reflect the respiratory cycles. Figure 3 shows the raw and filtered accelerometer data when using a second-order Butterworth low-pass filter (cut-off frequency is set to 0.5 Hz) and a TV filter respectively. As can be seen, the data processed by low-pass filter is still noisy due to some low-frequency noise, while the data filtered by the TV filter shows clean respiration signal.

Let $\mathbf{s} \in \mathbb{R}^{n\times 1}$ denote a series of raw accelerometer data in a sampling window, let $\tilde{\mathbf{s}} \in \mathbb{R}^{n\times 1}$ denote the filtered data, and let s_i and $\tilde{s}_i$ denote the i^{th} sample in $\mathbf{s}$ and $\tilde{\mathbf{s}}$ respectively. Then, the TV filter is to find an appropriate $\tilde{\mathbf{s}}$ which minimizes the following objective function:

$$J(\tilde{\mathbf{s}}, \mathbf{s}) = E(\tilde{\mathbf{s}}, \mathbf{s}) + \lambda V(\tilde{\mathbf{s}}) \tag{1}$$

where

$$E(\tilde{\mathbf{s}}, \mathbf{s}) = \frac{1}{2}\sum_{i=1}^{n}(\tilde{s}_i - s_i)^2 \tag{2}$$

$$V(\tilde{\mathbf{s}}) = \sum_{i=2}^{n}|\tilde{s}_i - \tilde{s}_{i-1}| \tag{3}$$

and $\lambda > 0$ is the regularization parameter.

In the objective function $J(\tilde{\mathbf{s}}, \mathbf{s})$, $E(\tilde{\mathbf{s}}, \mathbf{s})$ measures the closeness between the filtered data $\tilde{\mathbf{s}}$ and the raw data $\mathbf{s}$, the total variation $V(\tilde{\mathbf{s}})$ describes the fluctuation in $\tilde{\mathbf{s}}$, and λ controls how smoothing $\tilde{\mathbf{s}}$ is. If λ is too small, the first term in $J(\tilde{\mathbf{s}}, \mathbf{s})$ is dominant and the filtered data will be close to the original data but with very little noise removed (shown in Figure 4(b)); if λ is too large, the second term in $J(\tilde{\mathbf{s}}, \mathbf{s})$ is dominant and the filtered data will be smooth but less like the original data (shown in Figure 4(c)). As shown in Figure 4(d), when λ is set appropriately (λ is set to 5 in our system), the noise can be filtered out and the periodic movement of the wrist which are hidden in the noisy raw data can be extracted.

The optimization problem shown in Equation 1 can be solved by Majorization-Minimization algorithm [26] with the following iterations:

$$\begin{cases} \tilde{\mathbf{s}}_k = \mathbf{s} - D^\top \mathbf{z}_{k-1} \\ \mathbf{z}_k = clip(\mathbf{z}_{k-1} + \frac{1}{4}D\tilde{\mathbf{s}}_k, \frac{\lambda}{2}) \end{cases} \tag{4}$$

where k is the iteration index, $\mathbf{z} \in \mathbb{R}^{(n-1)\times 1}$,

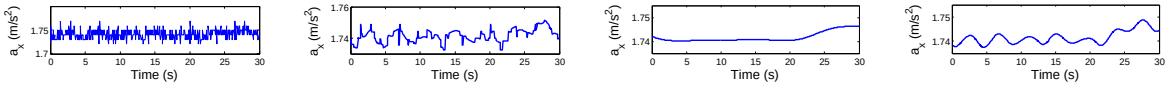

(a) Raw data (b) Filtered data ($\lambda = 1$) (c) Filtered data ($\lambda = 20$) (d) Filtered data ($\lambda = 5$)

Fig. 4. The raw and filtered accelerometer data (along X axis) collected from the wrist during sleep.

$$D_{(n-1)\times n} = \begin{bmatrix} -1 & 1 & & & \\ & -1 & 1 & & \\ & & \ddots & \ddots & \\ & & & -1 & 1 \end{bmatrix},$$

and the clip function $clip(\mathbf{x}, \beta)$ ($\mathbf{x} = [x_1, \ldots, x_i, \ldots, x_m]^\top$) is defined as

$$clip(\mathbf{x}, \beta) = \quad \mathbf{y} \quad | \quad \mathbf{y} = [y_1, \ldots, y_i, \ldots, y_m]^\top,$$
$$y_i = \begin{cases} x_i & \text{if } |x_i| \leq \beta \\ \beta \cdot sign(x_i) & \text{if } |x_i| > \beta \end{cases}.$$

As proven in [26], the optimization problem shown in Equation 1 is convex and the iteration will converge from any initialization. Thus, we can simply set $\mathbf{z}_0 = \mathbf{0}$ and solve the problem by using equation 4 iteratively to filter out noise in the raw accelerometer data.

3.2 Respiratory Rate Estimation

After the sampling windows are preprocessed, the noise is removed and the acceleration fluctuation due to respiration is preserved. Since the fluctuation in the preprocessed data is mainly caused by the repeated inhalation and exhalation process while a user breathes during sleep, the user's respiratory rate can be estimated as the frequency of the fluctuation. To calculate such frequency, we apply FFT to the filtered accelerometer data. Since FFT decomposes the time-domain signal into frequency components which make it up, for a quasi-periodic time-domain signal, there will be a strong frequency component in the corresponding FFT, which indicates the frequency of the time-domain signal. In SleepMonitor, the respiratory pattern is hidden in the accelerometer data, which is quasi-periodic after being preprocessed (red line in Figure 5(a)), and thus by applying FFT to the filtered accelerometer data, the respiratory rate (in Hertz) can be estimated as the frequency of the component with the largest magnitude (shown in Figure 5(b)).

As shown in Figure 5(a), although respiration during sleep only causes a slight movement of the wrist, due to the effect of gravity, the absolute values of the accelerations along X, Y and Z axis are much larger than 0. For each axis, compared with the slight acceleration fluctuation caused by respiration, the bias caused by gravity is much larger, resulting in a very strong DC component in the corresponding FFT, which can dominate all the other frequency components. In order to detect the respiration related frequency component, we simply remove the DC component in the FFT and set its magnitude to 0. In addition, since the human respiratory rate is less than 30 bpm (breaths per minute) [16], in FFT, we only consider the components with frequencies less than 0.5 Hz (i.e., 30 bpm).

3.3 Multi-axis Fusion

As shown in Figure 5, in total, 3 respiratory rate estimates are obtained from X, Y and Z axis. To obtain more accurate respiratory rate, we design a multi-axis fusion approach to fuse these estimates together.

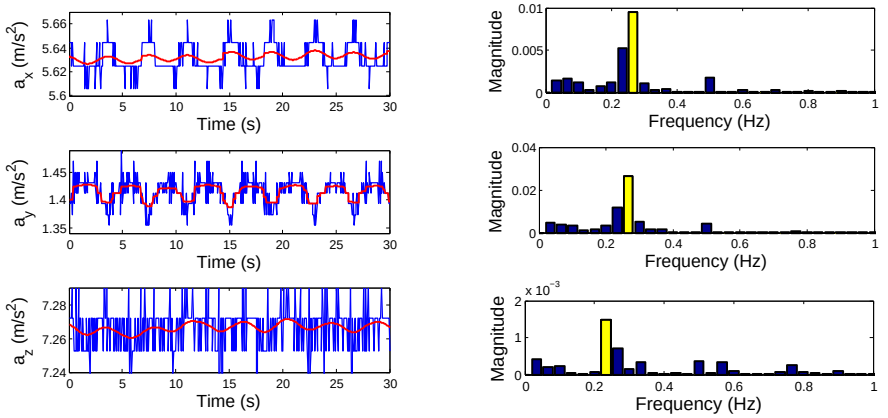

(a) The raw and filtered accelerometer data along three axes

(b) FFT of the filtered data along three axes

Fig. 5. The raw (blue line) and filtered (red line) accelerometer data along each axis and the respiratory rate estimation by applying FFT (with DC component removed) to the filtered data along each axis. The respiratory rate is estimated as the frequency with the largest magnitude (marked as yellow). The ground truth respiratory rate is 0.267 Hz (16 bpm). According to the figure, respiratory rates estimated with accelerations along X, Y and Z are 0.267 Hz, 0.267 Hz and 0.233 Hz respectively. A multi-axis fusion approach will be used to further improve the estimation accuracy (see Section 3.3).

To fuse estimates obtained from different axes together, a straightforward idea is to average these estimates. However, this may result in large error in some cases due to the possible inaccurate estimates along some axes. During a user's sleep, the smartwatch worn on his/her wrist may be at different positions (e.g., at the side of the body, on the belly or close to the pillow) with different poses (e.g., smartwatch screen facing up, down or other directions). Depending on the position and pose of the smartwatch, the intensity of the acceleration fluctuation caused by respiration and the noisiness of the accelerometer data along each axis may be different, which may result in different estimation accuracies along different axes. For example, in Figure 5, compared with Z axis, the acceleration fluctuations along X and Y axis are much stronger and their data is less noisy, and thus the respiratory rates obtained from X and Y-axis data are more accurate than that from Z-axis data. Therefore, simply averaging results from all three axes may cause large inaccuracy if the results from some axes are far away from the ground truth value.

In our system, to fuse respiratory rates estimated from different axes, instead of merely considering the three estimates obtained at a particular moment, we also exploit the historical information by implementing a Kalman filter (KF), where the respiratory rates calculated previously are used for prediction and the current estimates from 3-axis accelerometer data are used as the measurement in KF model for update. More specifically, in our Kalman filter, we predict the respiratory rate at time step t (i.e., the t^{th} sampling window in our system) as that of time step $t-1$ since the respiratory rates of the two continuous sampling windows are not likely to vary dramatically during sleep. Then we use the weighted arithmetic mean of the three respiratory rates obtained from X, Y and Z axis at time step t as the measurement and use the corresponding standard deviation as the variance of the measurement

noise in KF model (see Equation 10 and 11). Let $rr_{t|t}$ and $rr_{t|t-1}$ respectively denote the a posteriori and the a priori respiratory rate at time step t, the Kalman filter recursively updates the respiratory rate as follows:

$$rr_{t|t-1} = rr_{t-1|t-1} \tag{5}$$

$$p_{t|t-1} = p_{t-1|t-1} \tag{6}$$

$$k_t = \frac{p_{t|t-1}}{p_{t|t-1} + r_t} \tag{7}$$

$$rr_{t|t} = rr_{t|t-1} + k_t(rr_t^m - rr_{t|t-1}) \tag{8}$$

$$p_{t|t} = (1 - k_t)p_{t|t-1} \tag{9}$$

where $p_{t|t}$ and $p_{t|t-1}$ respectively represent the a posteriori and the a priori estimation error variances at time step t, k_t is the Kalman gain at t, the measurement rr_t^m is

$$rr_t^m = \sum_{ax \in \{x,y,z\}} w_t^{ax} \cdot rr_t^{ax}, \tag{10}$$

and the variance of the measurement noise r_t is

$$r_t = \sqrt{\sum_{ax \in \{x,y,z\}} [w_t^{ax} \cdot (rr_t^{ax} - rr_t^m)]^2}. \tag{11}$$

In Equation 10 and 11, w_t^{ax} is a normalization factor to describe the weight of rr_t^{ax} when calculating rr_t^m. Let var_t^{ax} denote the variance of the filtered accelerometer data along axis ax in the window at time step t, and the normalized weight w_t^{ax} is calculated as:

$$w_t^{ax} = \frac{var_t^{ax}}{var_t^x + var_t^y + var_t^z}$$

When the variance along an axis is greater, which means the wrist movement caused by respiration along that axis is stronger, the respiratory rate estimated from that axis should be more reliable. Thus, in our system, we assign weights to the estimates from different axes based on their corresponding variances.

As can be seen from Equation 8, the Kalman gain k_t is a weighting factor to decide the weights of the respiratory rate calculated at $t-1$ and the measurement at t in calculating the final respiratory rate at t. When k_t is large, the filter places more weight on the current measurement, and when k_t is small, the filter places more weight on the previous respiratory rate. According to Equation 7, k_t is related to the variance of the measurement noise r_t, which describes the reliability of the measurement at t. In our system, when the estimates from the three axes are more consistent, r_t is smaller and k_t is larger, resulting in more weight placed on the current measurement in Equation 8. Otherwise, the measurement is less reliable, k_t is smaller and $rr_{t|t}$ is closer to the previous respiratory rate. At the extremes, if the estimates from all three axes are the same, $r_t = 0$, $k_t = 0$ and $rr_{t|t} = rr_t^m$; if the estimates from all three axes are totally different, $r_t \to \infty$, $k_t \to 1$ and $rr_{t|t} \to rr_{t-1|t-1}$.

Kalman filter is recursive and works in real time. As shown in Figure 6, the respiratory rates obtained by using Kalman filter are more accurate than those obtained by simply averaging estimates from different axes. Even when the respiratory rates in two continuous sampling windows change suddenly, the respiratory rate calculated in Kalman filter can be corrected quickly by the new measurement values.

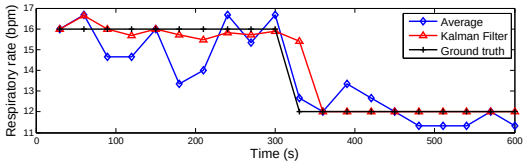

Fig. 6. The respiratory rates obtained by averaging estimates from all three axes and by using Kalman filter. A user is asked to change his respiratory rate while he is lying on the bed. The ground truth is obtained by asking the user to breathe to a metronome.

4 BODY POSITION MONITORING

Body position is an important physiological parameter in sleep study. In [7], body position data during sleep was collected to assess a user's sleep quality, and in [22] and [8], researchers monitored body positions during sleep to study how different body positions affected sleep-related diseases, such as breathing disorders and back pains. Based on the 3-axis accelerometer data collected from the smartwatch worn on a user's wrist, SleepMonitor can detect the body position at each time step as one of the four basic positions, including lying on the back, the chest, the left side and the right side (denoted as *supine*, *prone*, *left* and *right* respectively).

When a user sleeps in different positions, his/her wrist (either left or right wrist) is normally at different positions with different poses, and the wrist position and pose are related to the body position. Take the left wrist as an example: when sleeping in *supine*, the user is likely to put his/her wrist at the left side of the body or on the belly, with the palm facing down or sideways; when sleeping on *left*, the user is likely to put his/her wrist at the left side of the body or close to the pillow, with the palm facing up. Thus, body position can be detected by the position and pose of the wrist. In SleepMonitor, we first filter out sampling windows which contain body movements (illustrated in Section 3.1). Then, in the preserved sampling windows, we extract features which can reflect the wrist position and pose, and use various classification algorithms to train classifiers for body position detection.

Although movement patterns of the left and right wrist are different during sleep, the technique used for detecting body position is the same, regardless of the wrist. Only experimental parameters need to be adjusted when the smartwatch is worn on different wrists. To simplify our representation, in this section, the system is illustrated assuming that the smartwatch is worn on the left wrist.

4.1 Ground Truth

In order to train a classifier which can recognize body positions based on the extracted features, the ground-truth body position at each time step must be collected. To obtain the ground truth, two approaches are used in most of the existing works [5, 9, 14, 18, 30]. In [5, 9, 14], low-light cameras are utilized to monitor the users' sleep and their body positions are labeled based on the video data. In [18, 30], users are asked to sleep in particular body positions under supervision in lab environment and the ground truth is lab controlled. However, both have disadvantages. For the video based approach, it is difficult to recruit volunteers due to privacy concerns. Besides, to label a user's ground-truth body positions, the video recorded during the whole night should be played and watched, which costs much time and labor. For the lab-controlled approach, the body positions are not obtained while the users are sleeping. Since a user may sleep differently under supervision in lab environment from normal days, the experimental data collected using such approach may not accurately reflect the real body positions during sleep. To overcome these problems, we designed our own approach.

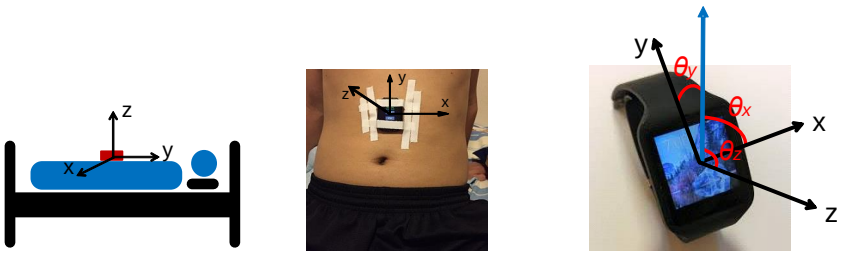

Fig. 7. The experimental setup to get the ground-truth body position.

Table 1. Angles between each axis and the upward direction in different body positions.

Body position	$(\theta_x, \theta_y, \theta_z)$	
	Ideal case	Our experiment
Supine	$(\frac{\pi}{2}, \frac{\pi}{2}, 0)$	$([\frac{\pi}{3}, \frac{2\pi}{3}], [\frac{\pi}{3}, \frac{2\pi}{3}], [0, \frac{\pi}{4}])$
Prone	$(\frac{\pi}{2}, \frac{\pi}{2}, \pi)$	$([\frac{\pi}{3}, \frac{2\pi}{3}], [\frac{\pi}{3}, \frac{2\pi}{3}], [\frac{3\pi}{4}, \pi])$
Left	$(\pi, \frac{\pi}{2}, \frac{\pi}{2})$	$([\frac{2\pi}{3}, \pi], [\frac{\pi}{3}, \frac{2\pi}{3}], [\frac{\pi}{4}, \frac{3\pi}{4}])$
Right	$(0, \frac{\pi}{2}, \frac{\pi}{2})$	$([0, \frac{\pi}{3}], [\frac{\pi}{3}, \frac{2\pi}{3}], [\frac{\pi}{4}, \frac{3\pi}{4}])$

As shown in Figure 7(a) and Figure 7(b), a smartwatch without band is attached to the user's belly while the user sleeps, and its pose varies when the user changes his/her body position. Based on the pose of the smartwatch, we can detect the user's body position correctly. Figure 7(c) shows the smartwatch's local coordinate system: when looking at the smartwatch screen in a normal way (e.g., checking the time), X axis points towards the right side of the screen, Y axis points towards the top of the screen, and Z axis is perpendicular to the screen and points towards the sky. Given such local coordinate system, the smartwatch's pose and the user's body position can be detected through the angles between each axis and the upward direction (denoted as θ_x, θ_y, θ_z respectively, shown in Figure 7(c)). For example, as shown in Figure 7(a) and Figure 7(b), the smartwatch is placed on the user's belly with its X axis pointing towards the left side of the body and its Y axis pointing towards the user's head. Ideally, if the user's torso is modeled as a rigid body, when the user is lying on the back, X and Y axis are perpendicular to the upward direction, Z axis points upwards, and $\theta_x = \frac{\pi}{2}$, $\theta_y = \frac{\pi}{2}$ and $\theta_z = 0$; when the user is lying on the left side, Y and Z axis are perpendicular to the upward direction, X axis points downwards, and $\theta_x = \pi$, $\theta_y = \frac{\pi}{2}$ and $\theta_z = \frac{\pi}{2}$. In practice, the angles measured for a particular body position may vary slightly around those in ideal case. For different body positions, Table 1 shows the corresponding angles in ideal case and their ranges measured in our experiment.

In order to use Table 1 to infer the user's body position, the angles between the smartwatch's three axes and the upward direction must be obtained. When a smartwatch is static, the total acceleration, which is calculated from the accelerations measured along X, Y and Z axis, is equal to the gravity in magnitude but points to the opposite. Since the gravity is always downwards, the total acceleration measured from a static smartwatch always points upwards. In this case, the angles between each axis and the upward direction can be calculated as:

$$\underset{ax\in\{x,y,z\}}{\theta_{ax}} = \arccos\left(\frac{a_{ax}}{\sqrt{a_x^2 + a_y^2 + a_z^2}}\right) \tag{12}$$

During sleep, except for some occasional body movements, the user's body remains static for most of the time and the accelerometer of the attached smartwatch is only affected by the gravity and respiration. Since the acceleration caused by the respiration is much smaller than that caused by gravity, the effect of the respiration can be ignored and we can use Equation 12 to calculate the angles and further detect the user's body positions. In our scheme, similar to that described in Section 3, the accelerometer data collected from the belly is segmented as the same size windows as those collected from the wrist, and the windows containing body movements are discarded. Then, for the preserved windows, we use Equation 12 to calculate the angles and use Table 1 to infer the corresponding body positions.

To verify our scheme, we used a camera to record a user's sleep for one night, and compared the ground truth obtained from the video data with that from our scheme (denoted as *Video* and *Acc* respectively). Figure 8 shows 30-minute data of the accelerations, the calculated angles and the body positions labeled using *Video* and *Acc*. As can be seen from Figure 8(d), all body positions labeled by *Acc* are consistent with those labeled by *Video*. Some short periods around the body movements are not labeled by *Acc* because it labels body positions based on 30-second-long windows. Except for these short periods, the ground truth obtained by *Video* and *Acc* are consistent for most of the time. Thus, we use the smartwatch attached to the user's belly to obtain ground truth in our experiment.

4.2 Feature Extraction and Classification

Extracting the proper features is important for classifying body positions. Since the position and pose of a user's wrist are related to his/her body position during sleep, given the accelerometer data collected in SleepMonitor, we are interested in extracting the features which can distinguish the different positions and poses of the smartwatch when the user is in different body positions.

We extract the standard deviation of the accelerometer data in a window as a feature to distinguish the wrist positions. When a user puts his/her wrist at different positions during sleep, the amplitudes of the wrist movement caused by respiration may be different, and thus the intensities of the fluctuations in the collected accelerometer data may be different. For example, as shown in Figure 9, the accelerometer data collected when the wrist is put on the belly fluctuates much stronger than that collected when the wrist is at the side of the body. Since standard deviation can be used to measure the intensity of the fluctuation in a signal, we use it as a feature to distinguish the wrist positions.

As illustrated in Section 4.1, the wrist pose can be described by the angle between each axis and the upward direction, which can be calculated by Equation 12. However, the angle is calculated based on each acceleration sample, which is not a window based feature. To extract a window based feature, we need to average all the calculated angles in a window. Since the sampling window processed here does not contain body movement and the wrist pose does not change very much, we can first average all the accelerations along each axis and then use Equation 12 to reduce the computation overhead. Further, since the wrist is almost static when the window is sampled, the magnitude of the total variation is equal to the gravity, which is a constant. Thus, we can ignore the denominator in Equation 12 and use the mean value of the accelerations along each axis as a feature to distinguish the wrist poses.

Besides mean and standard deviation, we also extract the 20th percentile, median, 80th percentile of the accelerometer data along each axis and the covariance and correlation between accelerometer data along every two different axes as features in our system.

After the features are extracted from each window, different machine learning techniques can be used to construct the body position classifier. In our system, we implement 4 commonly used classification

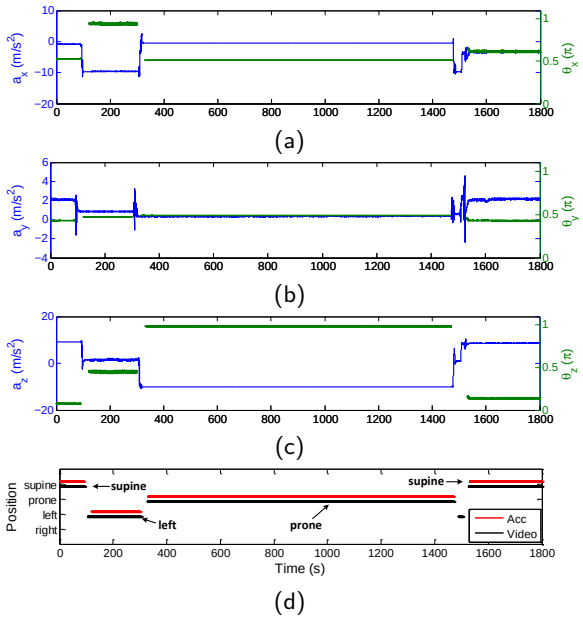

Fig. 8. The ground-truth body positions obtained by using *Video* and *Acc*. In *Acc*, angles between each axis and the upward direction (green line) calculated from the accelerometer data (blue line) are used to label body positions based on Table 1.

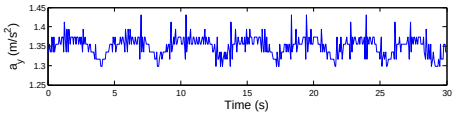

(a) At the side of the body ($std = 0.0257$).

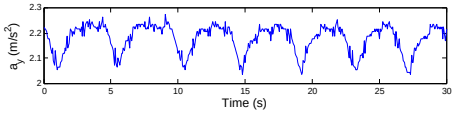

(b) On the belly ($std = 0.0558$).

Fig. 9. The standard deviation (*std*) of the Y-axis accelerometer data in a window when the smartwatch is at different positions.

algorithms: Naive Bayes (NB), Bayesian Network (BN), Decision Tree (DT) and Random Forest (RF), and compare their performances in Section 5.

5 PERFORMANCE EVALUATIONS

In this section, we evaluate the performance of SleepMonitor based on the data collected in real experiments.

5.1 Experimental Setup

To evaluate the performance of our system, we implemented SleepMonitor on two Android Wear based smartwatches: Sony Smartwatch 3 and Huawei Watch.

As illustrated in Section 4.1, a smartwatch without band can be attached to the user's belly to collect ground-truth data for body position. It can also be used to collect ground-truth data for respiratory rate. As shown in Figure 10, since respiration causes strong movement of the belly, in the corresponding accelerometer data, the fluctuation indicating respiratory cycles is also strong. Thus, the filtering and frequency analysis technique described in Section 3 can be applied to the accelerometer data collected from the belly to estimate respiratory rate accurately, and the results can be used as the ground truth. Different from the ground truth collection approaches in [15, 23], which are lab-controlled by asking users to artificially breathe to a metronome, our approach is able to collect the ground-truth respiratory rate when users breathe naturally during sleep with only a smartwatch attached to the belly.

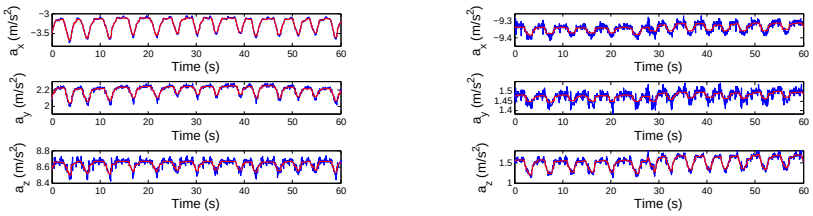

(a) *Supine* (b) *Left*

Fig. 10. The raw (blue line) and filtered (red line) accelerometer data collected from the belly while the user is sleeping in different positions.

Table 2. Overview of the experiment setup.

User	Nights	Gender	Age	Smartwatch (belly)	Smartwatch (wrist)	Respiration pattern
1	8	M	20-30	Sony	Huawei/Sony	Normal
2	4	F	30-40	Sony	Huawei/Sony	Normal
3	5	M	20-30	Sony	Huawei/Sony	Normal
4	4	F	20-30	Sony	Huawei/Sony	Normal
5	8	M	30-40	Sony	Huawei/Sony	Normal
6	6	M	20-30	Sony	Huawei/Sony	Normal
7	5	M	30-40	Sony	Huawei/Sony	Normal
8	3	F	30-40	Sony	Huawei/Sony	Normal
9	1	M	20-30	Sony	Huawei	Normal
10	2	M	20-30	Sony	Huawei	Normal
11	6	F	20-30	Sony	Huawei	Artificial
12	2	M	20-30	Sony	Huawei	Artificial
13	4	M	30-40	Sony	Huawei	Artificial
14	4	M	20-30	Sony	Huawei	Artificial
15	4	F	40-50	Sony	Huawei	Abnormal
16	4	M	50-60	Sony	Huawei	Abnormal

As shown in Table 2, we recruited 16 people and collected 70 nights of sleep data in our experiment. Among these users, user 15 and user 16 are sleep-impaired, and one of them has sleep apnea (denoted as *Abnormal* in the last column of Table 2). Users 11 to 14 are artificially irregular breathing users, who were asked to intentionally change respiratory rate for half an hour before falling asleep (denoted as *Artificial* in Table 2).

During the experiment, each user was given two smartwatches: one worn on the left wrist and the other one attached to the belly. These two smartwatches were synchronized before given to the user. After each night, we collected smartwatches from the user and read accelerometer data from both smartwatch's SD cards. The collected accelerometer data was processed separately to obtain experimental results and ground truth. The experiment has been approved by our IRB (Institutional Review Board).

5.2 Respiratory Rate

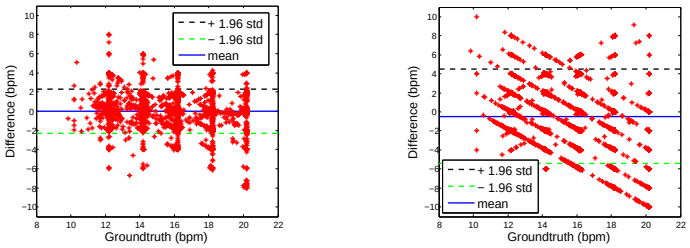

(a) SleepMonitor (b) BioWatch

Fig. 11. Bland-Altman plot for respiratory rate estimation using SleepMonitor and BioWatch respectively.

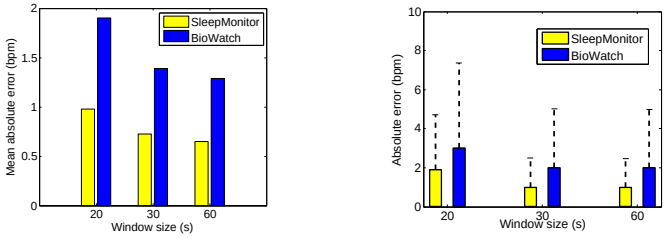

(a) (b)

Fig. 12. Effect of window size on respiratory rate estimation.

5.2.1 Comparison with BioWatch. We use Bland-Altman plot [3] to evaluate the estimation accuracy of SleepMonitor and compare it with that of BioWatch [13]. As shown in Figure 11, the Bland-Altman plot is modified a little bit by using the groundtruth as the X-axis and the estimation difference between the evaluated system and groundtruth as the Y-axis. In our comparison, we apply FFT to the accelerometer data along three axes and select the estimate from the one with the maximum magnitude in frequency domain as BioWatch's respiratory rate estimation.

As shown in Figure 11(a), when using SleepMonitor, most of the estimations are within the limits of agreement (i.e., interval between +1.96 *std* and -1.96 *std*). Comparing Figure 11(a) with Figure 11(b), the mean and standard deviation of the estimation difference between SleepMonitor and groundtruth ($mean = 0.0274$ and $std = 1.0235$) are much smaller than those between BioWatch and groundtruth ($mean = -0.4714$ and $std = 2.4841$), indicating that SleepMonitor estimates respiratory rate more accurately than BioWatch. This is because SleepMonitor leverages estimates from three axes and fuses them together to reduce errors, while BioWatch only uses the accelerometer data from one axis, without taking advantage of information from other axes.

5.2.2 Effect of window size. Figure 12 shows the mean absolute estimation error and the range of the absolute error when using SleepMonitor and BioWatch respectively under different window sizes. As discussed in Section 5.2.1, since the weighted arithmetic mean and Kalman filter are implemented in SleepMonitor to fuse estimates from different axes together, SleepMonitor performs much better than BioWatch under various window sizes. For both SleepMonitor and BioWatch, the mean absolute error decreases and the absolute error varies within a shorter range when the window size is increased.

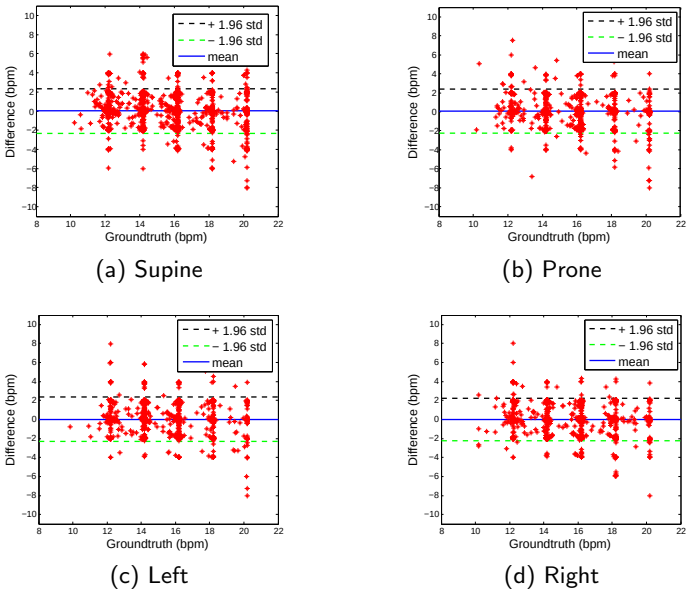

Fig. 13. Bland-Altman plot for respiratory rate estimation when using SleepMonitor in different body positions.

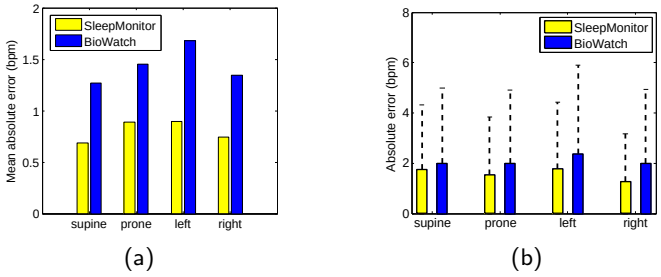

Fig. 14. Effect of body position on respiratory rate estimation.

This is because in both *SleepMonitor* and *BioWatch*, FFT is applied to the accelerometer data for respiratory rate estimation, and the frequency resolution of FFT is decided by the sampling rate and the number of samples used in FFT. At a fixed sampling rate (i.e., 16 Hz in our experiment), the frequency resolution is increased with more samples used in FFT (i.e., larger window), and thus the respiratory rate estimation is more accurate with a larger window. However, increasing the window size will need more computing resource and consume more power, which is limited on smartwatch. As show in Figure 12, when 30-second-long window is used, the estimation error falls to an acceptable range with a mean absolute error of 0.72 bpm. Thus, in our system, the window size is set to 30 seconds.

5.2.3 Effect of body position. Figure 13 shows the Bland-Altman plot for respiratory rate estimation when using SleepMonitor in different body positions. As shown in the figure, no matter in which position

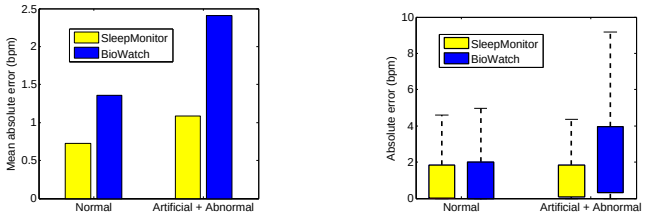

(a) (b)

Fig. 15. Effect of respiration pattern on respiratory rate estimation.

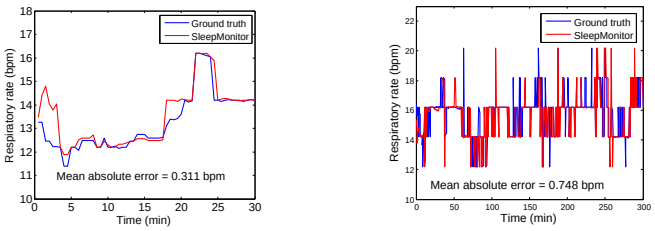

(a) User 11 (b) User 16

Fig. 16. Respiratory rate estimation for user 11 and 16.

the user sleeps, most of the estimation errors are within around ±2 bpm, indicating that SleepMonitor can estimate respiratory rate accurately regardless of the sleep position. Figure 14 shows the mean absolute error and the range of the absolute error of respiratory rate estimation when using SleepMonitor and BioWatch in different body positions. As can be seen, for both SleepMonitor and BioWatch, the estimation is more accurate in *supine* than in other body positions. This is because users are more likely to put the wrist on the belly in *supine* than in other body positions, which leads to stronger wrist movement and thus more accurate estimation. When users are in *right*, the left arm is likely to be put on the torso, leading to stronger movement on the left wrist than on the right. In our experiment, users wear smartwatches on their left wrists, and thus the fluctuation in the accelerometer data collected in *right* is stronger than that in the data collected in *left*, resulting in less estimation error in *right* than in *left*.

5.2.4 Effect of respiration pattern. To evaluate our system under different respiration patterns, we use the first half hour's sleep data from users 11 to 14 (artificially irregular breathing users) and the whole night's sleep data from user 15 and user 16 (sleep-impaired patients) as a dataset, and compare the respiratory rate estimation using this dataset (denoted as *Artificial + Abnormal*) with that using the sleep data from users 1 to 10 (denoted as *Normal*). As shown in Figure 15, although the estimation accuracy in *Artificial + Abnormal* is not as good as that in *Normal*, SleepMonitor can still estimate respiratory rate accurately when the respiration pattern varies, with a mean absolute error of 1.08 bpm. To study how SleepMonitor tracks rapid, unpredictable changes in breathing patterns while lying down, we use one night of user 11's sleep data and show the respiratory rate estimation at different time in Figure 16(a). In the figure, we only use the 30-minute sleep data when the user changes his respiratory rate. As can be seen, even if the user's respiratory rate varies, SleepMonitor can capture the variations

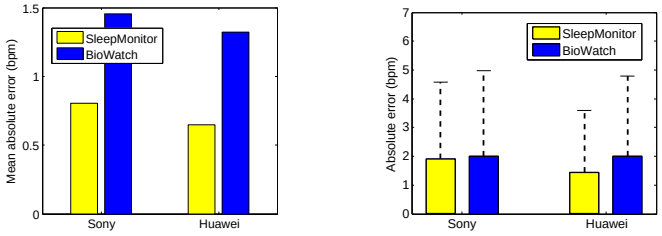

Fig. 17. Effect of device on respiratory rate estimation.

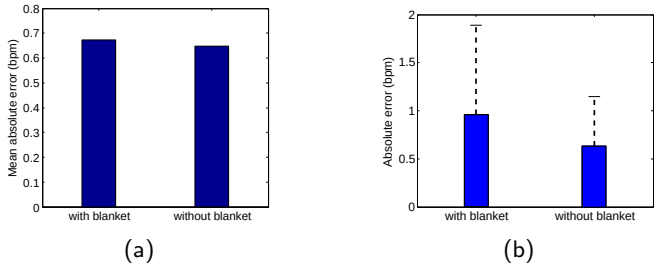

Fig. 18. Effect of blanket on respiratory rate estimation.

and return estimation with mean absolute error of 0.311 bpm. To understand SleepMonitor's performance for people with sleep diseases, we use one night of user 16's sleep data and show the respiratory rate estimation in Figure 16(b). As shown in the figure, for most of the time, SleepMonitor can estimate the respiratory rate with small errors. The mean absolute error based on the 300-minute-long sleep data is less than 1 bpm (i.e., 0.748 bpm).

5.2.5 Effect of device. To evaluate our system across devices, two smartwatches, Sony Smartwatch 3 and Huawei Watch (denoted as *Sony* and *Huawei* respectively), are tested and their performance on respiratory rate estimation is shown in Figure 17. As can be seen, although *Huawei* performs a little bit better than *Sony*, both of them can estimate respiratory rate accurately when using SleepMonitor, verifying that our system can work on different types of smartwatches for respiratory rate monitoring.

5.2.6 Effect of blanket. In practice, people use blankets during sleep. To evaluate the effect of using blanket on the accuracy of respiratory rate estimation, we asked users to sleep with blanket and without blanket for 20 minutes respectively and estimated the respiratory rate based on these two datasets. In each experiment, users slept in the supine position and put wrist on the belly. As shown in Figure 18, in both cases, SleepMonitor can estimate respiratory rate with high accuracy (less than 0.7 bpm). Although the accuracy with blanket is a little bit worse than that without blanket, there is no major difference, demonstrating that our system can work well in practice when people use blankets.

5.2.7 Effect of λ. The raw accelerometer data collected from wrist during sleep is very noisy and we need to filter out the noise before any further techniques can be applied. To implement the filter, as discussed in Section 3.1, an appropriate value for the regularization parameter λ is important. In

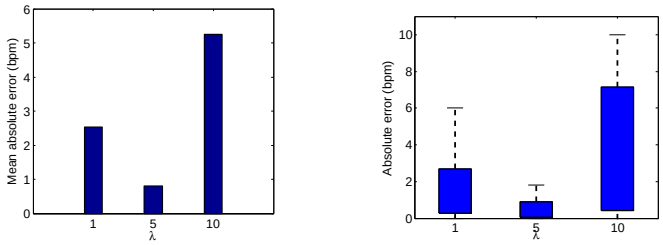

(a) (b)

Fig. 19. Effect of λ on respiratory rate estimation.

Table 3. The overall detection results using different classification algorithms.

Position	NB		BN		DT		RF	
	TPR	*FPR*	*TPR*	*FPR*	*TPR*	*FPR*	*TPR*	*FPR*
supine	0.70	0.24	0.84	0.18	0.93	0.08	0.95	0.05
prone	0.56	0.26	0.77	0.05	0.84	0.06	0.90	0.03
left	0.61	0.14	0.81	0.12	0.87	0.02	0.92	0.02
right	0.54	0.07	0.80	0.03	0.88	0.02	0.90	0.02

order to evaluate the effect of λ on respiratory rate estimation, we collect sleep data from users and set λ to different values. As shown in Figure 19, the mean absolute error when λ is set to 10 (i.e., too large) or 1 (i.e., too small) is larger than that when λ is set to 5, which is consistent with the analysis in Section 3.1. Compared with $\lambda = 10$, the respiratory rate estimated when $\lambda = 1$ is more accurate. This is because when λ is small, even some noise cannot be filtered, the respiration signal in the raw data is still preserved, which can be used to estimate respiratory rate, but when λ is large, the respiration signal may be removed as noise.

5.3 Body Position

5.3.1 Overall performance. We use *True Positive Rate* (TPR) and *False Positive Rate* (FPR) to evaluate the performance of our system on body position monitoring. After extracting features from sampling windows, we use Weka [11] to implement four commonly used classification algorithms, Naive Bayes, Bayesian Network, Decision Tree and Random Forest (denoted as *NB*, *BN*, *DT* and *RF* respectively), and the overall performance of each algorithm based on *leave-one-user-out* strategy is shown in Table 3. For a particular user, we use data collected from the other users to train classification models and use data from this user to test. The process is repeated to cross validate the classification algorithms across users.

As can been seen, *RF* outperforms all the other schemes, and 95% of *supine*, 90% of *prone*, 92% of *left* and 90% of *right* can be correctly classified using our extracted features. *DT* can also recognize different kinds of body positions with TPR larger than 84%. For all 4 classification algorithms, *supine* is classified more accurately than other body positions, but it also has lager FPR than others, indicating that many other body positions are misclassified as *supine*.

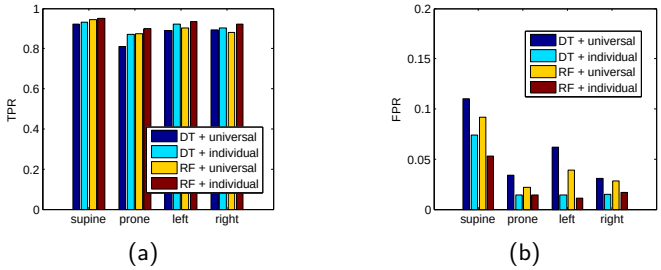

Fig. 20. TPR and FPR when using different training datasets.

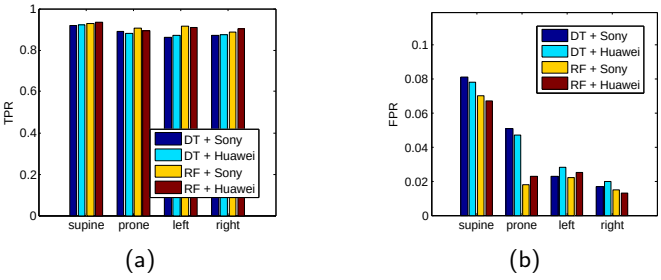

Fig. 21. TPR and FPR when using different devices.

5.3.2 Effect of training dataset. For a specific user, the classifiers can be trained by using different datasets: dataset collected from other users and dataset collected only from this user. To evaluate the effect of dataset on the performance of body position classification, we focus on a particular user and train model by using dataset from all the other users and dataset from this user (denoted as *universal* and *individual* respectively). For *individual*, we use the *leave-one-night-out* strategy to cross validate across nights. Each time, data collected from one night is used for testing and data collected from the remaining nights is used for training. This process is repeated so that each night's data is used exactly once as the testing data. Since *DT* and *RF* perform much better than *NB* and *BN*, we use *DT* and *RF* as the classification algorithms and the results are shown in Figure 20. As can be seen, for both *DT* and *RF*, *individual* achieves better performance than *universal*. More than 87% of different kinds of body positions can be correctly classified in *individual*, and the FPR is less than 7% in *individual*. This is because different users have different sleeping patterns, and compared with *universal*, the features extracted from a specific user in *individual* are more distinctive and more reliable to describe the relationship between the body position and the corresponding wrist position and pose.

5.3.3 Effect of device. Figure 21 shows the TPR and FPR of different body positions when using different classification algorithms and different types of smartwatches. Similar to Section 5.2.5, there is no major difference in the performance of *Sony* and *Huawei*. Both of them can correctly identify more than 87% of body positions when using *DT* and more than 88% when using *RF*. No matter which classification

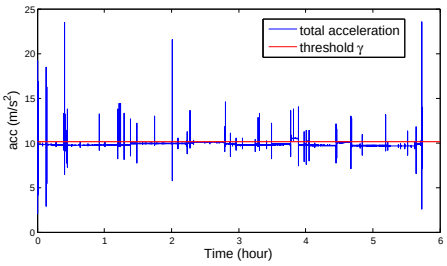

Fig. 22. The total acceleration collected from a user's wrist during his entire sleep.

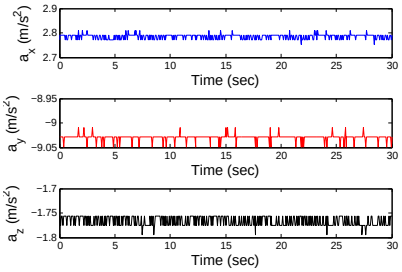

Fig. 23. The accelerometer data collected from a user's wrist when he wears the smartwatch loosely.

algorithm (i.e., *DT* or *RF*) is used, both *Sony* and *Huawei* can achieve FPR less than 9%. Therefore, SleepMonitor can work on different types of smartwatches for body position monitoring.

5.4 Discussions

SleepMonitor is designed to work in an unobtrusive way. The user only needs to start it by a simple click before sleep and stop it after getting up. Although the user may not go to motionless sleep immediately after starting SleepMonitor and may occasionally turn around during sleep, these events can be filtered out by using a predefined threshold γ (see Section 3.1) and they will not affect the respiratory rate estimation during motionless sleep. Figure 22 shows the total acceleration collected from a user's wrist during his entire sleep. As can be seen, for most of the time, the user is in motionless sleep that can be detected by threshold γ, demonstrating that our system is suitable for monitoring the respiratory rate during sleep.

SleepMonitor estimates a user's respiratory rate during sleep by capturing the slight periodic wrist movement caused by the repeated inhalation and exhalation process during respiration. However, there are cases when a user wears the smartwatch loosely, which makes it difficult to capture the wrist movement through smartwatch's accelerometer data. As shown in Figure 23, when a user wears the smartwatch loosely, we cannot detect respiration cycles since the accelerometer readings along three axes are too weak. We can use mean and variance-based thresholds to detect such situation and display a warning signal (or vibration) to let the user tighten the smartwatch.

6 CONCLUSIONS

In this paper, we designed SleepMonitor, a smartwatch based system which can monitor a user's respiratory rate and body position during his/her sleep by leveraging the accelerometer data collected from the wrist. Filtering and frequency analysis techniques are applied to estimate the respiratory rate as the frequency of the wrist movement along each axis caused by respiration during sleep. A multi-axis fusion approach is further designed to improve the estimation accuracy. Features which describe the wrist position and pose are extracted from the accelerometer data, and machine learning techniques are applied to detect the body position. We have implemented our system on Android Wear based smartwatches and evaluated its performance in real experiments. The results show that our system can monitor respiratory rate and body position during sleep with high accuracy under various conditions. Although our system is illustrated based on smartwatch, it can also work on other wrist-worn devices (e.g. Fitbit or Jawbone) since the only sensor used in SleepMonitor is 3-axis accelerometer, which is equipped on almost all wearable devices.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers whose insightful comments helped improve the presentation of this paper significantly. This work was supported in part by the National Science Foundation (NSF) under grants CNS-1421578 and CNS-1526425.

REFERENCES

[1] Heba Abdelnasser, Khaled A Harras, and Moustafa Youssef. 2015. UbiBreathe: A ubiquitous non-invasive WiFi-based breathing estimator. In *ACM MobiHoc*.

[2] 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.

[3] Douglas G Altman and J Martin Bland. 1983. Measurement in medicine: the analysis of method comparison studies. *The statistician* (1983), 307–317.

[4] Heba Aly and Moustafa Youssef. 2016. Zephyr: Ubiquitous Accurate multi-Sensor Fusion-based Respiratory Rate Estimation Using Smartphones. In *IEEE INFOCOM*.

[5] Armin Bidarian-Moniri, Michael Nilsson, Lars Rasmusson, John Attia, and Hasse Ejnell. 2015. The effect of the prone sleeping position on obstructive sleep apnoea. *Acta oto-laryngologica* 135, 1 (2015), 79–84.

[6] Daniel J Buysse, Charles F Reynolds, Timothy H Monk, Susan R Berman, and David J Kupfer. 1989. The Pittsburgh Sleep Quality Index: a new instrument for psychiatric practice and research. *Psychiatry research* 28, 2 (1989), 193–213.

[7] Gustavo Desouzart, Ernesto Filgueiras, and Rui Matos. 2015. Relationship between Postural Reeducation Technique During Sleep and Relaxation Technique in Sleep Quality. *Procedia Manufacturing* 3 (2015), 6093–6100.

[8] Gustavo Desouzart, Ernesto Filgueiras, Rui Matos, and Filipe Melo. 2016. Human Body–Sleep System Interaction in Young Adult Residence. In *Ergonomics in Design Methods & Techniques*. CRC Press, 335–357.

[9] Gustavo Desouzart, Rui Matos, Filipe Melo, and Ernesto Filgueiras. 2016. Effects of sleeping position on back pain in physically active seniors: A controlled pilot study. *Work* 53, 2 (2016), 235–240.

[10] Marian Haescher, Denys JC Matthies, John Trimpop, and Bodo Urban. 2016. SeismoTracker: Upgrade any Smart Wearable to enable a Sensing of Heart Rate, Respiration Rate, and Microvibrations. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 2209–2216.

[11] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H Witten. 2009. The WEKA data mining software: an update. *ACM SIGKDD explorations newsletter* 11, 1 (2009), 10–18.

[12] R Scott Harris, Tilo Winkler, Guido Musch, Marcos F Vidal Melo, Tobias Schroeder, Nora Tgavalekos, and José G Venegas. 2009. The prone position results in smaller ventilation defects during bronchoconstriction in asthma. *Journal of Applied Physiology* 107, 1 (2009), 266–274.

[13] Javier Hernandez, Daniel McDuff, and Rosalind W Picard. 2015. Biowatch: estimation of heart and breathing rates from wrist motions. In *Proceedings of the 9th International Conference on Pervasive Computing Technologies for Healthcare*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 169–176.

[14] Enamul Hoque, Robert F Dickerson, and John A Stankovic. 2010. Monitoring body positions and movements during sleep using wisps. In *ACM Wireless Health*.

[15] Ossi Kaltiokallio, Hüseyin Yiğitler, Riku Jäntti, and Neal Patwari. 2014. Non-invasive respiration rate monitoring using a single COTS TX-RX pair. In *IEEE IPSN*.

[16] Wilburta Q Lindh, Marilyn Pooler, Carol D Tamparo, Barbara M Dahl, and Julie Morris. 2013. *Delmar's comprehensive medical assisting: administrative and clinical competencies*. Cengage Learning.

[17] Jian Liu, Yan Wang, Yingying Chen, Jie Yang, Xu Chen, and Jerry Cheng. 2015. Tracking vital signs during sleep leveraging off-the-shelf wifi. In *ACM MobiHoc*.

[18] Jason J Liu, Wenyao Xu, Ming-Chun Huang, Nabil Alshurafa, Majid Sarrafzadeh, Nitin Raut, and Behrooz Yadegar. 2013. A dense pressure sensitive bedsheet design for unobtrusive sleep posture monitoring. In *IEEE PerCom*.

[19] Sicong Liu, Zimu Zhou, Junzhao Du, Longfei Shangguan, Han Jun, and Wang Xin. 2017. UbiEar: Bringing Location-independent Sound Awareness to the Hard-of-hearing People with Smartphones. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 2 (2017), 17.

[20] Alex Mariakakis, Megan A Banks, Lauren Phillipi, Lei Yu, James Taylor, and Shwetak N Patel. 2017. BiliScreen: Smartphone-Based Scleral Jaundice Monitoring for Liver and Pancreatic Disorders. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 2 (2017), 20.

[21] Rajalakshmi Nandakumar, Shyamnath Gollakota, and Nathaniel Watson. 2015. Contactless sleep apnea detection on smartphones. In *ACM MobiSys*.
[22] Arie Oksenberg and Donald S Silverberg. 1998. The effect of body posture on sleep-related breathing disorders: facts and therapeutic implications. *Sleep medicine reviews* 2, 3 (1998), 139–162.
[23] Ruth Ravichandran, Elliot Saba, Ke-Yu Chen, Mayank Goel, Sidhant Gupta, and Shwetak N Patel. 2015. Wibreathe: Estimating respiration rate using wireless signals in natural settings in the home. In *IEEE PerCom*.
[24] Yanzhi Ren, Chen Wang, Jie Yang, and Yingying Chen. 2015. Fine-grained sleep monitoring: Hearing your breathing with smartphones. In *IEEE INFOCOM*.
[25] Leonid I Rudin, Stanley Osher, and Emad Fatemi. 1992. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena* 60, 1 (1992), 259–268.
[26] Ivan W Selesnick and Ilker Bayram. 2010. Total variation filtering. *White paper* (2010).
[27] Xiao Sun, Zongqing Lu, Wenjie Hu, and Guohong Cao. 2015. SymDetector: detecting sound-related respiratory symptoms using smartphones. In *ACM UbiComp*.
[28] Xiao Sun, Zongqing Lu, Xiaomei Zhang, Marcel Salathé, and Guohong Cao. 2016. Infectious Disease Containment Based on a Wireless Sensor System. *IEEE Access* 4 (2016), 1548–1559.
[29] Shimrit Uliel, Riva Tauman, Michal Greenfeld, and Yakov Sivan. 2004. Normal polysomnographic respiratory values in children and adolescents. *CHEST Journal* 125, 3 (2004), 872–878.
[30] Meng-Chieh Yu, Huan Wu, Jia-Ling Liou, Ming-Sui Lee, and Yi-Ping Hung. 2012. Multiparameter sleep monitoring using a depth camera. In *International Joint Conference on Biomedical Engineering Systems and Technologies*. Springer, 311–325.