Classifying signals from a wearable accelerometer device to measure respiratory rate

Gordon B. Drummond[1], Darius Fischer[2], Margaret Lees[3], Andrew Bates[2], Janek Mann[2] and D.K. Arvind[2]

Affiliations: [1]Dept of Anaesthesia, Critical Care, and Pain Medicine, University of Edinburgh, Edinburgh, UK. [2]Centre for Speckled Computing, School of Informatics, University of Edinburgh, Edinburgh, UK. [3]Princess Margaret Eye Pavilion, Edinburgh, UK.

Correspondence: Gordon B. Drummond, Dept of Anaesthesia, Critical Care, and Pain Medicine, University of Edinburgh, Edinburgh Royal Infirmary, 51 Little France Crescent, Edinburgh, EH16 4SA, UK.
E-mail: g.b.drummond@ed.ac.uk

ABSTRACT

Background: Automatic measurement of respiratory rate in general hospital patients is difficult. Patient movement degrades the signal and variation of the breathing cycle means that accurate observation for ⩾60 s is needed for adequate precision.

Methods: We studied acutely ill patients recently admitted to a teaching hospital. Breath duration was measured from a triaxial accelerometer attached to the chest wall and compared with a signal from a nasal cannula. We randomly divided the patient records into a training (n=54) and a test set (n=7). We used machine learning to train a neural network to select reliable signals, automatically identifying signal features associated with accurate measurement of respiratory rate. We used the test records to assess the accuracy of the device, indicated by the median absolute difference between respiratory rates, provided by the accelerometer and by the nasal cannula.

Results: In the test set of patients, machine classification of the respiratory signal reduced the median absolute difference (interquartile range) from 1.25 (0.56–2.18) to 0.48 (0.30–0.78) breaths per min. 50% of the recording periods were rejected as unreliable and in one patient, only 10% of the signal time was classified as reliable. However, even only 10% of observation time would allow accurate measurement for 6 min in an hour of recording, giving greater reliability than nurse charting, which is based on much less observation time.

Conclusion: Signals from a body-mounted accelerometer yield accurate measures of respiratory rate, which could improve automatic illness scoring in adult hospital patients.

@ERSpublications
A machine learning method was developed to classify sections of breathing records from acutely ill patients wearing a small wireless motion sensor. This would allow accurate and automatic measurement, recording, and charting of respiratory rate. https://bit.ly/301P8XW

Introduction

To assess acute illness, scoring systems using measurements such as respiratory rate, heart rate, systolic blood pressure and oxygen saturation are often used in hospital settings [1]. Respiratory rate is a key observation in early recognition of acutely ill patients in hospital and is important in illness scoring, such as in COVID-19 infection [2] where it has prognostic value [3, 4]. Respiratory rate is not commonly measured automatically, although clinicians now recognise the need for reliable wearable monitoring devices to assist management of outbreaks of infectious disease [5].

Current measurements of respiratory rate need to be improved. Breaths are counted for a limited time [6], values recorded by different observers often differ [7] and charted values are not trustworthy [8].

Accurate measurement of respiratory rate with a wearable device has practical difficulties. Artefacts caused by patient movement can cause repeated false alarms. In surveillance monitoring, rapid changes in patient condition are likely to be less than in critical illness and management of artefacts may require a different strategy, such as signal selection.

A poorly recognised factor is that the breathing signal itself is not stable, which can affect comparisons of methods of measurement. Breath duration varies from breath to breath and measurements of rate made using short periods of counting, such as 30 or 60 s, are imprecise. Repeating observations made over 1 min gives results with an interquartile range of 3 breaths per min [6]. When comparing methods, exactly the same breaths should be sampled but previous studies failed to do this [9, 10]. Recent studies of automatic measurement of respiratory rate in low-care and high-dependency patients report similar limits of agreement, about ±4 breaths per min [11–13].

To monitor respiratory rate in a general hospital setting, we need sufficient reliable and accurate information to identify acute illness or deterioration [14, 15]. This should be possible by signal processing to select trustworthy sections of the signal. Although signal quality may sometimes be unreliable, a sample of less frequent but more trustworthy observations could provide adequate input values for a nursing chart and the illness scoring system, and more accurately identify acute illness or deterioration [16, 17].

We used a noninvasive monitor (Respeck) to measure respiration by sensing chest wall movement, transmitting the data by a Bluetooth low-energy connection to a bedside receiver. In patients after surgery, we obtained reliable measures of respiratory rate [18] but these patients were "settled" and inactive. We now assessed how the device performed when used in acutely ill medical patients. We developed an automatic classification process to detect periods when the signal was reliable by identifying features that were present when the signal corresponded with a second, reliable method of measuring breathing. Our comparison method was nasal cannula pressure, which measures gas flow at the nose [19].

The research question was whether the device could be trained to measure respiratory rate accurately and reliably in patients with acute medical illness. An accurate process would give values that agreed with a standard measure, within ±2 breaths per min. Measurements would be reliable if they were sufficient for adequate input into a hospital early warning chart. Currently, national UK recommendations [1] suggest respiratory rate should be measured on admission to hospital and then 12-hourly, and increased up to hourly in patients with severe illness.

Methods

Ethical approval

We wished to recruit patients who had just been admitted to hospital with acute illness. The Scotland A Research Ethics Committee considered the research (ref. 12/SS/0054) met the requirements of section 51 of the Adults with Incapacity (Scotland) Act with regard to research subjects who lack capacity to consent. We were thus not permitted to record patient data that could allow identification.

Patient recruitment

A trained research nurse worked with nursing staff in the acute medical receiving unit of a 570-bed teaching hospital. Patients with medical illnesses are admitted for initial assessment in one ward, and moved to a second ward if further investigation and treatment are required. Most patients spend about 4–8 h in the unit and no patient is in the unit for >24 h.

Patients who might be approached by the research nurse were identified by the clinical nursing staff and given an information leaflet. If the patient moved to the second ward, the study was explained to them, and they were asked if they wished to take part and recruited if they agreed. If possible, signals were to be recorded for 1 h or until the patient was prepared to leave the ward. Patients were allowed to discontinue wearing the devices if they wished.

Patient recordings

The Respeck sensor measures acceleration in three orthogonal axes. The accelerometer signals are digitised and transmitted by Bluetooth LE to a receiver (iPod) for storage and later retrieval and analysis. The Respeck is enclosed in a plastic capsule 45×38×13 mm and weighs 15.4 g. It is marked to allow a consistent orientation when attached to the patient. When attached as described below, a positive signal from the x-axis indicates a force from right to left, on the y-axis from caudal to cranial and on the z-axis from anterior to posterior. The device is sealed in a plastic bag and the bag is taped to the patient's skin using a conforming, perforated polyester fabric tape with acrylic adhesive (Mefix; Mölnlycke Health Care Limited, Dunstable, UK) to hold the capsule firmly on the body surface. The device is placed just under the rib margin, in the midclavicular line, and usually on the left side of the patient.

A single-use nasal cannula (Sleep Sense 15805-2; Medes Ltd, Radlett, UK) was placed below the nostrils to provide a comparison signal. The cannula was connected *via* a bacterial filter to a battery-powered pressure transducer (PTAF2; Respironics). This was placed beside the patient and connected to a base station, which transmitted the pressure signal wirelessly (Bluetooth LE) to the iPod receiver. All signals were digitised at 12.5 Hz.

Details of the reason for admission of the patient to hospital, and the year of birth, height, weight and sex of the patient were recorded, but no other patient data were kept. Recordings were transferred from the receiver to a secure computer for further analysis.

Signal analysis

The nasal flow signal in each patient record was replayed, without displaying the motion sensor signals, using proprietary software (Spike2, version 5.19; CED, Cambridge, UK). The breathing pattern was inspected by a skilled observer unaware of the identity of the patient. Sections of each record that showed clear breathing signals were identified. Breath onset was timed when the nasal pressure became less than a threshold set to be just below atmospheric and automatically marked using a "decrease through threshold" setting. The entire record was visually checked by the observer to avoid any false or missed breaths. After breath onsets had been measured, the patient records were randomly divided into two groups. The first group was used to develop and train a method to classify the Respeck signal (see later) and the second was used to test the classification model performance.

Respeck signal analysis

Gravity is the predominant component of acceleration sensed by the device. As the chest wall expands and changes shape, the angle of the device relative to gravity changes. The angle changes of the axis measurements are smoothed and combined. Previous studies showed a close correspondence between the Respeck signal and nasal air flow [18]. By setting upper and lower threshold limits for the Respeck, we measure the timing of respiratory phases when the signal successively crosses these threshold limits. The limits are determined dynamically by the signal amplitude. Phases of breathing are identified by measuring successive half-breath cycles, with the duration of inspiration taken as the shorter.

Record synchronisation and comparison

Nasal cannula and Respeck signal values were time stamped when stored on the iPod. The Respeck signal was shifted in time by 960 ms to correct for delay introduced by the smoothing function. The respiratory rate measurements from the nasal cannula and the Respeck signals were compared in 20-s windows, with successive windows overlapping by 10 s. These would contain about 5–7 breaths and allowed dynamic comparisons of the two signals. The instantaneous frequency of each breath, nasal cannula and Respeck, was calculated as:

$$\text{Frequency} = 60\left(\frac{1}{\text{breath duration (s)}}\right)$$

The breathing rate for the window was reported as the mean of the instantaneous rates of each complete breath cycle in the 20-s window. The accuracy of the Respeck measurement was defined as the modulus of the difference between nasal cannula and Respeck rates.

Features associated with accurate estimates of respiratory rate

To allow training of the selection method and subsequent testing, satisfactory records were blinded, randomly numbered and sorted for length. To ensure that the sample used for testing was adequate, the shortest four records were temporarily removed. The remainder were then separated into a sample of 50

for training the selection method and a sample of seven for final testing. The short records were then added back to the training group.

In each window of the Respeck signal, features were identified that could show whether the measurements could be trusted. These features were used to classify the Respeck signal windows and only windows that gave a close estimate of the breathing rate were selected [20]. A trusted period was defined as one in which the difference between nasal cannula and Respeck rate was no greater than a specified threshold value. There is an inverse relationship between this threshold and agreement of the estimate with the breathing rate, and the proportion of the record that could be selected. The threshold value was initially set to 1.0 breaths per min. For different values of threshold <1, features of the signals in all the windows were compared for each patient record in the training set. We inspected the distribution of these features to generate a list of those that would possibly classify Respeck signal windows as those with a breathing rate close to the nasal cannula breathing rate (table 1).

The most effective features were selected and ranked using recursive feature elimination and logistic regression. A limited number of features would save on computation, allowing implementation to run on the Respeck device itself. A neural network with two hidden layers was trained to select acceptable windows using a linear rectifier as the activation function [21, 22]. Using leave-one-out cross-validation for each patient, the model was trained for the rest of the patients in the set and then tested on the chosen patient. This prevents the model learning subject-specific patterns before testing. The features, ranked in descending order of effectiveness, were: the standard deviation of the breathing rates, the mean difference of the breathing rates, the number of rate values in the window, the mean of the peak intervals, the standard deviation of the breathing signal, and the difference between the maximum and minimum rates.

Statistical analysis

We used the modulus of the difference between nasal cannula and Respeck rates to assess agreement. For comparison with other reports, we compared the Respeck and nasal cannula measures, before and after selection of reliable signals, using the method of Bland and Altman for repeated measures in multiple subjects [23].

Results

Clinical information

We approached 96 patients for the study. Of these, 78 agreed to participate. In 12 data sets, the nasal pressure records were absent or of poor quality, and in five, the recording was too short to be useful (<10 min). Thus, 61 records were used for analysis. The mean±SD age of these patients was 61±18 years, their body mass index was 28±7 kg·m^{-2} and 35 were female. The most frequent presenting illness was respiratory disease, either exacerbation of chronic obstructive pulmonary disease or asthma (16 patients). Ischaemic heart disease or venous thromboembolism was suspected in 14 patients and 13 had either sepsis or urinary tract infection. The remainder had neurological conditions such as stroke (n=9), gastrointestinal

TABLE 1 Initial features considered for testing in the model	
Factors	**Comment or explanation**
Fourier transform of Respeck signal	Normalised as $\dfrac{\text{Fourier factors in range 9–24 breaths per min}}{\text{sum of all frequency factors}}$
Mean, median and standard deviation of individual breathing rates	Extreme values are less likely to be reliable. A smaller standard deviation reflects more regular breathing
Maximum–minimum rate	A large range of rates within the sampling window would be less trustworthy
Mean and median of differences	Mean and median of frequency differences between one breath and the next within the window
Standard deviation of axis acceleration values	A measure of large changes in axis signal, often associated with posture changes
Standard deviation, median and skew of activity levels	Activity level is computed as the length of the vector differences using the x, y and z coordinates
Standard deviation of Respeck signal	A large standard deviation indicates probable posture change
Number of valid breaths	A small number of detected breaths suggests the signal is unreliable
Standard deviation of Respeck amplitude peaks	Greater variation in the size of breaths in a window suggests unstable breathing
Mean and median of peak differences	Large differences between breath amplitudes suggest unstable breathing
Mean and median of peak intervals	Times between successive peaks in the window should be similar

These are derived from the waveform of the Respeck signal, the accelerometer signals and the breath timing measurements.

TABLE 2 Windows used in the study

	Sample	
	Used for training	**Used to test**
Total windows N	21 055	2223
No cannula or Respeck rate	603 (3)	29 (1)
No Respeck rate	618 (3)	67 (3)
No cannula rate	4844 (23)	363 (16)
Discarded due to undefined feature	575 (3)	94 (4)
Used for analysis	14 415 (68)	1670 (75)

Data are presented as n (%) unless otherwise stated.

disease (n=6) and other conditions (n=4). Many patients had comorbid conditions: these were cardiovascular in 32, respiratory in 20, diabetes in 11 and renal impairment in 12.

The total number of 20-s windows analysed was 23 278. Table 2 shows how these periods were used in the study.

The breathing rates of the patients, measured from the nasal pressure records, varied considerably. The breathing rate for each patient was summarised as a median duration for all the breaths measured in that patient's record. For all patients, median (interquartile range) breath duration was 2.80 (2.40–3.44) s. These represent breathing rates of 21 (17–25) breaths per min. Figure 1 shows a representative trace from the study. It shows a period of quiet breathing transiently affected by a change most prominent in one axis, which disturbs the Respeck signal, and would affect the matching between the Respeck rate and the nasal cannula rate.

Figure 2 shows the results obtained when the selection process was applied to the test data. Before selection, the median difference between Respeck and cannula rates was 1.25 breaths per min. After selection, the median difference was reduced to 0.46 breaths per min. Although this required removal of a median 58% of the time periods recorded, in all but one of the test subjects, the remaining time available for analysis would have been sufficient to give 1 min of respiratory rate measurements in every 10 min of data.
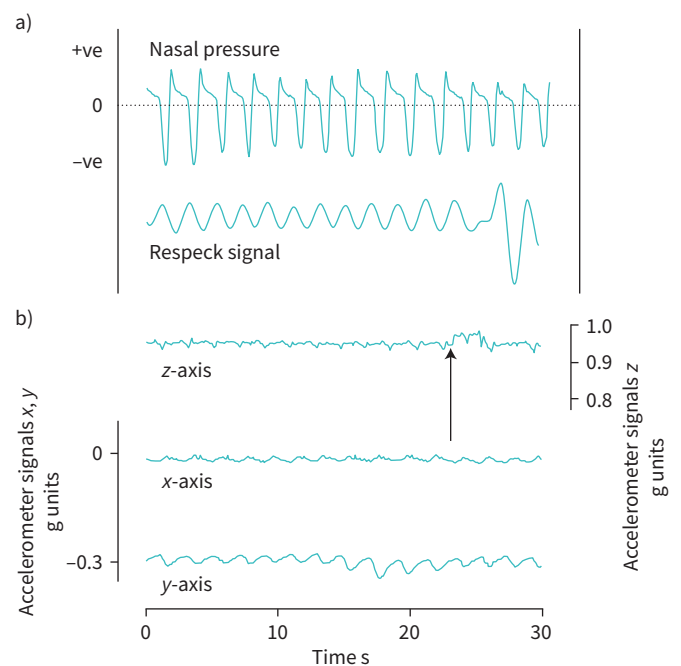

FIGURE 1 An example of a section of recording from a single subject. a) Nasal cannula pressure. A downward deflection (pressure decrease) indicates inspiration. b) Accelerometer forces. These are smoothed and combined into the Respeck signal. The arrow indicates the start of a patient movement, seen in one of the axis measures, that causes a large deflection of the Respeck signal. +ve: positive; −ve: negative.
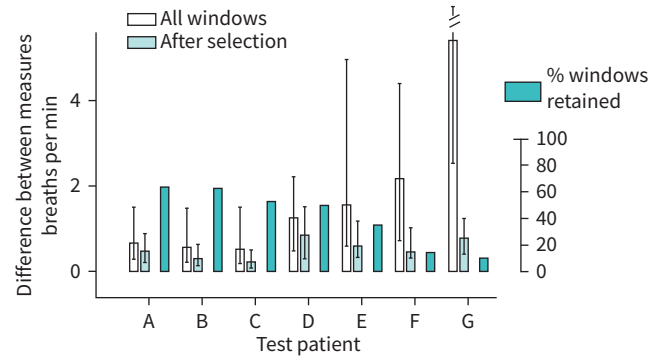
FIGURE 2 Effects of applying selection process to the test data. The six test patients are identified by letters, A–F. The test patients are ordered according to proportion of retained measurement windows. Patient G has a low proportion of windows retained.

Figure 3 shows the change in performance after selection. The limits of agreement ranged from −2.01 to 1.9 breaths per min and the confidence intervals of these limits were small. This contrasts with the greater limits found before signal selection.

Figure 4 compares two plots of rate estimates for patients A and G in the test sample (see figure 2). The record for patient G includes only 10% of acceptable windows but there remain sufficient estimates to record the respiratory rates in this 15-min period.

Discussion

A machine learning-based approach to Respeck data analysis has shown that useful respiratory rate information can be extracted from noisy Respeck signals in a challenging clinical setting. In the test data set, the respiratory frequency matched the reference values closely (median absolute difference of 0.46 breaths per min). The experimental setup to gather time-synchronised Respeck and reference data ensured comparisons over exactly contemporaneous periods. A previous analysis of the nasal cannula signals obtained from this study showed that repeated estimates of respiratory rate, made from the same record but at different times, could vary substantially [6]. The performance after signal selection is superior to results reported in previous studies [11–13, 24, 25].

The study has several strengths: a medically appropriate cohort of patients, studied in clinical circumstances, using an appropriate and accurate comparison signal. The respiratory signal was evaluated before the signal analysis was undertaken and with the other signal records concealed. The nasal cannula signal allows the exact timing of the breathing cycle to be identified easily and precisely, making it an attractive choice as a "gold standard" in comparison to other devices, which give a less faithful measure of respiration.
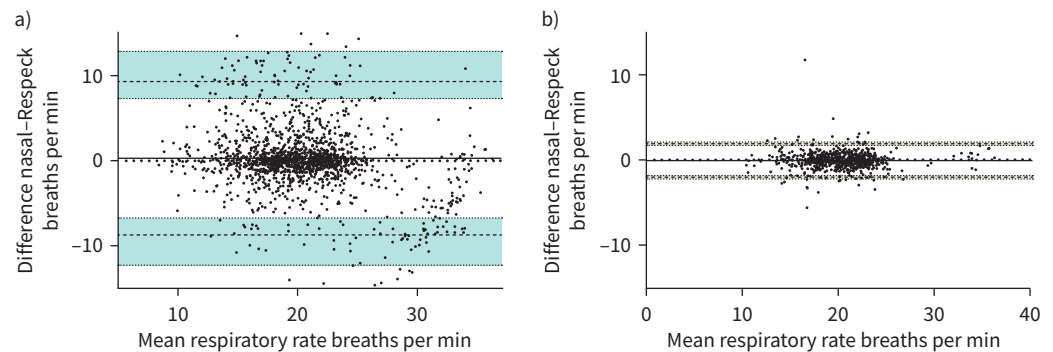

FIGURE 3 Bland and Altman plots of nasal−Respeck differences a) before and b) after selection of suitable signals, plotted in relation to the mean respiratory rate $\left(\dfrac{\text{nasal} + \text{Respeck}}{2}\right)$. The data are from all six validation subjects: each dot indicates one pair of compared values, nasal and Respeck measures, for a single observation window in one validation subject's data. Plots show bias (solid line) and limits of agreement (outer dashed lines) surrounded by their 95% confidence ranges (shaded). To allow comparable scales to be used in both panels, 31 of the 1670 data points from the "before selection" panel are not plotted.
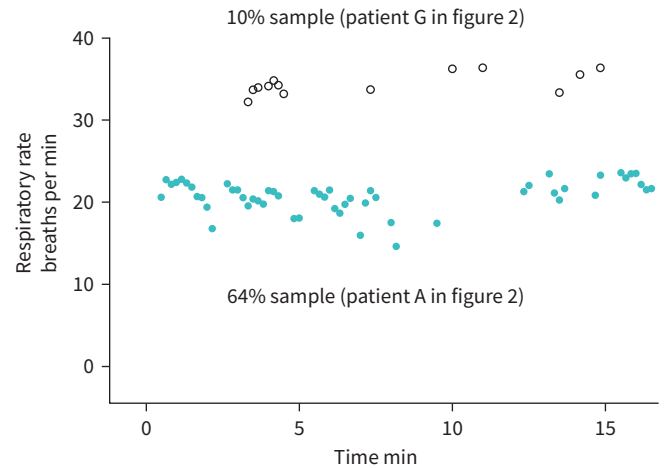
FIGURE 4 Plots of rate estimates from two representative patients shown for a 17-min time period (patient records are from A and G in the test sample, see figure 2). These are patients with extreme proportions of measurement windows that were chosen as accurate, from the whole window dataset for that patient. Even with a 10% acceptance rate, in patient G, these samples are sufficient to give a reasonable measure of the overall respiratory rate over the time period shown.

We had some technical difficulties in this study. We had to set up a new system for recording nasal pressure and this caused some early failures to record. Some of the nasal pressure records were of poor quality, probably caused by fluid within the cannula tubing. This could not be detected because the signal could not be checked as the recording was being made. Some records may have been degraded by predominant mouth breathing, and speech, coughing, cannula displacement or disconnection to allow ambulation can also interrupt nasal pressure measurements.

Signal analysis for model development was conducted without reference to the nasal pressure waveform, although onset times of breaths were used to measure the breath periods. We noted substantial heterogeneity in the breathing of the patients whose records were used to provide the training dataset. We considered that this would be helpful in training the classification process because the model incorporates these variations when considering the test data and future subjects.

The final model performed well on the test data. To allow comparison with other studies, we present this in a Bland and Altman analysis (figure 3). This showed a median rate difference of only 0.84 breaths per min (95% limits of agreement −2.01–1.90 breaths per min) and in 90% of the patients, the rate differences were <1.27 breaths per min. This degree of precision should be clinically acceptable for a scoring system. Very few reports assessing respiratory rate measuring devices define acceptable agreement explicitly but one study suggested that a difference <2 breaths per min indicates equivalence [26].

Previous work has been done to recognise respiratory signals that are "reliable". For example, some cardiac monitors measure respiration from an impedance signal measured using ECG chest leads, and body motion frequently makes this signal unreliable. A computer system of decision rules was developed to score reliability of this signal using a database of recordings from trauma victims [27]. Subsequent reanalysis of data using the selected signals improved prediction of outcome events [28]. Our work uses the same approach but has the strength of incorporating a concurrent reference measure, rather than depending on subjective identification of "trusted" sections of the original record.

Although vital signs monitoring in general wards appears to improve early detection of illness [29, 30], monitoring using continuous acquisition of data is not clearly more effective than when using intermittent chart-based systems [31]. Detecting deterioration may be more effective when the time course of values is used, rather than using threshold values [32, 33]. If so, regular and accurate measurement of respiratory rate may be particularly useful. More precise measurements may be more effective for scoring systems with narrow cutoff ranges [20–25].

Ideally, an automatic summary measure at 10-min intervals could display features such as a progressive increase in rate in a patient with sepsis or a decrease associated with excess opioid. Machine learning allowed us to select trustworthy signals from our monitor to do this, giving results that differ by <2 breaths per min from a reference device. This promising approach would allow automatic charting and provide input to illness scoring systems.

References

1   National Institute for Clinical Excellence. Acutely ill patients in hospital. Recognition of and response to acute illness in adults in hospital. 2007. www.nice.org.uk/guidance/CG50. Date last checked but not updated: January 2020. Date last accessed: 7 September 2020.

2   Huang C, Wang Y, Li X, et al. Clinical features of patients infected with 2019 novel coronavirus in Wuhan, China. Lancet 2020; 395: 497–506.

3   Cecconi M, Piovani D, Brunetta E, et al. Early predictors of clinical deterioration in a cohort of 239 patients hospitalized for Covid-19 infection in Lombardy, Italy. J Clin Med 2020; 9: 1548.

4   Burian E, Jungmann F, Kaissis GA, et al. Intensive care risk estimation in COVID-19 pneumonia based on clinical and imaging parameters: experiences from the Munich cohort. J Clin Med 2020; 9: 1514.

5   Ming DK, Sorawat S, Chanh HQ, et al. Continuous physiological monitoring using wearable technology to inform individual management of infectious diseases, public health and outbreak responses. Int J Infect Dis 2020; 96: 648–654.

6   Drummond GB, Fischer D, Arvind D. Current clinical methods of measurement of respiratory rate give imprecise values. ERJ Open Res 2020; 6: 00023-2020.

7   Lim WS, Carty SM, Macfarlane JT, et al. Respiratory rate measurement in adults – how reliable is it? Respir Med 2002; 96: 31–33.

8   Badawy J, Nguyen OK, Clark C, et al. Is everyone really breathing 20 times a minute? Assessing epidemiology and variation in recorded respiratory rate in hospitalised adults. BMJ Qual Saf 2017; 26: 832–836.

9   Bianchi W, Dugas AF, Hsieh YH, et al. Revitalizing a vital sign: Improving detection of tachypnea at primary triage. Ann Emerg Med 61: 37–43.

10   van Loon K, Breteler MJM, van Wolfwinkel L, et al. Wireless non-invasive continuous respiratory monitoring with FMCW radar: a clinical validation study. J Clin Monit Comput 2016; 30.

11   Hernandez-Silveira M, Ahmed K, Ang SS, et al. Assessment of the feasibility of an ultra-low power, wireless digital patch for the continuous ambulatory monitoring of vital signs. BMJ Open 2015; 5: e006606.

12   van Loon K, Peelen LM, van de Vlasakker EC, et al. Accuracy of remote continuous respiratory rate monitoring technologies intended for low care clinical settings: a prospective observational study. Can J Anaesth 2018; 65: 1324–1332.

13   Breteler MJM, KleinJan EJ, Dohmen DAJ, et al. Vital signs monitoring with wearable sensors in high-risk surgical patients: A clinical validation study. Anesthesiology 2020; 132: 424–439.

14   Buist M, Bernard S, Nguyen T V, et al. Association between clinically abnormal observations and subsequent in-hospital mortality: a prospective study. Resuscitation 2004; 62: 137–141.

15   McGaughey J, Alderdice F, Fowler R, et al. Outreach and Early Warning Systems (EWS) for the prevention of intensive care admission and death of critically ill adult patients on general hospital wards. Cochrane Database Syst Rev 2007; 3: CD005529.

16   Odell M, Victor C, Oliver D. Nurses' role in detecting deterioration in ward patients: systematic literature review. J Adv Nurs 2009; 65: 1992–2006.

17   Tirkkonen J, Ylä-Mattila J, Olkkola KT, et al. Factors associated with delayed activation of medical emergency team and excess mortality: an Utstein-style analysis. Resuscitation 2013; 84: 173–178.

18   Drummond GB, Bates A, Mann J, et al. Validation of a new non-invasive automatic monitor of respiratory rate for postoperative subjects. Br J Anaesth 2011; 107: 462–469.

19   Thurnheer R, Xie X, Bloch KE. Accuracy of nasal cannula pressure recordings for assessment of ventilation during sleep. Am J Respir Crit Care Med 2001; 164: 1914–1919.

20   Bishop CM. Pattern recognition and machine learning. New York, Springer, 2006.

21   Pedregosa F, Varoquaux G, Gramfort A, et al. Scikit-learn: machine learning in Python. J Mach Learn Res 2011; 12: 2825–2830.

22   sklearn.neural_network.MLPClassifier. http://scikit-learn.org/stable/modules/generated/sklearn.neural_network. MLPClassifier.html#sklearn.neural_network.MLPClassifier. No update: version scikit-learn 0.23.2. Date last accessed: 7 September 2020.

23   Bland JM, Altman DG. Agreement between methods of measurement with multiple observations per individual. J Biopharm Stat 2007; 17: 571–582.

24   Granholm A, Pedersen NE, Lippert A, et al. Respiratory rates measured by a standardised clinical approach, ward staff, and a wireless device. Acta Anaesthesiol Scand 2016; 60: 1444–1452.

25   Breteler MJM, Huizinga E, Loon KV, et al. Reliability of wireless monitoring using a wearable patch sensor in high-risk surgical patients at a step-down unit in the Netherlands: a clinical validation study. BMJ Open 2018; 8: e020162.

26   Smith I, MacKay J, Fahrid N, et al. Respiratory rate measurement: a comparison of methods. Br J Healthc Assist 2011; 5: 18–23.

27   Chen L, McKenna T, Reisner A, *et al.* Algorithms to qualify respiratory data collected during the transport of trauma patients. *Physiol Meas* 2006; 27: 797–816.

28   Chen L, Reisner AT, Gribok A, *et al.* Can we improve the clinical utility of respiratory rate as a monitored vital sign? *Shock* 2009; 31: 574–580.

29   Smith GB, Prytherch DR, Schmidt P, *et al.* Hospital-wide physiological surveillance – a new approach to the early identification and management of the sick patient. *Resuscitation* 2006; 71: 19–28.

30   Schmidt PE, Meredith P, Prytherch DR, *et al.* Impact of introducing an electronic physiological surveillance system on hospital mortality. *BMJ Qual Saf* 2015; 24: 10–20.

31   Cardona-Morrell M, Prgomet M, Turner RM, *et al.* Effectiveness of continuous or intermittent vital signs monitoring in preventing adverse events on general wards: a systematic review and meta-analysis. *Int J Clin Pract* 2016; 70: 806–824.

32   Zimlichman E, Szyper-Kravitz M, Shinar Z, *et al.* Early recognition of acutely deteriorating patients in non-intensive care units: assessment of an innovative monitoring technology. *J Hosp Med* 2012; 7: 628–633.

33   Kellett J, Wang F, Woodworth S, *et al.* Changes and their prognostic implications in the abbreviated VitalPAC Early Warning Score (ViEWS) after admission to hospital of 18,827 surgical patients. *Resuscitation* 2013; 84: 471–476.