



US 20170071484A1

(19) **United States**(12) **Patent Application Publication**
Strachan et al.(10) **Pub. No.: US 2017/0071484 A1**(43) **Pub. Date: Mar. 16, 2017**(54) **METHOD AND APPARATUS FOR
MONITORING PATIENT STATUS**(71) Applicant: **OBS MEDICAL LIMITED**,
Oxfordshire (GB)(72) Inventors: **Iain Guy David Strachan**, Oxfordshire
(GB); **Hamid Reza Mohseni**, Oxford
(GB)(21) Appl. No.: **15/310,652**(22) PCT Filed: **Mar. 31, 2015**(86) PCT No.: **PCT/GB2015/050993**

§ 371 (c)(1),

(2) Date: **Nov. 11, 2016**(30) **Foreign Application Priority Data**

May 13, 2014 (GB) 1408469.3

Publication Classification(51) **Int. Cl.****A61B 5/0205** (2006.01)**G06F 19/00** (2006.01)**A61B 5/00** (2006.01)(52) **U.S. Cl.**CPC **A61B 5/02055** (2013.01); **A61B 5/7264**
(2013.01); **A61B 5/7282** (2013.01); **G06F**
19/3431 (2013.01); **G06F 19/3437** (2013.01);
A61B 5/024 (2013.01)

(57)

ABSTRACT

A method and apparatus for monitoring the status of a patient on the basis of a plurality of vital signs measurements such as heart rate, breathing rate, blood pressure, temperature, oxygen saturation etc. At each time point a set of such vital signs measurements is compared to two statistical models, one being a population-based statistical model representing the probability distribution of measurements from a large number of individuals, and the second is a patient-specific statistical model based on measurements from the particular patient being monitored. The patient-based statistical model is initialized as identical to the population-based model, but then is updated in a Bayesian update technique so that it evolves towards the patient's individual state. First and second novelty indexes are calculated by comparing each status measurement to each of the two statistical models. A threshold on the novelty index from the population-based model can generate an alarm if the index indicates that the patient's status is outside a predefined normal region for the population. A threshold on the numerical index from the patient-specific model can indicate a change in status of the patient, which can be an improvement or a deterioration.




Fig. 1




Fig. 2




Fig. 3(a)




Fig. 3(b)




Fig. 3(c)




Fig. 3(d)




Fig. 3(e)




Fig. 3(f)




Fig. 4




Fig. 5




Fig. 6




Fig. 7(a)




Fig. 7(b)




Fig. 7(c)




Fig. 7(d)




Fig. 7(e)




Fig. 7(f)




Fig. 8




Fig. 9





Fig. 10






Fig. 11

Fig. 12





Fig. 13



METHOD AND APPARATUS FOR MONITORING PATIENT STATUS

[0001] The present invention relates to a method and apparatus for monitoring the status of a patient as measured by a plurality of vital signs.

[0002] It is common in a clinical setting, particularly a high dependency care unit or intensive care unit, to monitor continuously the condition or status of patients. Typically vital signs such as one or more of blood pressure, heart rate, skin or body temperature, blood oxygen saturation (for instance measured by pulse oximetry with a finger probe), breathing rate (for instance measured by electrical impedance pneumography) may be monitored. The heart rate, and sometimes breathing rate, may be obtained from one or more channels of an electrocardiogram (ECG), which may also be monitored continuously. It is also possible to derive from these primary signals some secondary parameters such as heart rate variability, S-T segment elevation/depression and so on. Normally the individual measurements are displayed to a clinician in a form of a number, and possibly a graphical trace, as illustrated, for example, in FIG. 11. However over recent years techniques for combining these individual parameters into a score which gives some overall indication of the patient's condition or status have been in use. To date, though, the majority of such scoring regimes in patient monitoring systems tend to be "one size fits all", in that the same scoring rules are applied irrespective of the subject. This is true of the modified early warning score (MEWS system), explained, for example in "The value of modified early warning score (MEWS) in surgical in-patients: a prospective observational study" by J Gardner-Thorpe, N Love, J Wrightson, S Walsh and N Keeling, (Ann. R. Coll. Surg. Engl. October 2006; 88(6): PP571-575). It is also true of the present applicant's own Visensia system explained in U.S. Pat. No. 7,031,857 B2 which uses an FDA-cleared population-based statistical model of patient condition against which current measurements are compared.

[0003] However, normal values for vital signs can vary considerably from one patient to another due to factors such as age, type of disease, medication, type of surgery, etc. Replacing the population-based statistical model with a personal model for an individual patient is problematic because no valid set of measurements for developing such a model exists for an individual, especially when they have just changed status, for example because of a surgical operation or developing a disease or other medical condition. In the applicant's own U.S. Pat. No. 7,031,857 B2 it was proposed that after an initial period of judging a patient's condition against a population-based statistical model, enough measurements on the individual patient would have been obtained to replace the population based model with a patient-specific one. However, in practice this is not possible as a patient's condition is changing and obtaining regulatory clearance for monitoring on the basis of a not-yet validated set of data is impossible. Therefore it is difficult to provide a scoring system which can provide a reliable and valid indication of variations from normality for an individual patient. Further it would require a long period of monitoring the patient once they have reached a stable state to provide a statistically reliable set of measurements from which a valid model can be generated. At the start of monitoring therefore, when the patient may be most at risk of sudden deterioration, a personalised model is not avail-

able. So it is not clear how to obtain a personalised model which is useful in a reasonably quick time.

[0004] The present invention provides an enhancement to a continuous patient monitoring system such as that disclosed in U.S. Pat. No. 7,031,857 B2 by comparing the status of a patient, as measured by a plurality of vital signs, simultaneously to two different statistical models, one being population-based and one being personal to the patient. The statistical model which is personal to the patient is obtained by applying machine learning techniques to an initial population-based model and updating it based on each current vital signs measurement from the patient as it is received. By simultaneously evaluating the patient's status against two statistical models an alert notification can be generated relative to the population-based model if a computed score representing the patient's condition breaches a predefined threshold (this model can be a regulatory approved model), thus providing safety for the patient if their condition deteriorates, but also a notification of a change in the patient's status compared to the patient-specific statistical model can also be generated. Such a change in status could be an improvement or deterioration for that particular patient, but which is still within an overall region of normality for the population, and it can thus provide an earlier indication of change.

[0005] In more detail, therefore, the present invention provides a method of monitoring the status of a patient comprising the steps of continuously measuring a plurality of patient vital signs to provide a succession of multi-parameter vital signs observations, comparing each successive multi-parameter vital signs observation to two statistical models, a first of the statistical models being a population-based reference model mapping the probability distribution of multi-parameter vital signs observations for a population of different individuals, and a second of the statistical models being a patient-specific model which is a probability distribution based at least partly on previous multi-parameter vital signs observations for the patient being monitored, calculating first and second numerical indices based respectively on the probability densities of the current multi-parameter vital signs observation obtained by comparison with the first and second statistical models, wherein an increased value of said indices represent an increase in abnormality with respect to the statistical model, outputting an alert notification if the first numerical index exceeds a first threshold, and outputting a status change notification if the second numerical index exceeds a second threshold.

[0006] The first threshold on the first numerical index in the population-based reference model effectively defines a global region of normality for the population. The region of parameter space occupied by the patient's current state would normally be expected to be smaller than the global region of normality and may be embedded within it. The second threshold effectively defines a region of normality personal to the individual patient and changes from this could reflect either an improvement or deterioration in the individual patient's condition depending on whether the change is accompanied by an increase or decrease in the first numerical index obtained by reference to the population-based model.

[0007] The vital signs measured may typically comprise heart rate, breathing rate, blood pressure, blood oxygen saturation and body temperature and at each time point the measured values form a multi-parameter vital signs obser-

vation for evaluation against the statistical models. Typically the various parameters are collected at different sampling rates, for example the heart rate may be calculated every few seconds whereas the blood pressure may be measured only once every 2 to 4 hours, and so to provide a continuous set of samples the lower sampling rate measurements are repeated in successive multi-parameter observations until they are re-measured.

[0008] Preferably the patient-specific statistical model is updated at each time step by using the most recent multi-parameter vital signs observation. However the updating step can be selectively enabled when the patient status is within the global normality region for the population, i.e. when the first numerical index does not exceed the first threshold. This can avoid updating the model based on noisy or spurious measurements. The inhibition on updating can be automatic based on the first numerical index, or can be under the control of a clinician. The threshold defining the region of global normality and the threshold for determining whether or not to update the patient specific model with an individual measurement can be different from each other. Thus it is possible for the clinician to selectively allow the patient-specific model to be updated even if the population-based model shows some degree of abnormality. In the case of one or more of the parameters in an observation being spurious or missing it is possible for this to be treated as a missing variable (for example replaced by its mean value) in determining whether to update the patient-specific model, but if it is decided to update then that the complete original observation is used to update the model. The decision on whether to update may be automatic, e.g. based on the degree of abnormality, or under the control of the clinician or operator.

[0009] The step of updating the patient-specific model can be stopped when the model becomes stable (for example when the parameters do not change significantly for a predetermined period of time, e.g. 20 or 30 minutes), and restarted when a comparison of the most recent multi-parameter vital signs observation (or some predefined number of them) with the patient-specific model indicates a change in patient status.

[0010] Preferably the patient-specific model is initialised to be identical to the population model, or to have the same statistical properties (e.g. mean and variance) as the population model such that it would give approximately the same values for the numerical indices. As patient-specific observations are received, this initialized model evolves to cover the area of normality for the particular patient being monitored.

[0011] The patient-specific model and the population-based model may be Kernel Density Estimators, such as a plurality of spherical Gaussian density functions with all kernels having equal weights. The patient's specific model may be updated by a Bayesian process which comprises calculating from the patient-specific model the likelihood of the most recent multi-parameter vital signs observation, multiplying the model's prior probability distribution by the calculated likelihood and then renormalizing it to form the Bayesian posterior probability distribution for the model, given the current sample, which then serves as the prior probability distribution for the next sample.

[0012] This update process for the patient-specific model may utilise particle filtering. This is a standard technique in which "particles" (being samples initially generated from

the population-based model by random sampling of it) have a noise term added to each of them and then the likelihood of each particle plus noise is computed. A plurality of the particles is then resampled with a probability of selection being set proportional to its computed likelihood. The resampled set of particles constitutes the new patient-specific model. As an alternative to the standard particle filtering approach, the new samples can be drawn by sampling from the nearest kernel (to the selected particle) in the population-based model. This has the effect that the patient-specific model is always bounded by the population-based model. If it is desired for certain patients that a patient specific model may lie outside the population-based region of normality, then the standard resampling approach should be used.

[0013] As mentioned above the status change notification generated with reference to the patient-specific model can correspond to an improvement or deterioration in the patient's condition and it thus provides extra information to clinical staff compared to the results from the population-based reference model. Preferably an alert based on the patient-specific model is only generated if the current observation or observations are associated with an increase in the first numerical index generated with reference to the population-based reference model. This means that an alert is only generated if it reflects deterioration in the patient's condition.

[0014] Preferably such an increase of the first numerical index is judged by comparing the value of the first numerical index obtained for the current multi-parameter vital signs observation to the value of the first numerical index that was obtained at some previous time, for example at the time when the patient-specific model became stable or when the monitoring process was started.

[0015] To avoid generating alerts or status change notifications when spurious or erroneous measurements are received (for example because of noisy signals), the notifications may only be generated if a threshold is exceeded by a filtered version of the index (e.g. for more than a predefined number of observations in succession or more than a predefined number of observations within a predetermined time period, though other filtered versions may be used). This effectively smoothes the output of the system.

[0016] The invention may be embodied in an apparatus for executing the method and such an apparatus may therefore comprise an input for receiving the patient's vital signs measurements, a memory for storing the statistical models, a processor programmed to execute the steps of comparing the measurement to the statistical models, calculating the numerical indices and comparing them to the thresholds, and to output the alert and status change notifications. The output may comprise a display for displaying the notifications in a graphical form, though audible notifications may also be generated. Typically the display can display the individual measured vital signs in a conventional manner together with the numerical indices and alert and status notifications.

[0017] The apparatus may be incorporated into a standard vital signs monitor or can be embodied in a programmed computer system which receives the vital signs measurements.

[0018] The invention will be further described by way of examples with reference the accompanying drawings in which:—

[0019] FIG. 1 conceptually illustrates regions of normality in statistical models for a population-based model and a patient-specific model;

[0020] FIG. 2 schematically illustrates the operation of a system according to an embodiment of the present invention;

[0021] FIG. 3(a) shows a time series of the population-based numerical index and FIGS. 3 (b) to (f) show a time series of five vital signs for a patient;

[0022] FIG. 4 illustrates the first two principal components of the global population model in one embodiment of the invention;

[0023] FIG. 5 illustrates the first two principal components of the centre of kernels for the population model together with the centre of Gaussian kernels for the patient-specific model in one embodiment of the invention;

[0024] FIG. 6 illustrates the numerical indices calculated from population-based, patient-specific and combined models of FIG. 5;

[0025] FIG. 7(a) to (f) show the population-based index (a) and the five vital signs (b) to (f) for a patient with a large artefactual observation;

[0026] FIG. 8 illustrates the corresponding first two principal components for the measurements of FIG. 7;

[0027] FIG. 9 illustrates the first two components for the measurements of FIG. 7 against the patient-specific model of FIG. 5;

[0028] FIG. 10 schematically illustrates a system for executing the method in one embodiment of the invention;

[0029] FIG. 11 illustrates a prior art combined vital signs display;

[0030] FIG. 12 is a flow diagram of the method steps for alert generation based on the population model in one embodiment of the invention; and

[0031] FIG. 13 is a flow diagram of the method steps for alert generation based on the patient-specific model in one embodiment of the invention.

[0032] FIG. 1 illustrates the concept of the operation of the monitoring system of the present invention once a patient-specific model has been learnt. The large ellipse represents schematically a multi-dimensional population model (for example the FDA-cleared Visensia model), represented as a two-dimensional schematic (in practice there are as many dimensions as measured parameters). By comparing a current observation to the population model as described in U.S. Pat. No. 7,031,857 B2 it can be judged whether the current patient state is inside or outside the boundary of the ellipse which represents normality. If the patient's status is outside then the model threshold is exceeded and if this persists for a certain number of observations or a certain period, an alert is generated calling for intervention by a member of a medical emergency team.

[0033] The small ellipse within the larger ellipse in FIG. 1 represents a patient-specific model that has been learnt during an initial phase of operation of the system. As will be explained in more detail below, on initiation of the system the patient-specific model has the same statistical properties as the global population-based model, but as data is acquired in real time, the updating of the patient's specific model causes its region of normality to shrink in size and define a region of normality specific to that patient. The smaller ellipse in FIG. 1 represents that region of normality for the specific patient. After a fixed period of learning, for example 20 minutes, or after the region of normality defined by the

patient-specific index has become stable, the patient-specific model can be frozen and then status change notifications generated by comparison of new vital signs measurements with the patient-specific model.

[0034] As illustrated in FIG. 2, therefore, in this embodiment of the invention, a current observation (preferably having been pre-processed so that the parameters are normalised and decorrelated, e.g. by using a standard whitening process) is compared to the population-based model 1 and a general alert 3 can be generated if the predefined normality threshold (represented by the large ellipse in FIG. 1) is exceeded. Thus a general alert derived from the population-based model can be generated at any time even immediately on starting monitoring. Once the patient-specific model has been learnt in the initial phase, a static snapshot is taken of the patient-specific model, and incoming observations (again preferably pre-processed as mentioned above) are compared to the static snapshot of the patient-specific model in order to judge whether the observation represents normality for the specific patient (represented by the small ellipse in FIG. 1) or a change of status. In the event of a status change, i.e. the reading lying outside a region of normality for that patient, a status change notification 9 can be generated, which may be according to Type A or Type B events as illustrated in FIG. 1. The current observation can also be used to update the dynamic patient-specific model 5, assuming that updating is enabled (such enablement can be with reference to the population-based model as discussed in more detail below).

Examples of Updating Patient-Specific Model

[0035] In this embodiment the population model (large ellipse in FIG. 1) is estimated by a kernel density estimator from a large and artefact-free data set (as previously disclosed in European Patent EP 1 389 948 B1). The patient-specific model can also be estimated by a kernel density estimator and is updated by machine-learning techniques in an online fashion as explained below. Thus the population-based and patient-specific statistical models are probability density functions which give a probability density p for each combination of measured vital signs y (HR, BP, BR, SpO₂, Temp). The probability density functions in this embodiment are the sum of a plurality of Gaussian functions known as kernels, in a multi-dimensional parameter space (i.e. each point in that space represents a particular set of normalised vital signs measurements).

[0036] In one embodiment of the invention, the approach for updating the patient-specific model of vital signs is to employ particle filtering (sequential Monte Carlo technique), though by imposing some assumptions other machine-learning techniques such as Kalman Filtering or Unscented Kalman Filtering can be used (see: Beyond the Kalman Filter: Particle Filters for Tracking Applications, Branko Ristic, Artech House Publishers, 2004, ISBN 158053631X, or The unscented Kalman filter for nonlinear estimation, Wan E. A, and Van der Merwe, R, Adaptive Systems for Signal Processing, Communications, and Control Symposium 2000. AS-SPCC. The IEEE 2000).

[0037] A particle filter is a non-Gaussian density estimator, which is formulated in the Bayesian framework. Bayesian learning estimates the distribution of interest (posterior distribution), using the prior distribution and the likelihood of the current observation (more details on dynamic Bayesian updating can be found in, for example, Arulampalam, M.

S.; Maskell, S.; Gordon, N.; Clapp, T., "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," IEEE Transactions on Signal Processing, vol. 50, no. 2, pp. 174, 188, February 2002).

[0038] As mentioned above, the population distribution is used as the prior distribution for estimation of the posterior distribution for a specific patient. This allows measurement artefacts to be handled better and reduces the updating rate when the vital signs are well-away from the normal values. It also assists a better decision about the improvement or deterioration of the patient, for example, when the vital signs are moving away from the centre of normality of the patient model, but they are moving towards the centre of the normality of the population model as illustrated by arrow B in FIG. 1.

[0039] Thus starting from the population distribution, to obtain the patient-specific distribution, suppose that the population distribution, i.e., the prior, is represented using a kernel density estimator such as Parzen window [see, e.g., U.S. Pat. No. 7,031,857B2]. Let N_p be the number of kernels, e.g. 400 in this embodiment, each of which is a Gaussian function

$$\mathcal{G}(x; \mu^j, C^j) = \frac{1}{\sqrt{(2\pi)^d |C^j|}} e^{-\frac{1}{2}(x-\mu^j)^T C^{j-1} (x-\mu^j)}$$

with mean μ^j , $j=1, \dots, N_p$ and covariance C^{N_p} , respectively. Here d is the dimensionality of data. It is also possible to have different covariance matrices for each kernel, as in the Gaussian mixture models (GMM), but here for the sake of simplicity of notation, we assume that the covariance matrices are the same for all kernels. The population distribution that the process starts from is therefore given by:

$$p(x_i; \mathcal{P}) = \frac{1}{N_p} \sum_{j=1}^{N_p} \mathcal{G}(x_i; \mu^j, C^j) \quad (1)$$

[0040] So the probability density p for a particular set of values x_t is the sum of the value of each of the Gaussian kernel. Furthermore, suppose that the vital signs and their noisy observations at time t are denoted by x_t and y_t , respectively, we have:

$$y_t = x_t + v_t \quad (2)$$

where v_t is the additive and zero-mean noise. Based on the Bayesian framework, we are interested in the estimation of the posterior distribution of the vital signs denoted by $p(x_t | y_t; \mathcal{P})$ i.e. the probability of state x_t given measurement y_t for patient $\mathcal{P}$. Suppose also that x_t and y_t obey the following Markovian state space:

$$x_t \sim p(x_t | x_{t-1}; \mathcal{P}) \quad (3)$$

$$y_t \sim p(y_t | x_t; \mathcal{P}) \quad (4)$$

i.e. that the values of x_t are distributed according to $p(x_t | x_{t-1}; \mathcal{P})$ and the observations y_t are distributed according to $p(y_t | x_t; \mathcal{P})$. The posterior distribution, i.e. the patient-specific distribution, is assumed to be non-Gaussian, and thus N

$\mathcal{P}$ (e.g. 1000 in this embodiment) Gaussian kernels, each with potentially a different weight w_t^i are used to approximate this distribution:

$$p(x_t | y_t; \mathcal{P}) = \sum_{i=1}^{N_p} w_t^i \mathcal{G}(x_t; x_t^i, C^{\mathcal{P}}) \quad (5)$$

where w_t^i , $i=1, \dots, N_p$ are kernel weights and $\mathcal{G}(x_t; x_t^i, C^{\mathcal{P}})$ is the Gaussian function centered at x_t^i with covariance matrix $C^{\mathcal{P}}$. Thus, instead of the Dirac delta function as in standard particle filter which is known as selective importance resampling (SIR), we use Gaussian kernels for approximation of posterior distribution, and therefore the regularised particle filter is employed. The regularised particle filter is the same as the SIR filter except it has a modified resampling technique to take into account the kernel used in the approximation of the posterior density [for more details see: Arulampalam, M. S.; Maskell, S.; Gordon, N.; Clapp, T., "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," IEEE Transactions on Signal Processing, vol. 50, no. 2, pp. 174, 188, February 2002].

[0041] Based on the above definitions and assumptions, an algorithm using a generic particle filter may be given by Algorithm 1.

Algorithm 1

FOR $i = 1:N_p$:

 Sample X_0^i from equation (1) and set the weight of each $w_0^i = \frac{1}{N_p}$

END FOR [this gives us an initial model of posterior distribution consisting of, e.g. 1000, Gaussian kernels each centered at a particular value of x_0^i and each of equal weight]

FOR $t = 1:T$ [at each receipt of a new observation]

 FOR $i = 1:N_p$

 sample from importance density $x_t^i \sim \pi(x_t | x_{t-1}^i, y_t)$

 calculate a new weight for each of the particle:

$$w_t^i \propto w_{t-1}^i \frac{p(y_t | x_t^i) p(x_t^i | w_{t-1}^i)}{\pi(x_t^i | w_{t-1}^i, y_t)}$$

 [see discussion below for how to calculate these terms]

 END FOR

 normalize the weights $w_t^i = \frac{w_t^i}{\sum w_t^i}$ to have unit sum

 calculate $N_{eff} = \frac{1}{\sum w_t^i{}^2}$ number of effective samples as a measure

 for snapshots [It can happen that some of the weights are so low that the number of kernels with a sufficient weight to contribute to the model is too low. This resampling has the effect of introducing more relevant kernels into the patient-specific model and removing outlying kernels which have a very low weight.]

 IF $N_{eff} < N_T$

 resample $\{x_t^i, w_t^i\}$ such that $w_t^i = \frac{1}{N_p}$

 ENDIF

ENDFOR

[0042] In this algorithm, $\pi(x_t | x_{t-1}^i, y_t)$ is the importance density and the measurement transition probability is also given by:

$$p(y_t | x_t^i; \mathcal{P}) \approx \mathcal{G}(y_t | 0, C_v) \quad (6)$$

where $\mathcal{G}(y_t | 0, C_v)$ is the distribution of the additive noise v_t , which has been assumed to be distributed independently and identically. In this embodiment, the importance density is chosen to be equal to the prior density $\pi(x_t^i | x_{t-1}^i, y_t) \approx p(x_t^i | x_{t-1}^i)$, and therefore the update procedure for weights is simplified to (as in SIR filter):

$$w_t^i \propto w_{t-1}^i p(y_t | x_t^i) \quad (7)$$

[0043] Optionally, the state transition probabilities can be approximated by the population model:

$$p(x_t^i|x_{t-1}^i; \mathbb{P}) \approx p(x_t|x_{t-1}; \mathcal{P}) \quad (8)$$

[0044] Here, we assume that the transition density is the same as the population density model given in equation (1). We may further approximate the state transition probability by:

$$p(x_t|x_{t-1}^i; \mathcal{P}) \approx \mathcal{G}(x_t; \mu_{\mathcal{P}}^{i,inst}, C_{\mathcal{P}}) \quad (6)$$

where $\mu_{\mathcal{P}}^{i,inst} = \arg\max_{\mu^j} \mathcal{G}(x_{t-1}^i; \mu^j, C_{\mathcal{P}})$, for $j=1, \dots, N_{\mathcal{P}}$. That is to say, to sample x_t^i from the population model, we sample from the kernel in the population model that is nearest to the particle x_{t-1}^i , or we consider the nearest kernel as the transition distribution. Equations (8) and (9) have the desirable effect that the patient-specific model's region of normality is always bounded by the population-based model. If a clinician decides, for a particular patient, that their region of normality lies outside the population-based region of normality, then this further approximation can be omitted.

[0045] Note that the measurement equation (2) is linear, but the transition of the state is non-Gaussian (nonlinear), since it is based on the population distribution, which is a non-Gaussian density. However, for a faster implementation, one may approximate the state transition by a Gaussian distribution and effectively employ the Kalman filter, which assumes Gaussian distribution for both initial and transition distributions. Based on the above assumptions the implementation the generic algorithm given in Algorithm 1 is converted to Algorithm 2.

Algorithm 2

- 0) To get the initial patient-specific model randomly sample the population model to get $N_{\mathcal{P}}$ particles (e.g. $N_{\mathcal{P}} = 1000$) each having an associated weight $w_0^i = \frac{1}{N_{\mathcal{P}}}$ centered at x_0^i in parameter space, a "model" as defined by equation (5).
 - 1) At time point t receive an observation y_t .
 - 2) Construct the likelihood according to the equation (6) and update the weights in the patient specific model according to (7) (note the particle position in parameter space x_0^i does not change, just the weight w_t^i).
 - 3) Normalize the weights $w_t^i = \frac{w_t^i}{\sum w_t^i}$.
 - 4) Calculate number of inefficient samples and if they are less than a predefined threshold go to next step, otherwise go to step 1.
 - 5) If it is desired for the patient-specific model to remain bounded by the population-based model, then resample from the population model (equation 8) i.e. for each particle find the nearest kernel in the population model and take a random sample from that kernel to be the new particle (equation 9). This changes the position of the particles (centre of the kernels), while the weights remain the same. This resampling also is independent from the observations.
 - 6) Otherwise (ie if the patient-specific model may go outside the bounds of the population-based model) then use a standard resampling technique to have an un-weighted approximate density by eliminating particles having low importance weights and by replicating particles having high importance weights. This resampling makes all particles have the same weights.
 - 7) The particles with new positions and new equal weights approximate the snapshot patient specific model.
-

[0046] As in the standard particle filter, the algorithm has some free parameters including the covariance of the process noise, the covariance of the measurement noise (if we assume they are Gaussian) and the number of effective samples. These parameters can be set empirically, according

to the available data set in hand. A simple approach is to set the process noise covariance equal to the covariance of the kernel in the patient-specific model, and to set the measurement noise covariance equal to the covariance of the kernel of the population model.

First Numerical Index

[0047] A numerical index such as novelty score, e.g. the Visensia index (VSI) used by the applicant, is a measure for the wellness of the patient, and in U.S. Pat. No. 7,031,857 B2 it was defined as a VSI score using equation (1) by:

$$VS \mathbf{I}_{\mathcal{F}}(x_t) = -\log p(x_t; \mathcal{F}) + \log p(\mu_{t,\mathcal{F}}; \mathcal{F}) \quad (7)$$

where $\mu_{t,\mathcal{F}}$ is the mean of $p(x_t; \mathcal{F})$. The term $\log p(\mu_{t,\mathcal{F}}; \mathcal{F})$ is added to make $VS \mathbf{I}_{\mathcal{F}}(x_t)$ zero at the centre of the distribution. This same index is used in this embodiment as the first numerical index which is based on the population-based model. A threshold on this index defines the boundaries of the region regarded as normal and illustrated as the large ellipse in FIG. 1. The population model is fixed so the mean, the second term in equation (10) does not change and can be pre-calculated from the model. As illustrated in FIG. 12 obtaining the first numerical index simply involves taking the incoming multi-parameter observation y_t (HR, BP, BR, SpO₂, Temp) at step 120, preprocessing it as mentioned above at step 122, then for this purpose the state x_t is considered to be the same as the observation y_t as indicated at step 124. This is then compared in step 125 to the population model of equation (1) to read off its probability density $p(x_t|y_t)$ and this is inserted this into equation (10) to calculate $VS \mathbf{I}_{\mathcal{F}}$ in step 126. In step 128 an alert can be generated if $VS \mathbf{I}_{\mathcal{F}}$ or a filtered version of it is outside a threshold, i.e. outside the large ellipse of FIG. 1.

Second Numerical Index

[0048] To obtain a second numerical index, based on the patient specific model, two approaches, as explained below, can be used to calculate the novelty score and therefore generate relevant alerts.

Approach 1

[0049] Similar to equation (10), a novelty (VSI) score for the patient-specific model can be defined (which sets the boundaries of the small ellipse in FIG. 1):

$$VS \mathbf{I}_{\mathcal{P}}(x_t) = -\log p(x_t|y_t; \mathcal{P}) + \log p(\mu_{t,\mathcal{P}}; \mathcal{P}) \quad (8)$$

where $\mu_{t,\mathcal{P}}$ is the mean of $p(x_t|y_t; \mathcal{P})$. This definition sets the boundaries of the small ellipse in FIG. 1.

[0050] Thus to obtain the second numerical index, as illustrated in FIG. 13, the incoming multi-parameter observation y_t (HR, BP, BR, SpO₂, Temp) received at step 130, is pre-processed in step 132, and this is compared in step 134 to the current snapshot of the patient-specific model of equation (5) to obtain its probability density $p(x_t|y_t; \mathcal{P})$ inserting this into equation (11) in step 136. Again the second term based on the mean of the model is fixed (once the patient-specific model is frozen its mean would be recalculated when it is updated) and can be pre-calculated.

[0051] If the value of personalised VSI _{$\mathcal{P}$} exceeds a threshold this is used to operate a status change notification in step 140 which can be displayed or transmitted to a clinician, who then determines whether the patient has deteriorated or improved. In the case of deterioration, as shown in arrow A in FIG. 1, this would provide an earlier

warning than the standard population model as it will be clear from FIG. 1 that a change as shown by arrow A will operate a status change notification before an alarm is operated on crossing the boundary of the large ellipse. In the case of improvement, as shown by arrow B in FIG. 1, the system can permit a new stage of machine learning, to track the new and improved state.

[0052] With an A type notification, patient deterioration is indicated, as the global VSI score will have increased with respect to the mean VSI score for the learned model. This type of notification will likely occur earlier than an eventual Visensia alert, because the state has not yet exceeded first (population) threshold—the outer ellipse in FIG. 1;

[0053] With a B type notification, the global VSI score will have decreased—indicative of an improvement of the patient's condition.

[0054] Steps 138 and 139 illustrate automatic use of the global VS $I_{\mathcal{F}}$, to determine whether an alert should be generated based on the learned (patient-specific) VS $I_{\mathcal{P}}$ score going beyond its threshold.

[0055] It is envisaged that either type of notification would summon a clinician to make an assessment of what has changed. In the event of a type A alert, this could warrant escalation of the intervention, and will have therefore provided an earlier warning than the original system. In the event of a B type alert, the option exists for the clinician to authorise the system to re-engage the learning process, in order to track the patient's improvement.

[0056] In another scenario the patient state might change along a trajectory that leads to only small changes (+ or -) in VSI. In this case, the system can also be allowed to learn the patient's new state.

Approach 2

[0057] As a more sophisticated alternative, the VSI for personalised monitoring can consider both population and patient models, which leads to the second approach. By assuming that these distributions are independent, we may write:

$$p(x_i; \mathcal{P}, \mathcal{P}) \propto p(x_i; y_i; \mathcal{P}) p(x_i; \mathcal{P}) \quad (9)$$

[0058] This leads to the definition of the personalised VSI for a specific patient $\mathcal{P}$:

$$VS_{\mathcal{F}}(x_i) \propto -\log p(x_i; \mathcal{P}, \mathcal{P}) + \log p(\mu_{t, \mathcal{P}\mathcal{P}}; \mathcal{P}, \mathcal{P}) \quad (10)$$

where $\mu_{t, \mathcal{P}\mathcal{P}}$ is the mean of $p(x_i; \mathcal{P}, \mathcal{P})$. Same as above, the term $\log p(\mu_{t, \mathcal{P}\mathcal{P}}; \mathcal{P}, \mathcal{P})$ has been added to make VS $I_{\mathcal{F}}(x_i)$ zero at the centre (mean) of the distribution. It is clear that if the values of both population and patient probability density values decrease, which is a sign of deterioration, the VSI increases, whereas if the probability density values increase, which is a sign of improvement, the VSI decreases.

[0059] Based on equation (12), and according to the equations (1) and (5), after resampling we have:

$$p(x_i; \mathcal{P}, \mathcal{P}) \propto \frac{1}{\Sigma_{i=1}^{N_{\mathcal{P}}} \Sigma_{j=1}^{N_{\mathcal{P}}}} \mathcal{G}(x_i; x_i^j, \mathbf{C}_{\mathcal{P}}) \frac{1}{\Sigma_{i=1}^{N_{\mathcal{P}}} \Sigma_{j=1}^{N_{\mathcal{P}}}} \mathcal{G}(x_i; \mu^j, \mathbf{C}_{\mathcal{P}}) \quad (11)$$

[0060] Since the product of two Gaussian probability density functions (pdf) is Gaussian itself, this pdf is a Gaussian mixture model. By algebraic manipulation of equation (14), we obtain:

$$p(x_i; \mathcal{P}, \mathcal{P}) \propto \Sigma_{i=1}^{N_{\mathcal{P}}} \Sigma_{j=1}^{N_{\mathcal{P}}} \mathcal{G}(x_i; \mu_{t, \mathcal{P}\mathcal{P}}^{ij}, \mathbf{C}_{\mathcal{P}\mathcal{P}}) \quad (12)$$

where

$$\mu_{t, \mathcal{P}\mathcal{P}}^{ij} = (\mathbf{C}_{\mathcal{P}}^{-1} + \mathbf{C}_{\mathcal{P}}^{-1})^{-1} (\mathbf{C}_{\mathcal{P}}^{-1} x_i^j + \mathbf{C}_{\mathcal{P}\mathcal{P}}^{-1} \mu^j), \quad \mathbf{C}_{\mathcal{P}\mathcal{P}}^{-1} = (\mathbf{C}_{\mathcal{P}}^{-1} + \mathbf{C}_{\mathcal{P}}^{-1})^{-1} \quad (13)$$

[0061] If it is assumed that the covariance is spherical, the above mean basically is the average of the means of population and patient pdf weighted according to their variance. Therefore, the personalised Visensia index is defined by:

$$VSI_{\mathcal{P}}(x_i) = -\log p(x_i | y_i; \mathcal{P}) p(x_i; \mathcal{P}) + \log p(\frac{1}{\Sigma_{i=1}^{N_{\mathcal{P}}} \Sigma_{j=1}^{N_{\mathcal{P}}}} \mu_{t, \mathcal{P}\mathcal{P}}^{ij}; \mathcal{P}, \mathcal{P}) \quad (14)$$

which can be used as a measure for the wellness of the patient.

[0062] Thus to obtain the second numerical index the incoming multi-parameter observation y_i (HR, BP, BR, SpO₂, Temp), is compared to the population model of equation (1) to read off its probability density $p(x_i; \mathcal{F})$ with respect to the population and to the patient-specific model of equation (5) to obtain its probability density $p(x_i; y_i; \mathcal{P})$ with respect to the patient's own status inserting this into equation (17) This is illustrated in step 135 in FIG. 13. Again the second term based on the mean of the two models is fixed (once the patient-specific model is frozen its mean would be recalculated when it is updated) and can be pre-calculated. If the combined index (or a smoothed version of it) exceeds a specified threshold, then an alert is generated at step 136.

Artefact Removal

[0063] Artefacts such as motion of probe detachment, or movement can provide spurious measurements and would result, generally, in a high instantaneous VSI value in the population model (which would normally be smoothed out to prevent false alerts). To stop updating the model based on the noisy measurements, the artefact-free data may be selected by clinicians or automatically. Automatic selection of the acquired data would be on the basis of the population model. For instance, acquired vital sign sets can be excluded by the following rules:

[0064] 1. Initially anything that would generate an alert based on the population model is excluded from the online learning.

[0065] 2. After a short time, sequences of VSI values have been acquired with a distribution about some mean value. At this point, further data points can be excluded based on this distribution.

[0066] 3. If exclusion on the above basis persists, this may be indicative of the fact that the patient's status is changing, and the learning would be re-initialised to step 1.

[0067] It may be the case that a patient's state is stable, but the vital signs are so abnormal that they would always generate a VSI alert, for example, if the heart rate is over 150, but the patient is deemed by a clinician to be "stable". With the learning system proposed above, the specific model would never learn. However, the acceptance for online learning criterion can be modified by allowing the operator (e.g. clinician) to treat certain vital signs (in this case the Heart Rate) to be treated as "missing variables" for the

purpose of accepting the measurements. A simple method for doing this is to replace that variable's value by the population mean value.

[0068] Therefore, in this approach when the vital signs are considerably small or large, the system stops updating the weights of kernels, as these parameters could be because of the artefact in the data set. This is not however a problem, in the updating example based on the particle filtering, the particles are sampled from the prior distribution.

EXAMPLES

[0069] In this section, we provide some of examples of applying the invention to real vital signs measurements. FIGS. 3(a) to (f) are examples of the population-based index (a) and of five vital signs; heart rate (HR), oxygen saturation (SpO₂), blood pressure (BP), temperature (Temp) and breathing rate (BR) (b)-(f). The "current observation" has been shown by a solid vertical line. FIG. 4 shows the first two principal components of the population-based statistical model, where the centre of the kernels is shown by circles, and the principal components of the current observation are shown by the large asterisk. In this example, the observations were first normalised (or they can be whitened) and we set the number of particles equal to 1000, the observation and transition noises were assumed to be white Gaussian with powers equal to 0.1 and 1, respectively.

[0070] The first two principal components of the centre of the kernels modelled by particles are shown again in FIG. 5 for population-based statistical model (circles), together with the centre of the Gaussian kernels (particles) for the patient-specific statistical model using filled triangles. The mean or mode of kernels represented by triangles can be considered as the normal value of vital signs for that specific patient at a given time.

[0071] The novelty scores calculated with respect to the population based statistical model (equation 11) patient (equation 13) and combined (equation 17) are shown in FIG. 6. If the patient and population novelty scores increase, the combined novelty score increases as well. The combined novelty score can be less than the patient-specific model based novelty score, if the vital signs move enough towards the centre of the population-based model. FIG. 7(a) to (f) shows an example of the population-based index and vital signs when the current observation is well-outside of the boundaries of the population-based statistical model (star in the top right corner in FIG. 8). In this case the patient-specific model (circles) still is inside of the population model (triangles), see FIG. 9. As explained above, this is because of the proposed state transition which lowers the impact of the outliers on the learning procedure. Note that there are some outliers outside of the population model, which are the result of sampling procedure which adds some noise to the new samples.

[0072] The invention may be embodied in a dedicated vital signs monitor or may be embodied in a computer system. In both cases the main parts of the system are, as illustrated in FIG. 10, an input 101 for receiving the five vital signs, a processor 102 for processing them and comparing them to the statistical models stored in memory 103 to produce numerical indices, and an output 104 which may include a display 105 for displaying the numerical indices, normally together with the individual vital signs themselves. The processor 102 is also responsible for establishing and

then updating the patient-specific model in line with the update procedure described above.

1. A method of monitoring the status of a patient comprising the steps of:

continuously measuring a plurality of patient vital signs to provide a succession of multi-parameter vital signs observations,

comparing each successive multi-parameter vital signs observation to two statistical models, a first of the statistical models being a population-based reference model mapping the probability distribution of multi-parameter vital signs observations for a population of different individuals, and a second of the statistical models being a patient-specific model which is a probability distribution based at least partly on previous multi-parameter vital signs observations for the patient being monitored,

calculating first and second numerical indices based respectively on probability densities of the current multi-parameter vital signs observation obtained by comparison with the first and second statistical models, wherein an increased value of said index represents an increase in abnormality with respect to the statistical model,

outputting an alert notification if the first numerical index exceeds a first threshold, and

outputting a status change notification if the second numerical index exceeds a second threshold.

2. A method according to claim 1 further comprising the step of updating the patient-specific model using the most recent multi-parameter vital signs observation.

3. A method according to claim 2 wherein step of updating the patient-specific model only takes place when the comparison of the most recently received multi-parameter vital signs observation with the population-based model indicates that the observation is within a predefined normal region of the model.

4. A method according to claim 1 wherein step of updating the patient-specific model is stopped when the patient-specific model is stable, and restarted when comparison of the most recent multi-parameter vital signs observation with the patient-specific model indicates a change in patient status.

5. A method according to claim 1 wherein the status change notification is only generated if said change corresponds to an increase in the first numerical index, generated with reference to the reference model, thereby indicative of a deterioration in the patient's condition.

6. A method according to claim 5 wherein the increase of the first numerical index is judged by comparing the value of the first numerical index obtained for the current multi-parameter vital signs observation to the value of the first numerical index obtained when the patient-specific model became stable.

7. A method according to claim 1 wherein the patient specific model is initialized to be identical to the population model or to have the same characteristics as the population model.

8. A method according to claim 1 wherein the patient-specific model and the population-based model are both Kernel Density Estimators.

9. A method according to claim 8 wherein the Kernel density estimator for the population-based model consists of a plurality of spherical Gaussian density functions having equal weights.

10. A method according to claim 2, wherein the patient specific model is updated by a Bayesian process comprising calculating from the patient-specific model the likelihood of the most recent multi-parameter vital signs observation, multiplying the model's probability distribution by the calculated likelihood and renormalizing it to form an updated probability distribution for the model.

11. A method according to claim 10 wherein the patient-specific model consists of a plurality of particles, being samples initially generated from the population-based model by random sampling, and the Bayesian update process comprises adding a noise term to each of the particles, computing each particle's likelihood on receipt of a multi-parameter vital signs observation, and then resampling the plurality of particles wherein the probability of selection of the particle is proportional to its computed likelihood.

12. A method according to claim 3 wherein updating of the patient-specific model is selectively allowable by a clinician even if the population-based model is showing abnormality.

13. A method according to claim 11 wherein resampling of the plurality of the particles is performed by sampling the nearest kernel in the population-based model to each particle wherein the probability of selection of the particle is proportional to its computed likelihood, thus ensuring that the patient specific model remains bounded by the population based model.

14. A method according to claim 12 wherein the clinician can instruct the measurement of one or more vital signs to be ignored, these being treated as missing variables in determining whether to update the patient-specific model.

15. A method according to claim 3, wherein the pre-defined region of normality for determining whether to update the patient-specific model is different from the region defined by the threshold for determining whether to generate an alert notification.

16. A method according to claim 1 wherein an alert or status change notification is only generated if the respective threshold is exceeded for more than a predefined number of multi-parameter vital signs observations which are successive or within a predetermined time period.

17. A method according to claim 16, wherein the alert is triggered by thresholding a combined measure derived from both the population-based model and the patient-specific model.

18. A method according to claim 1 wherein the vital signs measured comprise heart rate, breathing rate, blood pressure, blood oxygen saturation and body temperature.

19. An apparatus for executing the method of claim 1, the apparatus comprising an input for receiving the patient's vital signs measurements, a memory storing the statistical models, a processor programmed to execute the comparing and calculating steps and an output to output the alert and status change notifications.

20. Apparatus according to claim 18 wherein the output comprises a display for displaying the notifications, the display being adapted to display the measured vital signs and at least one of the first and second numerical indices.

* * * * *

专利名称(译)	用于监测患者状态的方法和设备		
公开(公告)号	US20170071484A1	公开(公告)日	2017-03-16
申请号	US15/310652	申请日	2015-03-31
[标]申请(专利权)人(译)	OBS医疗		
申请(专利权)人(译)	OBS医药有限		
当前申请(专利权)人(译)	OBS医药有限		
[标]发明人	STRACHAN IAIN GUY DAVID MOHSENI HAMID REZA		
发明人	STRACHAN, IAIN GUY DAVID MOHSENI, HAMID REZA		
IPC分类号	A61B5/0205 G06F19/00 A61B5/00		
CPC分类号	A61B5/02055 A61B5/7264 A61B5/7282 G06F19/3431 A61B5/14542 A61B5/024 A61B5/0816 A61B5/021 G06F19/3437 A61B5/02405 A61B5/0245 A61B5/0452 A61B5/14551 G06F19/00 G16H50/30 G16H50/50 Y02A90/26		
优先权	2014008469 2014-05-13 GB		
外部链接	Espacenet USPTO		

摘要(译)

一种用于基于多个生命体征测量来监测患者状态的方法和设备，所述生命体征测量例如心率，呼吸速率，血压，温度，氧饱和度等。在每个时间点，一组这样的生命体征测量值是与两个统计模型相比，一个是基于人群的统计模型，表示来自大量个体的测量的概率分布，第二个是基于来自被监测的特定患者的测量的患者特异性统计模型。基于患者的统计模型被初始化为与基于群体的模型相同，但随后以贝叶斯更新技术更新，以使其朝向患者的个体状态发展。通过将每个状态测量值与两个统计模型中的每一个进行比较来计算第一和第二新颖性指数。如果指数指示患者的状态在群体的预定义正常区域之外，则来自基于群体的模型的新颖性指数的阈值可以产生警报。来自患者特异性模型的数值指数的阈值可以指示患者状态的变化，这可以是改善或恶化。

