

RSS-based Respiratory Rate Monitoring using Periodic Gaussian Processes and Kalman Filtering

Roland Hostettler, Ossi Kaltiokallio, Hüseyin Yiğitler,
Simo Särkkä, and Riku Jäntti

This is a post-print of a paper published in *European Signal Processing Conference (EUSIPCO)*.
When citing this work, you must always cite the original article:

R. Hostettler, O. Kaltiokallio, H. Yiğitler, S. Särkkä, and R. Jäntti, “RSS-based respiratory rate monitoring using periodic Gaussian processes and Kalman filtering,” in *25th European Signal Processing Conference (EUSIPCO)*, Kos, Greece, August 2017

DOI:

10.23919/EUSIPCO.2017.8081208

Copyright:

Copyright 2017 EURASIP. First published in the Proceedings of the European Signal Processing Conference (EUSIPCO-2017) in 2017, published by EURASIP.

RSS-based Respiratory Rate Monitoring using Periodic Gaussian Processes and Kalman Filtering

Roland Hostettler*, Ossi Kaltiokallio[†], Hüseyin Yiğitler[†], Simo Särkkä*, and Riku Jäntti[†]

*Department of Electrical Engineering and Automation, Aalto University, Espoo, Finland

E-Mail: {roland.hostettler, simo.sarkka}@aalto.fi

[†]Department of Communications and Networking, Aalto University, Espoo, Finland

E-Mail: {ossi.kaltioallio, huseyin.yigitler, riku.jantti}@aalto.fi

Abstract—In this paper, we propose a method for respiratory rate estimation based on the received signal strength of narrowband radio frequency transceivers. We employ a state-space formulation of periodic Gaussian processes to model the observed variations in the signal strength. This is then used in a Rao-Blackwellized unscented Kalman filter which exploits the linear substructure of the proposed model and thereby greatly improves computational efficiency. The proposed method is evaluated on measurement data from commercially available off the shelf transceivers. It is found that the proposed method accurately estimates the respiratory rate and provides a systematic way of fusing the measurements of asynchronous frequency channels.

I. INTRODUCTION

In the past years, interest toward ubiquitous health monitoring has increased and it has been envisioned that the advances in health monitoring technologies enable future smart homes that would continuously monitor the vital signs of the residents. This information could be used for: i) increasing our health-awareness; ii) diagnostic purposes by medical professionals; and iii) detecting aberrations that need immediate clinical attention. The respiration rate is one essential indicator of a person's health and hence, monitoring it is important [1], [2].

Traditional approaches for measuring the respiratory rate require body contact in which an instrument is attached to the subject's body to measure for example respiratory airflow, chest/abdominal movement, CO₂ concentration, or blood oxygen saturation [2]. Such sensors are unsuitable in home healthcare applications since the person's mobility is restricted, an elderly suffering from dementia may forget to wear a sensor, and infants may remove them. An alternative is to use non-invasive methods that do not rely on body contact and popular approaches include vision-based techniques [3] and radio frequency (RF)-based methods [4]–[6], a research topic we explore in this work.

The propagation delay of the RF signal that reflects from the breathing person's moving chest cavity is quasi-periodic in time [4]. Measuring the propagation delay is possible for example with ultra-wideband radios [4] and frequency modulated carrier wave radars [6]. These works show that respiration rate can be estimated with high accuracy and the technologies can also be used to estimate the heart rate. As a drawback, these systems rely on customized technology and require a large bandwidth.

Narrowband receivers on the other hand can not resolve individual multipath components. However, such devices are cheap and received signal strength (RSS) measurements are ubiquitously available in nearly all radios. In [5], it is shown that sinusoidal changes in the propagation delay also cause periodic variations in the RSS and breathing rate estimation is possible if many transceivers [5] or frequency channels [7] are used. These systems measure the RSS between different wireless links, then the power spectral density (PSD) is calculated over a 15 s–60 s time window for each link, the PSDs are averaged over all links, and the maximum of the mean PSD is used as the breathing rate estimate. This approach is brute force and does not take into account that the RSS measurements for each link are sampled asynchronously and that breathing is not perfectly sinusoidal but also contains higher order harmonics.

In this paper, we develop a recursive, non-parametric approach to respiratory rate tracking. Specifically, we model the oscillatory nature of breathing using a state-space formulation of periodic Gaussian processes and subsequently apply a Rao-Blackwellized unscented Kalman filter for estimation. The method can easily handle asynchronous and missing (e.g. due to packet loss) measurements from multiple sources such as multiple radio channels and does not rely on spectral analysis. The proposed method is similar to the ones proposed in [8], [9] in the context of electrocardiography and oximetry-based respiratory rate monitoring, which use batch-based Gaussian processes. However, the method proposed herein distinguishes itself by using a state-space formulation and Kalman filtering which has the beneficial properties mentioned earlier.

II. MODEL

The chest's motion due to breathing causes changes in the properties of the radio path, which reflects as a change in the RSS. Due to the quasi-periodic nature of this motion, the changes in RSS are quasi-periodic too. These changes can be modeled using, for example, Fourier series or (quasi-)periodic Gaussian processes (GP). In this work, the latter approach is chosen due to its flexibility and robustness [10], [11].

The measured RSS for channel i ($i = 1, \dots, C$) at time t_n can thus be written as

$$y_{i,n} = g_i(t_n) + r_{i,n} \quad (1)$$

where $r_{i,n}$ denotes the measurement noise that is modeled as independent, identically distributed white Gaussian noise as $r_{i,n} \sim \mathcal{N}(0, R_i)$ [12]. Furthermore, $g_i(t)$ is a GP according to

$$g_i(t) \sim \mathcal{GP}(0, k(\tau)) \quad (2)$$

where $\mathcal{GP}(\mu(x), k(x, x'))$ denotes a Gaussian process with mean function $\mu(x)$ and covariance kernel $k(x, x')$, and $\tau = t - t'$. For periodic GPs, a suitable covariance function is the canonical periodic covariance function [11], [13] given by

$$k(\tau) = \sigma^2 \exp\left(-\frac{2 \sin^2\left(\frac{\omega\tau}{2}\right)}{l^2}\right) \quad (3)$$

where σ^2 is the variance magnitude, ω the frequency of the periodicity, and l the GPs length scale.

It has been shown that such periodic GPs can be written in the form of an infinite-dimensional state-space system as [11]

$$\dot{\gamma}_j(t) = A_j \gamma_j(t) + B_j \eta_{\gamma_j}(t) \quad (4a)$$

$$g_i(t) = \gamma_0(t) + \sum_{j=1}^{\infty} H_j \gamma_j(t) \quad (4b)$$

where $\gamma_j(t)$ are 2×1 vectors, $\gamma_0(t)$ is a scalar, $H_j = [1 \ 0]$, and $\eta_{\gamma_j}(t)$ is a white noise process with spectral density S_{γ_j} . In practice, the infinite sum over j is not realizable and the series has to be truncated at some upper bound J .

The matrices A_j , B_j , and S_{γ_j} are determined by the covariance kernel $k(\tau)$. For the kernel in (3) and $j > 0$, these are

$$A_j = \begin{bmatrix} 0 & -j\omega \\ j\omega & 0 \end{bmatrix}, \quad B_j = I_2, \quad (5a)$$

$$S_{\gamma_j} = 2\kappa_j^2 I_2, \quad \kappa_j^2 = 2\sigma^2 \exp(-l^{-2}) \mathcal{I}_j(l^{-2}) \quad (5b)$$

where $\mathcal{I}_j(x)$ is the j th order modified Bessel function of the first kind. Discretization of (4)-(5) yields the discrete-time model

$$\gamma_{j,n} = F_j \gamma_{j,n-1} + q_{\gamma_j,n} \quad (6)$$

with $q_{\gamma_j,n} \sim \mathcal{N}(0, Q_{\gamma_j})$,

$$F_j = \begin{bmatrix} \cos(j\omega\Delta t) & -\sin(j\omega\Delta t) \\ \sin(j\omega\Delta t) & \cos(j\omega\Delta t) \end{bmatrix}, \quad (7a)$$

$$Q_{\gamma_j} = S_{\gamma_j} \Delta t, \quad (7b)$$

and $\Delta t = t_n - t_{n-1}$ for $j > 0$.

For the DC component $\gamma_0(t)$, we have that

$$\dot{\gamma}_0(t) = \eta_0(t), \quad (8a)$$

$$\kappa_0^2 = \sigma^2 \exp(-l^{-2}) \mathcal{I}_0(l^{-2}), \quad (8b)$$

and the discretized version is

$$\gamma_{0,n} = \gamma_{0,n-1} + q_{\gamma_0,n} \quad (9)$$

with $q_{\gamma_0,n} \sim \mathcal{N}(0, Q_{\gamma_0})$, and Q_{γ_0} as in (7).

The frequency $\omega = 2\pi f$ of the GP is normally constant and known (or optimized as one of the GP's hyperparameters). However, here it corresponds to the breathing rate which is the time-varying quantity of interest. Thus, we introduce an

additional state for the breathing frequency $f(t)$. Since $f(t) \geq 0$, we model this as geometric Brownian motion with zero drift. Let $\nu(t) = \log(f(t))$, then, the stochastic differential equation for $\nu(t)$ is given by [14]

$$\dot{\nu}(t) = -\frac{1}{2} S_f^2 + \eta_f(t) \quad (10)$$

where $\eta_f(t)$ is a white noise process with spectral density S_f . Discretization of (10) yields

$$\nu_n = \nu_{n-1} - \frac{1}{2} S_f^2 \Delta t + q_{f,n} \quad (11)$$

and $q_{f,n} \sim \mathcal{N}(0, Q_f)$ with $Q_f = S_f \Delta t$. Note that by introducing a time-varying frequency, the covariance function becomes time-dependent and thus, the GP is non-stationary.

Collecting all the coefficients for the i th channel as

$$\begin{aligned} \gamma_{i,n} &= [\gamma_{0,n} \ \gamma_{1,n}^\top \ \cdots \ \gamma_{J,n}^\top]^\top, \\ \mathbf{q}_{i,n} &= [q_{\gamma_0,n} \ q_{\gamma_1,n}^\top \ \cdots \ q_{\gamma_J,n}^\top]^\top, \\ \mathbf{F}_i &= \text{blkdiag}(1, F_1, \dots, F_J), \\ \mathbf{Q}_i &= \text{blkdiag}(Q_{\gamma_0}, Q_{\gamma_1}, \dots, Q_{\gamma_J}) \end{aligned}$$

the complete discrete-time model for all C channels is

$$\begin{bmatrix} \nu_n \\ \gamma_{1,n} \\ \vdots \\ \gamma_{C,n} \end{bmatrix} = \mathbf{f}(\nu_{n-1}) + \mathbf{F}(\nu_{n-1}) \begin{bmatrix} \gamma_{1,n-1} \\ \vdots \\ \gamma_{C,n-1} \end{bmatrix} + \begin{bmatrix} q_{f,n} \\ \mathbf{q}_{1,n} \\ \vdots \\ \mathbf{q}_{C,n} \end{bmatrix} \quad (12a)$$

$$y_{i,n} = \mathbf{G}_i x_n + r_{i,n} \quad (12b)$$

with

$$\mathbf{f}(\nu_{n-1}) = \begin{bmatrix} \nu_{n-1} - \frac{S_f^2 \Delta t}{2} \\ 0_{C(2J+1) \times 1} \end{bmatrix}, \quad \mathbf{F}(\nu_{n-1}) = \begin{bmatrix} 0_{1 \times C(2J+1)} \\ I_C \otimes \mathbf{F}_i \end{bmatrix},$$

$$\mathbf{G}_i = [0_{1 \times (i-1)(2J+1)} \ 1 \ 1_{1 \times J} \otimes [1 \ 0] \ 0_{1 \times (C-i)(2J+1)}]$$

where \otimes denotes the Kronecker product and the dependency of \mathbf{F}_i on $\omega_{n-1} = 2\pi \exp(\nu_{n-1})$ is implicit. Furthermore, we can define the covariance matrix of the complete process noise vector as $\mathbf{Q} = \text{blkdiag}(Q_f, \mathbf{Q}_1, \dots, \mathbf{Q}_C)$.

Note that while the logarithm of the breathing rate ν_n is a shared state for all channels there is one individual GP for each channel, accounting for the different magnitudes and phases. Also, the measurements $y_{i,n}$ may be completely asynchronous, depending on the communication protocol.

III. ESTIMATION

A. State Estimation

The state-space model (12) can readily be used together with any standard state estimation algorithm such as Kalman or particle filters (and smoothers). However, the model is linear in all states but ν_n for the dynamics and completely linear in the measurements. Thus, in order to not waste computational resources, a Rao-Blackwellized variant should be used to efficiently exploit the linear substructure. We use a Rao-Blackwellized unscented Kalman filter here (see [15], [16]).

The idea is to partition the state vector x_n into states that enter the model non-linearly (x_n^n) and linearly (x_n^l) such that

$$x_n = \begin{bmatrix} x_n^n \\ x_n^l \end{bmatrix} \quad \text{and} \quad P_{n|n} = \begin{bmatrix} P_{n|n}^n & P_{n|n}^{nl} \\ (P_{n|n}^{nl})^\top & P_{n|n}^l \end{bmatrix}$$

with $x_n^n = \nu_n$ and $x_n^l = [\gamma_{1,n}^\top \ \gamma_{2,n}^\top \ \dots \ \gamma_{C,n}^\top]^\top$. This structure can now be exploited to obtain a reduced-complexity filtering algorithm summarized in Algorithm 1 (see, e.g., [15], [16] for details). Here, the $M = 2N + 1$ sigma-points ($N = 1$ is the dimension of the non-linear subspace) and their weights as obtained from the unscented transform are given by [17]

$$x_n^{n,m} = \begin{cases} \hat{x}_{n|n}^n & m = 1 \\ \hat{x}_{n|n}^n + [\delta x]_{m-1} & m = 2, \dots, N+1 \\ \hat{x}_{n|n}^n - [\delta x]_{m-1} & m = N+2, \dots, M \end{cases} \quad (13a)$$

$$w_{\mathcal{M}}^1 = \frac{\lambda}{N + \lambda}, \quad (13b)$$

$$w_{\mathcal{C}}^1 = \frac{\lambda}{N + \lambda} + (1 - \alpha^2 + \beta), \quad (13c)$$

$$w_{\mathcal{M}}^m = w_{\mathcal{C}}^m = \frac{\lambda}{2(N + \lambda)}, \quad m = 2, \dots, M \quad (13d)$$

where $[\cdot]_m$ denotes the m th column of the matrix argument, $\delta x = \sqrt{(N + \lambda)P_{n|n}^n}$, $w_{\mathcal{M}}^m$ and $w_{\mathcal{C}}^m$ are the weights of the mean and covariance, respectively, $\lambda = \alpha^2(N + \kappa) - N$, and α , β , and κ are the parameters of the unscented transform.

It is important to point out that the Kalman filter together with the model (12) easily accommodate asynchronous measurement updates. This is especially useful for the problem considered here as the RSS measurements arrive asynchronously due to the channel hopping nature of the communication protocol.

Also note that by using Rao–Blackwellization, the number of sigma-points is reduced to 3 as compared to using an unscented Kalman filter without Rao–Blackwellization which would require $2[C(2J + 1) + 1] + 1$ sigma-points.

B. Preprocessing

Because of the limited dynamic range of low-cost, off-the-shelf radio transceivers and the small amplitude of the RSS variations, significant quantization is sometimes observed. In particular, the RSS is often only reported in terms of integer values and the signal has to be preprocessed before the Kalman filter introduced above can be employed. A pragmatic solution to this is to recover the underlying signal (12b) by applying a low-pass filter to the quantized signal. For this purpose, a 5th order elliptical FIR low-pass filter with cutoff and stop-band frequencies $f_p = 1$ Hz and $f_s = 1.2$ Hz, respectively, passband ripple $R_p = 0.05$ dB and stop-band attenuation $R_s = 40$ dB has been used in this work. Strictly speaking, this does not exactly recover (12b) (e.g. the measurement noise will be correlated), but we found this to generally work well.

C. Initialization

A well initialized filter is key to achieve good performance. The crucial states here are the frequency f_n (or rather its

Algorithm 1 Rao–Blackwellized Unscented Kalman Filter Iteration

1) Prediction

- a) Calculate the sigma-points $\{x_{n-1}^{n,m}\}_{m=1}^M$ according to (13)
- b) Calculate the orthogonalization [15]

$$L_{n-1} = (P_{n-1|n-1}^{nl})^\top (P_{n-1|n-1}^n)^{-1}$$

$$\tilde{x}_{n-1}^{l,m} = \hat{x}_{n-1|n-1}^l + L_{n-1}(x_{n-1}^{n,m} - \hat{x}_{n-1|n-1}^n)$$

$$\tilde{P}_{n-1}^l = P_{n-1|n-1}^l - L_{n-1}P_{n-1|n-1}^n L_{n-1}^\top$$

- c) Propagate the state

$$x_n^m = f(x_{n-1}^{n,m}) + F(x_{n-1}^{n,m})\tilde{x}_{n-1}^{l,m}$$

- d) Predict the state and covariance

$$\hat{x}_{n|n-1} = \sum_{m=1}^M w_{\mathcal{M}}^m x_n^m$$

$$P_{n|n-1} = \sum_{k=1}^K w_{\mathcal{C}}^m [(x_n^m - \hat{x}_{n|n-1})(x_n^m - \hat{x}_{n|n-1})^\top + F(x_{n-1}^{n,m})\tilde{P}_{n-1}^l (F(x_{n-1}^{n,m}))^\top + Q]$$

2) Measurement update

$$S_n = G_i P_{n|n-1} G_i^\top + R_{i,n}$$

$$K_n = P_{n|n-1} G_i^\top S_n^{-1}$$

$$\hat{x}_{n|n} = \hat{x}_{n|n-1} + K_n (y_{i,n} - G_i \hat{x}_{n|n-1})$$

$$P_{n|n} = P_{n|n-1} - K_n S_n K_n^\top$$

logarithm ν_n) and the DC components $\gamma_{0,n}$. If the frequency is initialized poorly, there is a big risk that the filter will not converge to the correct frequency — which is a common challenge in recursive frequency estimation. Fortunately, the range of the breathing rate for healthy humans is relatively narrow, approximately between 12 bpm and 18 bpm (0.2 Hz–0.3 Hz) [18]. Thus, the initial breathing rate is set to the mean of these two, that is, $\hat{\nu}_{0|0} = \log(15/60)$ with covariance such that one-sigma covers the whole range, that is, $\text{Cov}\{\nu_0\} = \frac{1}{4}(\log(18/60) - \log(12/60))^2$. However, a suitable initial breathing rate could also be determined from a small batch of initialization data.

Proper initialization of the DC component $\gamma_{0,n}$ helps to speed up convergence. Since the amplitudes of the oscillations are generally relatively small compared to the mean RSS, a suitable initialization is simply the first RSS measurement $y_{i,1}$.

The remaining states (the coefficients of the higher orders) are not critical. Good initialization will help to speed up convergence, but initializing them to zero will not affect the performance significantly. Thus, these are set to zero.

IV. RESULTS

A. Experiment Setup

The breathing rate of a person is monitored using two IEEE 802.15.4 compliant narrowband transceivers that are deployed




Fig. 1. Example of the measured RSS signal for channel $i = 1$ (top), the estimated breathing rates for the proposed method (---) and the spectrum estimation-based approach (.....) (middle) and both methods' absolute error (bottom) for 14 bpm.

on opposite sides of a bed, two meters apart from each other. In the experiments, the person is lying still in between the transceivers and breathing at a predefined rate. The transceivers communicate on $C = 16$ frequency channels covering a band between 2.405 GHz and 2.480 GHz and the sampling rate for each channel is on average 32 ms. The transceivers are equipped with low-resolution A/D converters and the RSS is measured with ± 1 dB resolution. The experimental procedure follows that of experiment no. 1 in [7] and the reader is referred to that work for more details about the experimental setup.

In Section IV-B, the model parameters are set to $S_f = \sigma^2 = (1 \times 10^{-3})^2$, $l = 0.1$, $J = 4$, and $R_i = 0.25^2$. The filter parameters are $\alpha = 1$, $\beta = 0$, and $\kappa = 1$. Furthermore, the prefilter discussed in the previous section is applied to each channel. The proposed method is benchmarked against a Kalman filter-based spectrum estimation method [19] which is applied for each radio channel individually and yields similar performance as the commonly used batch-based PSD estimation methods. If multiple channels are used, the estimated spectra are averaged before the person's breathing rate is estimated as done in [7]. The breathing rate is estimated by taking the frequency of the highest peak in the range of 0.1 Hz to 1 Hz, excluding the strong DC term. The frequency grid spacing was chosen to be $\Delta f = 0.01$ Hz.

B. Results and Discussion

Fig. 1 and Fig. 2 show two examples of the RSS signal together with the estimated breathing rate as well as the absolute value of the estimation error $e_{i,n} = \hat{f}_n - f_n$ of the proposed method and the spectrum estimation-based method when using a single channel (channels 1 and 3, respectively). For channel 1, the RSS shows clear oscillations and both methods converge in approximately the same time to the true rate of 14 bpm.




Fig. 2. Example of the measured RSS signal for channel $i = 3$ (top), the estimated breathing rates for the proposed method (---) and the spectrum estimation-based approach (.....) (middle) and both methods' absolute error (bottom) for 14 bpm.

However, the spectrum-based method shows a jumpy behavior during convergence which is due to how the rate is estimated in this approach (selecting the peak frequency). For the proposed method, the estimate converges more smoothly.

For channel 3 (Fig. 2), the measured RSS exhibits much less clear oscillations of the base frequency. This causes problems in the spectrum-based method since the higher harmonics are stronger in this case. Using this approach, it is not straightforward to resolve this and the method does not converge. The GP-based method on the other hand copes with this since there is no requirement for the base frequency to exhibit the highest magnitude. Thus, the estimate converges as quickly as for channel 1.

Fig. 3 shows the estimated breathing rate and the absolute estimation error for the case when all channels are used simultaneously where the GP-based method simply performs asynchronous updates while the spectra for the different channels are averaged as described above. Again, both methods converge almost equally quickly. An important difference between the methods is that the spectrum-based approach actually exhibits a small and constant steady-state error. This is because of this method being bound to the discrete frequency bins which in this case are separated by 0.01 Hz. However, 14 bpm correspond to 0.233 Hz and thus, the closest bin is at 0.23 Hz. The GP method does not suffer from this limitation.

Finally, Fig. 4 illustrates the time-averaged root mean squared error (RMSE) $e_{\text{RMSE}} = \sqrt{\frac{1}{T} \sum_n e_{i,n}^2}$ for different breathing rates using all channels (similar figures could be shown for each of the 16 individual channels but that would not provide additional insight). During the transient period (0 s–30 s), the proposed method attains a significantly lower RMSE compared to the reference method. The main reason for this is the large error due to the jumpy behavior as observed in Fig. 3,




Fig. 3. Example of the estimated breathing rate (top) and the absolute errors (bottom) using all $C = 16$ channels for the proposed method ($- \cdot - \cdot$) and the spectrum estimation-based approach ($\cdot \cdot \cdot \cdot$) for 14 bpm.




Fig. 4. Time-averaged RMSE for the different breathing rate during convergence (0 s – 30 s; top) and after convergence (> 30 s; bottom) for the proposed method (■) and the spectrum estimation method (■).

which inflates the RMSE. It is also worth pointing out that the GP approach has higher RMSEs for 12 bpm, 18 bpm, and 20 bpm compared to 14 bpm and 16 bpm. This is due to the initialization of the breathing rate (15 bpm) which affects the convergence time and results in the higher RMSE for these cases. After convergence, the RMSE is much lower and of the same magnitude for both methods. The main factor contributing to the higher RMSE is how close the closest frequency bin is for the spectrum-based method, which determines the steady-state error. Note that for 12 bpm (0.2 Hz), the closest frequency bin coincides with the true breathing rate exactly and zero RMSE is achieved for the spectrum method.

V. CONCLUSIONS

In this paper, we presented a novel method for respiratory rate monitoring based on periodic Gaussian processes and Kalman filtering motivated by the variations in RSS measurements of narrowband radio transceivers. It has been shown that the method is capable of estimating the breathing rate recursively and optimally fusing the different channels, without requiring windowing or batch-processing. Furthermore, missing measurements do not pose a problem for the proposed method

since there is no requirement of equispaced sampling. Despite the fact that the method has been introduced in the light of RSS measurements, it can readily be generalized to any type of measurements where the respiratory rate manifests in quasi-periodic signals (e.g. ultra-wideband transceivers or chest-mounted accelerometers).

ACKNOWLEDGMENT

Financial support of the Academy of Finland under grants no. #266940, #295080, and #299099 is gratefully acknowledged.

REFERENCES

- [1] M. A. Cretikos, R. Bellomo, K. Hillman, J. Chen, S. Finfer, and A. Flabouris, "Respiratory rate: the neglected vital sign," *Medical Journal of Australia*, vol. 188, no. 11, pp. 657–659, 2008.
- [2] F. Q. AL-Khalidi, R. Saatchi, D. Burke, H. Elphick, and S. Tan, "Respiration rate monitoring methods: A review," *Pediatric pulmonology*, vol. 46, no. 6, pp. 523–529, 2011.
- [3] H.-Y. Wu, M. Rubinstein, E. Shih, J. Guttag, F. Durand, and W. Freeman, "Eulerian video magnification for revealing subtle changes in the world," *ACM Transactions on Graphics*, vol. 31, no. 4, July 2012.
- [4] S. Venkatesh, C. R. Anderson, N. V. Rivera, and R. M. Buehrer, "Implementation and analysis of respiration-rate estimation using impulse-based UWB," in *IEEE Military Communications Conference*, Oct 2005, pp. 3314–3320.
- [5] N. Patwari, J. Wilson, S. Ananthanarayanan, S. K. Kasera, and D. R. Westenskow, "Monitoring breathing via signal strength in wireless networks," *IEEE Transactions on Mobile Computing*, vol. 13, no. 8, pp. 1774–1786, Aug 2014.
- [6] F. Adib, H. Mao, Z. Kabelac, D. Katabi, and R. C. Miller, "Smart homes that monitor breathing and heart rate," in *33rd ACM Conference on Human Factors in Computing Systems*, 2015, pp. 837–846.
- [7] O. Kaltiokallio, H. Yigitler, R. Jantti, and N. Patwari, "Non-invasive respiration rate monitoring using a single COTS TX-RX pair," in *13th International Symposium on Information Processing in Sensor Networks*, April 2014, pp. 59–69.
- [8] M. A. F. Pimentel, D. A. Clifton, L. Clifton, and L. Tarassenko, "Probabilistic estimation of respiratory rate using Gaussian processes," in *35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, July 2013, pp. 2902–2905.
- [9] M. A. F. Pimentel, P. H. Charlton, and D. A. Clifton, "Probabilistic estimation of respiratory rate from wearable sensors," in *Wearable Electronics Sensors: For Safe and Healthy Living*, S. C. Mukhopadhyay, Ed. Springer International Publishing, 2015, pp. 241–262.
- [10] S. Särkkä, A. Solin, A. Nummenmaa, A. Vehtari, T. Auranen, S. Vanni, and F.-H. Lin, "Dynamic retrospective filtering of physiological noise in BOLD fMRI: DRIFTER," *NeuroImage*, vol. 60, no. 2, pp. 1517–1527, 2012.
- [11] A. Solin and S. Särkkä, "Explicit link between periodic covariance functions and state space models," in *Proceedings of the 17th International Conference on Artificial Intelligence and Statistics*, 2014, pp. 904–912.
- [12] H. Yigitler, R. Jantti, and N. Patwari, "On log-normality of RSSI in narrowband receivers under static conditions," *IEEE Signal Processing Letters*, vol. 24, no. 4, pp. 367–371, April 2017.
- [13] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [14] B. Øksendal, *Stochastic Differential Equations: An Introduction with Applications*, 6th ed. Springer, 2010.
- [15] M. R. Morelande and B. Moran, "An unscented transformation for conditionally linear models," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, April 2007, pp. 1417–1420.
- [16] J.-O. Nilsson, "Marginalized Bayesian filtering with Gaussian priors and posteriors," *ArXiv e-prints*, July 2016, arXiv:1603.06462v2.
- [17] E. A. Wan and R. Van Der Merwe, "The unscented Kalman filter for nonlinear estimation," in *Adaptive Systems for Signal Processing, Communications, and Control Symposium*, 2000, pp. 153–158.
- [18] K. E. Barrett, S. M. Barman, S. Boitano, and H. Brooks, *Ganong's Review of Medical Physiology*, 25th ed. McGraw-Hill, 2015.
- [19] Y. Qi, T. P. Minka, and R. W. Picard, "Bayesian spectrum estimation of unevenly sampled nonstationary data," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, May 2002, pp. 1473–1476.