35th Annual International Conference of the IEEE EMBS
Osaka, Japan, 3 - 7 July, 2013

Probabilistic Estimation of Respiratory Rate using Gaussian Processes

Marco A.F. Pimentel, David A. Clifton, Lei Clifton, and Lionel Tarassenko

Abstract— **The presence of respiratory information within the electrocardiogram (ECG) signal is a well-documented phenomenon. We present a Gaussian process framework for the estimation of respiratory rate from the different sources of modulation in a single-lead ECG. We propose a periodic covariance function to model the frequency- and amplitude-modulation time series derived from the ECG, where the hyperparameters of the process are used to derive the respiratory rate. The approach is evaluated using data taken from 40 healthy subjects each with 2 hours of monitoring, containing ECG and respiration waveforms. Results indicate that the accuracy of our proposed method is comparable with that of existing methods, but with the advantages of a principled probabilistic approach, including the direct quantification of the uncertainty in the estimation.**

Index Terms—**Respiratory rate, Gaussian processes.**

I. INTRODUCTION

Respiratory rate has been shown to be an important indicator of patient deterioration [1], [2], and extreme values of respiratory rate are associated with an increased risk of adverse events in hosptial patients [1], [2], [3]. The importance of assessing respiratory rate has led to its inclusion in most numerical patient assessment systems, often termed *early warning scores* or EWS [4], the use of which is widespread. Such systems typically consist of the application of univariate scoring criteria to observational physiological variables (including the vital signs) and produce a cumulative score that can, if it exceeds a predefined threshold, lead to identification of physiological deterioration in acutely-ill hospital patients.

While automated techniques exist for measuring respiratory rate, they usually require the use of equipment which might interfere with natural breathing, such as spirometry, or might be uncomfortable for the patient, such as measurement via a band that encircles the chest. The signals acquired from conventional methods, including impedance plethysmography (IP), are often unusable as a result of a poor signal-to-noise ratio and are sensitive to frequent movement artefact [5].

The ECG signal, recorded for most acutely ill patients, has been considered as a source of potentially reliable respiratory information. Respiratory activity may cause the ECG to be modulated in two fundamental ways: *R-peak amplitude* (RPA) modulation, which is caused by the movement of the chest due to the filling and emptying of the lungs (which in turn causes a rotation of the electrical axis of the heart and the consequent modulation of the amplitude of the ECG) [6], and *respiratory sinus arrhythmia* (RSA), which is a frequency modulation, corresponding to a variation in heart rate that occurs throughout the respiratory cycle [6], [7].

Much previous work exists concerning the estimation of respiratory rate from ECG or other signals, such as the photoplethysmogram or PPG, and the arterial blood pressure waveform or ABP [8]. These approaches are based on either the RPA- or RSA-modulated signals (or a combination of both), and use a variety of algorithms based on spectral analysis [9], the continuous wavelet transform [10], neural networks [11], and autoregressive models [12], [13]. Small errors (around 1 to 2 breaths per minute) between estimates derived from these signals and reference respiratory rate values have been reported [9], [12], [13] for studies of healthy volunteers. The main drawback of these approaches, however, is that they provide a point estimate of the respiratory rate. The uncertainty associated with the estimated value cannot be directly quantified, due to the nature of the algorithms employed. The failure of existing methods to estimate respiratory rate accurately in actual patients, rather than healthy volunteers, motivates the probabilistic approach.

We propose a method that uses the framework of Gaussian process (GP) regression to extract respiratory rate from the different sources of modulation in a single-lead ECG. This brings all of the advantages of a principled, probabilistic approach: our uncertainty in the estimation is directly quantified; incompleteness, noise, and artefact may be handled in a robust manner; and the output may consist of a predictive posterior distribution, rather than a single estimate - this is useful if the estimate of the respiratory rate is to be used as the input to a subsequent probabilistic inference system, where knowing the full distribution of the input is more informative than a point estimate. Finally, due to the generative nature of the approach, it is possible to generate data from the model, which can be useful for estimating the behaviour of the respiratory rate during periods of missing data.

II. METHODOLOGY

For the preliminary analysis described in this paper, we tested our method using the Physiobank/Physionet *Fantasia* database [14], [15]. The latter consists of data acquired from two subgroups of volunteers: 20 "young" (21 - 34 years old) and 20 "elderly" (60 - 85 years old) healthy subjects who underwent 120 minutes of continuous supine rest (while watching the film "Fantasia"). Continuous single-lead ECG

M.A.F. Pimentel, D.A. Clifton, L. Clifton, and L. Tarassenko are with the Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford, Old Road Campus Research Building, Headington, Oxford OX3 7DQ, UK. (*e-mail: marco.pimentel@eng.ox.ac.uk)

978-1-4577-0216-7/13/$26.00 ©2013 IEEE

and respiration (IP) signals were collected. Each subgroup of subjects includes equal numbers of men and women.

Respiratory rate was computed using methods to be described later, with windows of data of 1 minute duration, with successive windows having 50 s overlap (i.e., a new estimate is produced every 10 s).

A. Extracting the respiratory waveforms

ECG beat detection was performed using the Hamilton and Tompkins algorithm [16]. The amplitude of the resulting series of R-peaks was determined in order to derive the RPA waveform. The intervals between successive R-peaks were also calculated to derive the R-R time series, which corresponds to the RSA waveform.

B. Calculation of respiratory rate from the RPA and RSA waveforms

1) Proposed method: The proposed approach to extract respiratory rate from the RPA and RSA waveforms is based on Gaussian process regression. This offers a framework for performing inference using time-series data, in which a probability distribution over a *functional space* is constructed. We consider the RPA and RSA waveforms to be two separate functions, from which we can perform inference using the GP framework. Our approach is particularly suited to the analysis of data that may be sampled at irregular intervals, as with the R-peaks and R-R intervals contained in the RPA and RSA waveforms, respectively.

The GP is a stochastic process [17] that expresses a dependent variable $\mathbf{y}$ in terms of an independent variable $\mathbf{x}$, via a latent function $\mathbf{f}(\mathbf{x})$. This function can be interpreted as being a probability distribution over functions,

$$\mathbf{y} = f(\mathbf{x}) \sim \mathcal{GP}\left(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')\right) \tag{1}$$

where $m(\mathbf{x})$ is the mean function of the distribution and k is a covariance function which describes the coupling between two values of the independent variable as a function of the (kernel) distance between them. The covariance function encodes our assumptions concerning the structure of the time series that we wish to model [17]. Valid covariance functions can take a variety of forms, with the constraint that they are positive semi-definite.

In the case considered by this paper, $\mathbf{x}, \mathbf{y}$ will be the times of the observations and the values of the observations, respectively, comprising univariate pairs (x, y). Denoting $r = \parallel x_p - x_q \parallel$ as the (Euclidean) distance between two independent variables, x_p and x_q, we propose a periodic covariance function:

$$k(r) = \sigma_0^2 \exp\left[-\frac{\sin^2((2\pi/P_L)r)}{2\lambda^2}\right] \tag{2}$$

in which the hyperparameters σ_0 and λ give the amplitude and length-scale of the latent function. P_L is the length of the period and is the key parameter for estimation of the respiratory rate. Selection of k represents our prior knowledge of typical waveforms from which respiration may be derived, which are expected to be periodic, and where the period can be related to the respiratory rate. We assume that the observations are corrupted by additive i.i.d. Gaussian noise with variance component ε^2. Thus, the full covariance function is given by

$$V(x_p, x_q) = k(x_p, x_q) + \varepsilon^2 \delta\Big(\parallel x_p - x_q \parallel \Big) \tag{3}$$

where δ is the Kronecker delta, for which $\delta = 1$ if $p = q$, and $\delta = 0$ otherwise. Here, ε is the noise variance. The values of the hyperparameters are learned from univariate respiration waveforms which consist of n observations, $D = \{(x_i, y_i) \mid i = 1, ..., n\}$. The x_i and y_i points represent the independent and dependent variable values respectively.

The natue of the GP is such that, conditional on observed data, predictions can be made about the function values $f(x_*)$ at any ("test") location of the index set x_*. The distribution of the values of f at point x_* is Gaussian, with mean and covariance given by

$$f_* | x_*, \mathbf{x}, \mathbf{y} \sim \mathbf{N}\Big(\overline{f}_*, Var[f_*]\Big) \tag{4}$$

in which $\mathbf{x}, \mathbf{y}$ are the "training data", D.

The above allows us to determine the following predictive equations for GP regression, for which we assume the mean function m to be zero,

$$\overline{f}_* = m(x_*) + \mathbf{k}(x_*, x_*)^\top \mathbf{V}(\mathbf{x}, \mathbf{x})^{-1}\big(\mathbf{y} - m(\mathbf{x})\big) \tag{5}$$

$$Var[f_*] = \mathbf{k}(x_*, x_*) - \mathbf{k}(\mathbf{x}, x_*)^\top \mathbf{V}(\mathbf{x}, \mathbf{x})^{-1} \mathbf{k}(\mathbf{x}, x_*) \tag{6}$$

In the above, we use boldface terms $\mathbf{k}, \mathbf{V}$ to refer to the matrix equivalents of k, V defined earlier. For the purposes of this investigation, empirical selection of suitable prior values of the hyperparameters was $\sigma_0 = 1$, $\lambda = 1$, and $\varepsilon = 0.01$. For the period P_L, a set of appropriate priors was selected $P_L = (1.5...10)$, which corresponds to plausible respiratory rate values of 6 to 40 breaths per minute.

A full Bayesian treatment of GP regression requires integration over the posterior distribution of the hyperparameters. Even though most calculations in the GP regression framework are analytically tractable, the integral over the posterior of the hyperparameters often is not. The integration over the posterior of the hyperparameters $p(\theta|D)$, with $\theta = \{\sigma_0, \lambda, P_L, \varepsilon\}$, can be approximated by a point via the *maximum a posteriori* (MAP) estimate

$$\hat{\theta} = \arg\max_\theta p(\theta|D) \tag{7}$$

$$= \arg\min_\theta \Big[-\log p(D|\theta) - \log p(\theta) \Big] \tag{8}$$

In this approximation, the distribution over the hyperparameters is assigned a point mass at the mode of the posterior, allowing the marginal distribution of the latent function to be approximated by $p(f|D) \approx p(f|D, \hat{\theta})$. This approach is computationally attractive. The grid search approximation to the full integral over the posterior distributions of the hyperparameters follows the work of Rue et al. [18], in which the posterior mode $\hat{\theta}$ is first located by maximising the log-posterior distribution $\log p(\theta|y)$, and the shape of the log-posterior is approximated with a Gaussian, the covariance

of which is the inverse of the negative Hessian at the mode (more details may be found in [17], [18]).

For the work described in this paper, for each 60 s window, we first normalise the RPA and RSA time series using a zero-mean, unit-variance transformation. A Gaussian process is then fitted to each of the waveforms, using the procedure described above to obtain an estimate of both the value and uncertainty of the respiratory rate value (directly from the distribution over the period, P_L). The estimate with lower uncertainty (i.e., where the posterior distribution had lowest variance) was chosen as the final estimate of the respiratory rate for that window.

2) Autoregressive model: The performance of our proposed method was compared to that of the autoregressive model, a parametric spectral analysis technique that has been successively applied to this problem [12], [13]. This method requires equi-spaced samples, and so the time series of R-R intervals and R-peaks were resampled at 4 Hz, and the RPA and RSA waveforms were obtained using linear interpolation. Each of the waveforms was then filtered using an FIR band-pass filter with cut-off frequencies of 0.1 and 0.6 Hz (equivalent to respiratory rates of 6-36 breaths per min). Following previous methods, the steps involved in estimating respiratory rates are as follows for each of the RPA and RSA waveforms: (i) fit an AR model to each 60 s window of resampled data; (ii) reject all poles with frequencies corresponding to respiratory rates outside the range of physiologically-plausible values (6-36 breaths per min); (iii) keep all poles with magnitude of at least 95% of the remaining highest-magnitude pole; and (iv) select the pole remaining that has the smallest angle. The frequency associated with this pole is the estimate of respiratory rate for this waveform. Finally, the respiratory rate corresponding to the pole with the highest magnitude (extracted from the two waveforms) was selected as the final respiratory rate for that window.

C. Reference respiratory rate

We calculated the reference respiratory rate from the IP signal in the database using both an extrema detection algorithm and an AR-based method [13]. With the latter, we down-sampled the respiration signal to 4 Hz, after applying an anti-aliasing filter, and then applied a 0.1-0.6 Hz FIR band-pass filter. We then fitted the resulting waveform to an AR model and identified the respiratory pole using the same method as described in the previous section. Only those reference waveforms for which the agreement between both extrema- and AR-based estimates was within 3 breaths per min were retained ("valid windows"). As a result, 72% of the available windows were deemed to be "valid". This approach ensures only the highest quality reference values are considered by potentially eliminating regions of low IP quality.

III. Results and Discussion

Over the entire database, the mean reference respiratory rate was 18.3 ± 2.9 breaths per min (17.9 ± 2.8 for the young subjects and 18.8 ± 3.0 for the elderly subjects).

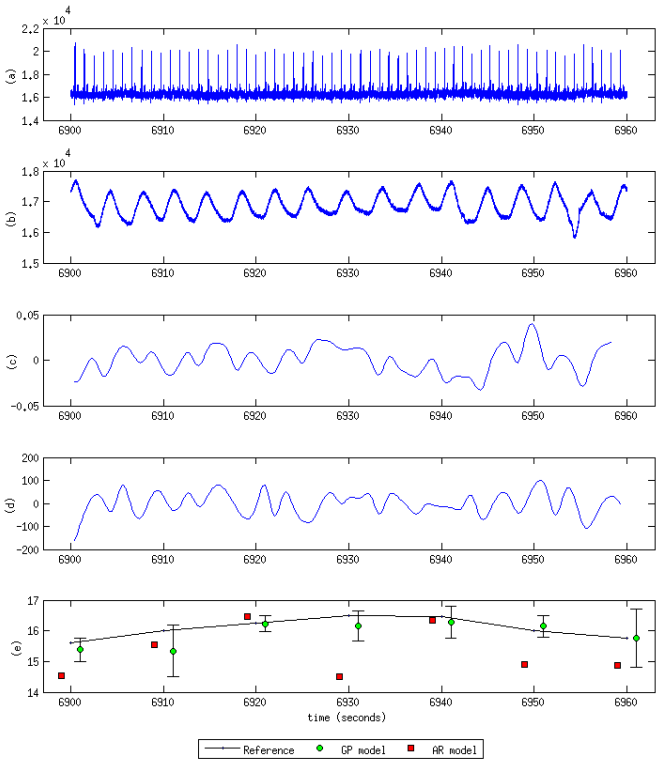

Fig. 1. Extraction of respiratory rate from single-lead ECG. (a) ECG signal; (b) Reference (IP) respiration signal; (c) RSA waveform from the R-R intervals time series; (d) RPA waveform from the R-peaks time series; (e) Final respiratory rate estimates from the AR and GP-based methods.

Figure 1 shows an example from a 1-minute window of data. In general, as illustrated in the figure, we observe that the values extracted using the AR method and the proposed GP method are close to the corresponding reference respiratory rate. However, using the GP-based approach, we can explicitly quantify our uncertainty in the estimated value by computing the variance of the posterior distribution drawn from the related period length parameter (P_L) of the GP. The uncertainty of an estimate may be due to the presence of noise in the derived respiration waveform (which in turn may be caused by a bad performance of the beat detector), which precludes an accurate estimation of the respiratory rate. In such cases, the precision (inverse of the variance) of the estimate is very low.

The performances of the AR and proposed methods were assessed by calculating the mean absolute error (MAE) in breaths per min, $\mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n} \mid \hat{y}_i - y_{ref,i} \mid$, where n is the number of valid windows over the entire database of both groups (young and elderly subjects), $\hat{y}_i$ is the estimated respiratory rate (mean posterior value in the case of the GP-based method) and $y_{ref,i}$ is the reference respiratory rate for window i. Table I shows MAE for different ranges of respiratory rates. While both methods show similar performance, both performed better for the young group of patients. As the respiratory rate increases (or decreases), the estimation errors also increase. Crucially, we note that the proposed method

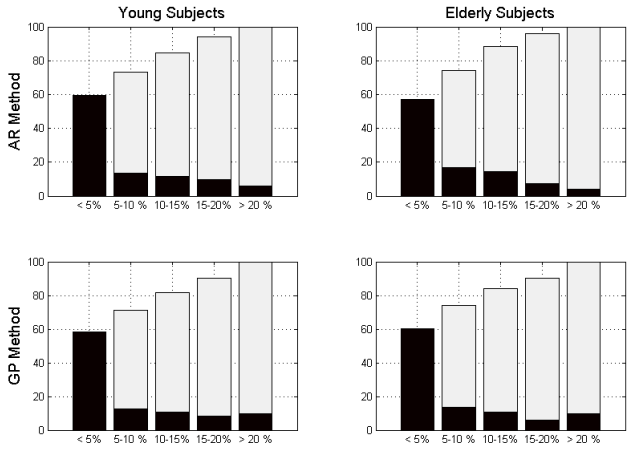

Fig. 2. Histograms showing the percentage of windows for different ranges of the percentage error E (X-axis). The shaded area shows the cumulative percentage of data.

is more accurate for higher respiration rates in the elderly, which is the typical target population.

To assess further the performance of our proposed method, we calculated the percentage of valid windows for different ranges of the percentage error (see Figure 2), which is given by $\mathsf{E} = \frac{\mathsf{MAE}}{|\mu_{ref}|} \times 100$, where μ_{ref} is the mean of the reference respiratory rates over each of the two patient groups. This is a useful metric since the significance of MAE is different depending on the actual respiratory rate. We note that both methods perform similarly.

IV. Conclusion

We have proposed a novel probabilistic approach for extracting respiratory rate from time-series sensor data using Gaussian processes. The method is able to give not only a point estimate of the breathing rate, but, for the first time, a measure of uncertainty of the estimate. By applying this technique to a database of 40 healthy subjects, we have demonstrated that it is possible to match the performance of the existing state-of-the-art, while bringing the benefits of a probabilistic framework.

Open-source code implementing this work may be obtained from http://www.robots.ox.ac.uk/∼davidc.

ACKNOWLEDGMENT

MAF Pimentel was supported by the RCUK Digital Economy Programme (Oxford Centre for Doctoral Training in Healthcare Innovation) and FCT - *Fundação para a Ciência e Tecnologia*. DAC was supported by the Centre of Excellence in Personalised Healthcare funded by the Welcome Trust and EPSRC. LC was supported by the NIHR Oxford Biomedical Research Centre.

Range of reference RR	Young subjects		Elderly subjects	
	AR	GP	AR	GP
All data	1.26	1.31	1.62	1.59
≤ 12 bpm	1.66	1.68	1.52	1.57
$12 - 16$ bpm	1.09	1.21	1.35	1.29
$16 - 20$ bpm	1.21	1.19	1.12	1.22
> 20 bpm	1.90	1.71	2.09	1.82

TABLE I

Mean absolute error in breaths per min (bpm)

References

[1] M. Cretikos, J. Chen, K. Hillman, R. Bellomo, S. Finfer, A. Flabouris *et al.*, "The objective medical emergency team activation criteria: A case-control study," *Resuscitation*, vol. 73, no. 1, pp. 62–72, 2007.

[2] J. Kause, G. Smith, D. Prytherch, M. Parr, A. Flabouris, K. Hillman *et al.*, "A comparison of antecedents to cardiac arrests, deaths and emergency intensive care admissions in australia and new zealand, and the united kingdom the ACADEMIA study," *Resuscitation*, vol. 62, no. 3, pp. 275–282, 2004.

[3] J. Fieselmann, M. Hendryx, C. Helms, and D. Wakefield, "Respiratory rate predicts cardiopulmonary arrest for internal medicine inpatients," *Journal of General Internal Medicine*, vol. 8, no. 7, pp. 354–360, 1993.

[4] H. Gao, A. McDonnell, D. Harrison, T. Moore, S. Adam, K. Daly, L. Esmonde, D. Goldhill, G. Parry, A. Rashidian *et al.*, "Systematic review and evaluation of physiological track and trigger warning systems for identifying at-risk patients on the ward," *Intensive Care Medicine*, vol. 33, no. 4, pp. 667–679, 2007.

[5] V. Larsen, P. Christensen, H. Oxhøj, and T. Brask, "Impedance pneumography for long-term monitoring of respiration during sleep in adult males," *Clinical Physiology*, vol. 4, no. 4, pp. 333–342, 1984.

[6] G. Clifford, F. Azuaje, P. McSharry *et al.*, *Advanced methods and tools for ECG data analysis*. Artech House, 2006.

[7] J. Hirsch and B. Bishop, "Respiratory sinus arrhythmia in humans: how breathing pattern modulates heart rate," *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 241, no. 4, pp. H620–H629, 1981.

[8] N. Shamim, M. Atul, C. Gari D *et al.*, "Data fusion for improved respiration rate estimation," *EURASIP Journal on advances in signal processing*, vol. 2010, 2010.

[9] K. Chon, S. Dash, and K. Ju, "Estimation of respiratory rate from photoplethysmogram data using time–frequency spectral estimation," *IEEE Transactions on Biomedical Engineering*, vol. 56, no. 8, pp. 2054–2063, 2009.

[10] D. Clifton, J. Douglas, P. Addison, and J. Watson, "Measurement of respiratory rate from the photoplethysmogram in chest clinic patients," *Journal of Clinical Monitoring and Computing*, vol. 21, no. 1, pp. 55–61, 2007.

[11] A. Johansson, "Neural network for photoplethysmographic respiratory rate monitoring," *Medical and Biological Engineering and Computing*, vol. 41, no. 3, pp. 242–248, 2003.

[12] J. Lee and K. Chon, "Respiratory rate extraction via an autoregressive model using the optimal parameter search criterion," *Annals of Biomedical Engineering*, vol. 38, no. 10, pp. 3218–3225, 2010.

[13] C. Orphanidou, S. Fleming, S. Shah, and L. Tarassenko, "Data fusion for estimating respiratory rate from a single-lead ECG," *Biomedical Signal Processing and Control*, 2012.

[14] Goldberger, A. L. et al., "Physiobank, physiotoolkit, and physionet," *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000 (June 13).

[15] N. Iyengar, C. Peng, R. Morin, A. Goldberger, and L. Lipsitz, "Age-related alterations in the fractal scaling of cardiac interbeat interval dynamics," *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, vol. 271, no. 4, pp. R1078–R1084, 1996.

[16] J. Pan and W. Tompkins, "A real-time QRS detection algorithm," *IEEE Transactions on Biomedical Engineering*, no. 3, pp. 230–236, 1985.

[17] C. Rasmussen and C. Williams, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 1.

[18] H. Rue, S. Martino, and N. Chopin, "Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 71, no. 2, pp. 319–392, 2009.