

Quantifying Respiratory Variation with Force Sensor Measurements

Joonas Paalasmaa, Lasse Leppäkorpi, and Markku Partinen

Abstract—Measuring the variation of the respiratory rate makes it possible to analyze the structure of sleep. The variation is high when awake or in REM sleep, and decreases in deep sleep. With sleep apnea, the respiratory variation is disturbed. We present a novel method for extracting respiratory rate variation from indirect measurements of respiration. The method is particularly suitable for force sensor signals, because, in addition to the respiratory phenomenon, they typically contain also other disturbing features, which makes the accurate detection of the respiratory rate difficult. Respiratory variation is calculated by low-pass filtering a force sensor signal at different cut-off frequencies and, at every time instant, selecting one of them for the determination of respiration cycles. The method was validated with a single-night reference recording, which showed that the proposed method detects the respiratory variation accurately. Of the 3421 calculated respiration cycle lengths, 95.9% were closer than 0.5 seconds to the reference.

I. INTRODUCTION

THE standard way to evaluate sleep quality and diagnose sleep disorders is polysomnography (PSG). In order to record and score sleep according to international guidelines, measurement of electroencephalography, electromyography and electrooculography is needed. PSG does give detailed information about sleep, but is inconvenient and expensive. Due to the problems of PSG, it is only suited for measuring sleep during a single night, and recording of sleep during several nights at home is difficult. Potential uses of such long-term monitoring are pharmacological intervention studies, and follow-up of patients with different disorders.

Recently, various sensors and analysis methods have been developed for the long-term measurement of sleep, including a bed post sensor [1], a pillow sensor [2] as well as different kinds of mattress sensors [3], [4]. Many of the methods include measurement of respiratory activity, because respiratory dysfunction is associated with sleep apnea. Respiratory function correlates also with the internal structure of healthy sleep. This correlation can be utilized for analyzing the structure of sleep based on respiratory measurements.

Various authors have presented methods for inferring sleep stages based on respiratory measurements [5], [6], [7], [8]. These methods are based on measuring the respiratory variation, which is high when awake or in REM sleep, and decreases in deep sleep. Karlen et al. [5] used

Fourier transformation in the quantification of respiratory variation. Chung et al. [6] calculated the mean respiration rate every 30 seconds. Redmond et al. [7] and de Chazal et al. [8] computed lengths of individual breaths. The latter method is preferable, because it treats every respiration cycle individually rather than averaging over a fixed time window.




Fig. 1. A signal excerpt showing three respiration cycles of a low-pass filtered force signal and an airflow pressure reference signal.

Existing methods for detecting respiration cycle lengths are not well-suited for analyzing all force sensor signals, because some such signals do not have a single repeating pattern at the respiratory frequency, but can have a more complex morphology. That is visualized in figure 1, where, in addition to the respiration phenomenon at 4-second intervals, there is a positive deflection between consecutive respiration cycle peaks. The methods proposed in [7], [8], [9], for example, might detect twice the real respiratory frequency when the disturbing deflections are strong enough. It may be that some force signals do not have a challenging morphology, but we do have noticed the phenomenon with various force sensors. These sensors measure the respiratory activity either from below the mattress or a bedpost.

Detecting respiration cycles from abdominal or airflow measurements is less difficult, as they typically follow the respiration cycle somewhat monotonically and have a single deflection per respiration cycle.

The method proposed in this paper addresses the challenging respiration cycle morphology of force sensor signals by low-pass filtering the signal at different cut-off frequencies and, at every time instant, selecting one of them to be used for the determination of respiration cycle lengths.

II. METHODS

The purpose of the method proposed below is to extract the respiration cycle onset times and respiration cycle lengths from a force sensor signal. The method can be described with the four successive steps below. The steps are visualized with a flowchart in figure 2.

Manuscript received April 14, 2011.

J. Paalasmaa is with Beddit.com Ltd, Espoo, Finland, with Department of Computer Science, University of Helsinki, Helsinki, Finland, and with Department of Signal Processing, Tampere University of Technology, Tampere, Finland (email: joonas.paalasmaa@beddit.com).

L. Leppäkorpi is with Beddit.com Ltd, Espoo, Finland.

M. Partinen is with Helsinki Sleep Clinic, Vitalmed Research Center, Helsinki, Finland and with Department of Clinical Neurosciences, Helsinki University Central Hospital, Helsinki, Finland.

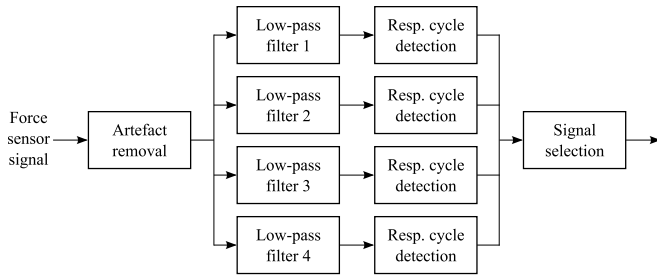


Fig. 2. A flowchart description of the method.

Step 1: Removing movement artefacts

Before the respiration cycles are analyzed, parts of the signal that contain movement artefacts are discarded from further processing. The signal is split into 10-second windows and the peak-to-peak value (difference between maximal and minimal signal value) of each window is calculated. Windows with a peak-to-peak value over twice the average, and the preceding and following 15 seconds, are discarded.

Step 2: Low-pass filtering the signal

The respiration signal is low-pass filtered at cut-off frequencies 0.154 Hz, 0.22 Hz, 0.33 Hz and 0.5 Hz, which result in signals s_1 , s_2 , s_3 and s_4 , respectively.

The cut-off frequencies are selected based on two assumptions: the respiratory frequency is in range 0.1 Hz . . . 0.5 Hz, and the potentially disturbing deflections are at a frequency that is over 1.5 times the respiratory frequency. The assumptions are valid for example for the upper signal in figure 1, where the frequency of the deflection is about twice the respiratory frequency.

Each cut-off frequency equals the previous frequency times 1.5, which guarantees that at least one of the low-pass filtered signals contains the respiration frequency without the disturbing higher-frequency deflections. For example, when the force signal of figure 1 (respiration frequency around 0.25 Hz) is filtered with the four filters, the respiratory frequency remains in signals s_3 and s_4 , but only s_4 contains the disturbing phenomenon, whose frequency is around 0.5 Hz. As s_3 contains only the respiratory frequency without the disturbing deflections, it is easy to detect respiration cycles from it.

The above cut-off frequencies allow measuring respiration cycles with lengths from around 2 to 10 seconds. Adding more low-pass filtered signals below 0.154 Hz and above 0.5 Hz would extend that range.

Step 3: Detecting respiration cycle lengths

After low-pass filtering, the peaks and troughs of signals s_1 , s_2 , s_3 and s_4 are detected by simply finding the local maxima and minima of each signal. Based on them, the respiration cycle lengths and amplitudes are calculated for each signal. The length is defined as the interval between two successive peaks, and the amplitude as the difference of the signal values at the peak and trough.

Step 4: Choosing correct respiration cycle lengths

The previous step results in preliminary respiration cycle lengths for each low-pass filtered signal. In this step, a final sequence of respiration cycle lengths is compiled from the four signals based on the stability of respiration cycle amplitudes in each signal.

The measurement time is divided into three-second intervals. For each interval, the respiration cycle lengths are selected from the low-pass filtered signal whose respiration cycle amplitudes exhibit least variability in the last five respiration cycles before the interval. The variability is calculated as the maximal absolute difference between two successive log-amplitudes, or

$$variability = \max_{2 \leq i \leq 5} |\log(a_i) - \log(a_{i-1})|,$$

where a_1, a_2, \dots, a_5 are the last five respiration cycle amplitudes before the interval. We found that five was a suitable number of respiration cycles for estimating the variability of amplitude. More test subjects are needed in order to find out if that applies to a larger population.

The amplitude variability criterion is based on the assumption that a signal that contains frequencies up to the respiratory frequency is more stable in its amplitude than a signal that also contains higher-frequency disturbing phenomena. See figure 1, where the amplitudes of the force sensor signal alternate between high and low, thus having a high variability. When frequencies above the respiratory frequency (0.25 Hz) are removed with low-pass filtering, the amplitude variation decreases.




Fig. 3. The synthetic signal without the noise component is shown in the topmost plot and signals s_1 , s_2 , s_3 and s_4 in the lower plots. The peak locations are shown with vertical lines.

When the cut-off frequency is below the respiratory frequency, only unsystematic low-frequency phenomena remain. They have a high amplitude variation, so the signal is not selected for the determination of the respiration cycle lengths. See figure 3, where only unsystematic low-frequency fluctuations remain in signal s_1 after 100 seconds.

III. MATERIAL AND RESULTS

A. Synthetic data

The method was evaluated with a synthetic data set to show its theoretical performance. A 300-second signal was created as the sum of the following components, so that it would resemble the signal morphology of the force sensor signal of figure 1:

- A sinusoid whose frequency changes linearly from 0.083 to 0.5 Hz
- A sinusoid whose frequency changes linearly from 0.166 to 1.0 Hz
- Normally distributed independent and identically distributed noise with unit variance

The synthetic signal corresponds to the situation where the respiration cycle length changes from 12 to 2 seconds in 5 minutes. The second sinusoid corresponds to the disturbing upward deflection between respiration peaks and the noise to physiological variability.




Fig. 4. The preliminary respiration cycle lengths from signals s_1 , s_2 , s_3 and s_4 are shown as continuous lines in the four first plots. The thick gray line shows the correct respiration length and the black diamonds are the respiration cycle lengths that were selected from each signal in step 4 of the method. The amplitude variation indexes are shown in the lowermost plot with a logarithmic vertical scale.

The method was applied to the synthetic signal and it succeeded in detecting the correct respiration cycle lengths over the whole signal. The results are shown in figures 3 and 4.

It can be seen in the figures how the correct respiration cycle length is found from the signal where the higher-frequency sinusoid has been filtered out, but not yet the lower-frequency sinusoid.

B. Force sensor data with airflow reference

An overnight polysomnography recording was acquired from a 29-year old female patient at the Helsinki Sleep Clinic, Helsinki, Finland¹. The patient had no limitations of respiratory flow during the night. A 5.5-hour section of the signal was selected for further analysis based on the

¹The study was approved by the Coordinating Ethics Committee of the District Hospital of Helsinki and Uusimaa, Helsinki, Finland.

quality of the airflow pressure channel, which was used as the reference. Other portions of the recording had issues with airflow measurement.

The respiration cycle lengths of the reference were detected by locating the rising zero crossings of the airflow signal, and validated by manual inspection.

The method proposed in this paper was applied to a force sensor signal which was acquired simultaneously with the airflow signal from the bedpost of the patient's bed. The method of signal acquisition is described in [10]. The movement suppression procedure discarded 27% of the recording time, which left 3421 respiration cycles that are not disturbed by movement. The results are shown in figures 5 and 6.




Fig. 5. The respiration cycle lengths that have been selected in step 4 of the method from signals s_1 , s_2 , s_3 and s_4 are plotted above.




Fig. 6. The upper plot shows the respiration cycle lengths detected from the force sensor signal with the proposed method. The reference is plotted in the lower plot.

A relatively large portion of the recording time was discarded, because the patient's sleep is disturbed by frequent limb movements. Two excerpts visualizing the quality of the signals is shown in figure 7.

The accuracy of the method was evaluated quantitatively by estimating the distribution of the deviations of the calculated respiration cycle lengths from the reference. 86.5% of the values deviate less than 0.25 seconds from the reference, 95.9% less than 0.5 seconds and 98.5% less than 1 second.

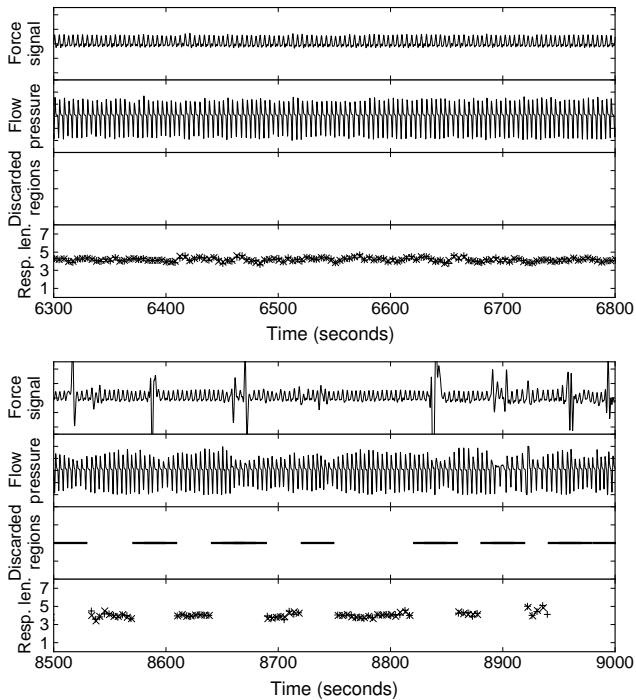


Fig. 7. Two excerpts of the signals acquired at the sleep clinic. The upper excerpt contains no disturbing limb movements, whereas a large portion of the lower excerpt has been discarded due to movement artefacts. The reference respiration cycle lengths are shown as straight crosses and the calculated cycle lengths as diagonal crosses in the lowermost plots.

A graphical description of the deviation is given in figures 8 and 9.




Fig. 8. A histogram density estimate of the deviation of the calculated respiration lengths from the reference.

IV. DISCUSSION

A novel method for detecting respiration cycle lengths from a force sensor signal is presented. The method was evaluated with a synthetic signal and an airflow reference acquired from a sleeping person. The results show that the method is able to follow the respiration rate of the reference accurately and with few incorrect values.

The presented method makes it possible to quantify the variation of the respiration rate precisely. The variation of respiration changes by sleep stage, so the method can work as a building block for respiratory sleep staging methods. For example, the respiration variability features described by [7] can be extracted from force sensor measurements with the proposed method.




Fig. 9. A Bland-Altman plot [11] of the detected respiration cycle lengths.

The method does quantify respiratory variation precisely in the one tested case of healthy breathing. Breathing during sleep can become disturbed, for example, with sleep apnea, and it is unclear how the proposed method would work with disturbed breathing. It is possible that the increased variability of respiration caused by sleep apnea can be detected with the proposed method, but more data is needed to investigate that further.

The selection of the low-pass filtered signal from which the respiration cycle lengths are taken is based on a simple amplitude variation criterion, and that is a part of the method which requires further work and validation with a larger set of test data.

REFERENCES

- [1] M. Brink, C. H. Müller, and C. Schierz, "Contact-free measurement of heart rate, respiration rate, and body movements during sleep," *Behavior Research Methods*, vol. 38, no. 3, pp. 511–521, 2006.
- [2] J. Cha, H. Choi, J. Shin, and K. Lee, "Unconstrained respiration and heart rate monitoring system based on a PPG pillow during sleep," in *30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2008, pp. 3224–3226.
- [3] J. L. Jacobs, P. Embree, M. Gleib, S. Christensen, and P. K. Sullivan, "Characterization of a novel heart and respiratory rate sensor," in *26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, vol. 1, 2004, pp. 2223–2226.
- [4] Y. Chee, J. Han, J. Youn, and K. Park, "Air mattress sensor system with balancing tube for unconstrained measurement of respiration and heart beat movements," *Physiological Measurement*, vol. 26, no. 4, pp. 413–422, 2005.
- [5] W. Karlen, C. Mattiussi, and D. Floreano, "Adaptive Sleep/Wake classification based on cardiorespiratory signals for wearable devices," in *IEEE Biomedical Circuits and Systems Conference*, 2007, pp. 203–206.
- [6] G. Chung, B. Choi, J. Lee, J. Lee, D. Jeong, and K. S. Park, "REM sleep estimation only using respiratory dynamics," *Physiological Measurement*, vol. 30, no. 12, pp. 1327–1340, 2009.
- [7] S. Redmond, P. de Chazal, C. O'Brien, S. Ryan, W. McNicholas, and C. Heneghan, "Sleep staging using cardiorespiratory signals," *Somnologie*, vol. 11, no. 4, pp. 245–256, 2007.
- [8] P. de Chazal, E. O'Hare, N. Fox, and C. Heneghan, "Assessment of sleep/wake patterns using a non-contact biomotion sensor," in *30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2008, pp. 514–517.
- [9] J. M. Kortelainen and J. Virkkala, "PCA model for recording respiration and posture with multichannel BCG sensor in bed mattress," in *4th pHealth Conference*, Chalkidiki, Greece, Jun. 2007.
- [10] J. Paalasmaa, "A respiratory latent variable model for mechanically measured heartbeats," *Physiological Measurement*, vol. 31, no. 10, pp. 1331–1344, 2010.
- [11] J. M. Bland and D. G. Altman, "Statistical methods for assessing agreement between two methods of clinical measurement," *Lancet*, vol. 327, no. 8476, pp. 307–310, 1986.