

Integrated monitoring and analysis for early warning of patient deterioration

L. Tarassenko^{1*}, A. Hann¹ and D. Young²

¹Department of Engineering Science, Parks Road, University of Oxford, Oxford OX1 3PJ, UK.

²Nuffield Department of Anaesthetics, University of Oxford, Adult Intensive Care Unit, John Radcliffe Hospital, Headley Way, Oxford OX3 9DU, UK

*Corresponding author. E-mail: lionel.tarassenko@eng.ox.ac.uk

Recently there has been an upsurge of interest in strategies for detecting at-risk patients in order to trigger the timely intervention of a Medical Emergency Team (MET), also known as a Rapid Response Team (RRT). We review a real-time automated system, BioSign, which tracks patient status by combining information from vital signs monitored non-invasively on the general ward. BioSign fuses the vital signs in order to produce a single-parameter representation of patient status, the Patient Status Index. The data fusion method adopted in BioSign is a probabilistic model of normality in five dimensions, previously learnt from the vital sign data acquired from a representative sample of patients. BioSign alerts occur either when a single vital sign deviates by close to ± 3 standard deviations from its normal value or when two or more vital signs depart from normality, but by a smaller amount. In a trial with high-risk elective/emergency surgery or medical patients, BioSign alerts were generated, on average, every 8 hours; 95% of these were classified as 'True' by clinical experts. Retrospective analysis has also shown that the data fusion algorithm in BioSign is capable of detecting critical events in advance of single-channel alerts.

Br J Anaesth 2006

Keywords: critical care; data fusion; early warning; signal analysis

Current ward care of critically ill patients is suboptimal;^{3 18} the consequences of which may be costly in terms of time, resources and patient outcomes. As many as 80% of ward patients have physiological parameters outside normal ranges within the 24 h preceding intensive care unit (ICU) admission,^{10 11} while up to three-quarters of ward patients who are admitted to the ICU have had at least one potentially life-threatening antecedent factor in the 8 h beforehand.^{13 20}

The incidence of abnormal physiological parameters before cardiac arrest has been reported to be as follows:²⁵ tachypnoea (58%), tachycardia (54%), altered mental state (46%), arterial hypotension (46%), poor urine output (29%), pyrexia (13%), arterial hypertension (8%) and hypothermia (4%). Others²² have indicated that the most frequent clinical deterioration seen before cardiac arrest is impaired respiratory or mental function, with the respiratory rate elevated well above normal in a majority of patients. In a study in which the physiological parameters of patients before cardiac arrest were compared with those of 'normal' patients,⁷ the occurrence of respiratory rate rising above 27 bpm at least once during a 72-h period was found to have a sensitivity of 0.54 and a specificity of 0.83 in predicting cardiac arrest. A study of patients with head injuries reported avoidable factors contributing to death in 30% of patients, most

of these being failure to recognize hypotension and respiratory difficulty leading to hypoxia.¹⁴

It is clear that the failure to respond to patient deterioration promptly and appropriately can lead to increased morbidity and mortality,^{17 26} increased requirement for intensive care^{16 18} and elevated costs.¹⁸ As a result, there has been an upsurge of interest in strategies for detecting at-risk patients in order to trigger the timely intervention of a medical emergency team (MET), also known as a rapid response team (RRT).^{1 2 5} These approaches are based on the premise that early recognition of physiological abnormalities coupled with the rapid intervention of suitably trained staff may result in an improvement in functional outcome or mortality rate.⁸

METs rely either on alert limits on single parameters or on a scoring system (early-warning scores—EWS) to trigger the call to the patient's bedside. The original EWS system¹⁹ was developed as a simpler version of the Acute Physiology and Chronic Health status Evaluation system which is used to evaluate the severity of illness in ICU. The EWS scoring system is usually based on physiological variables such as heart rate, respiratory rate and blood pressure as well as a measure of patient alertness.²⁶ The scores increase as the patient diverges from normality, and values above a set threshold mandate a call to the MET.

A prospective study to validate such an EWS system calculated the scores of 709 medical emergency admissions for up to 5 days and analysed the relationship with the outcome.²⁶ A maximum score of 5 or more was associated with increased risk of death, ICU admission and high-dependency unit admission. A later study by the same team monitored 1695 medical acute admissions using the same EWS system, with supporting protocols to trigger medical and critical care review. It showed no change in the outcome of acute medical admissions, or in the outcome of those patients highlighted to be at risk, although a trend towards earlier ICU referral and admission was observed.²⁷ More promising results were obtained when another version of EWS was combined with call-out criteria for a critical care outreach service.²¹ This led to a reduction in ICU emergency admission rate, with shorter stay and lower mortality for the emergency patients.

In spite of the benefits recorded in most of the observational studies with EWS systems, there has not really been any attempt to automate the process of calculating the scores. The closest perhaps is the use of a 'logic module' to analyse real-time data from multiple devices in the ICU.²³ The concept behind this was the introduction of logic in an attempt to differentiate clinically insignificant readings from those related to physiological trends detrimental to the patient. The user firstly selected a set of physiological parameters to be trended, for example the difference in heart rate variability over 1 min in the last minute and that from 3 min ago. Then a threshold was chosen to classify the trend as one of the following: above the threshold, below the threshold, or unknown because of missing values. The result of each test was given a score, and the sum of the scores was compared with a second threshold to determine whether or not an alert should be generated. The logic module was evaluated by comparing its alerts with those generated from single parameters ('single-channel alerts') and classified as true by ICU staff. Results from an analysis of 120 h of ICU data showed an increase in positive predictive value [TP/(TP+FP), where TP=true positive and FP=false positive] from 0.03 using single-channel alerts to 0.32 when using the module.

The choice of thresholds and score values in the EWS scoring systems is based on clinical experience alone, and the systems are difficult to validate because of the large number of parameters.⁴⁹ EWS are usually determined from 4-hourly observations²⁴ at best and their calculation increases the workload of nursing staff. It has been suggested that the gap between observations contributed in part to the failure of the 23-hospital MERIT study¹² to show a statistically significant improvement in patient outcome as a result of the introduction of METs. For example, the Institute of Healthcare Improvement's review of the MERIT study¹⁵ concludes that 'monitoring of patient status and vital signs may be a vulnerable part of the MET/RRT process. The RRT cannot be called if no-one is noticing the patients' changing status. Classical monitoring methods and

schedules may not be up to the task of properly supporting and triggering the RRT capability.'

With all of the above in mind, we review here a *data-driven*, rather than rule-based, automated system, BioSign, which tracks patient status in *real-time* by combining (or fusing) the patient's vital sign data acquired from monitors on the general ward. The five parameters of heart rate, breathing rate, blood pressure, arterial oxygen saturation (SaO_2) and skin temperature are evaluated every 5 s, except for the blood pressure which is measured every 30 min using oscillometry with an inflatable cuff placed over the medial artery. As a result of *learning* a model of normality directly from these data, we do not have the problem of extracting and/or validating the rules which would have to be used in an expert-system implementation of EWS.

BioSign—overview

BioSign fuses the five vital signs in order to produce a single-parameter representation of patient status, the patient status index (PSI). The data fusion method adopted in BioSign is a probabilistic model of normality in five dimensions, previously *learnt from the vital sign data* acquired from a representative sample of patients (in this case, high-risk adult patients), the 'training data set'. The model of normality is stored in BioSign and used to evaluate the probability that the vital sign parameters acquired from the patient being monitored can be considered to be normal, with respect to those in the training data set. An alert is generated whenever the vital sign parameters are abnormal enough to be outside the 'envelope of normality' and cause the PSI, whose value increases with abnormality, to cross the alerting threshold.

BioSign has been developed for use as a real-time, early-warning system for triggering the intervention of a MET or RRT. It can be integrated within existing patient monitors or used as part of a central station (in critical care, in anaesthesia or on the general ward) to facilitate the effective monitoring of greater numbers of patients without increasing the burden on staff.

BioSign—constructing the model of normality

Vital sign data were collected from 150 general-ward patients at the John Radcliffe Hospital in Oxford, between 2001 and 2003. These patients were connected to a multi-parameter monitor for, on average, 24 h per patient. The training data set therefore included approximately 3500 h of vital sign data. Ethics approval was granted before the data collection exercise started and informed consent was obtained in every case. Patients were drawn, but not exclusively, from the following 'high-risk' patient groups:

- patients monitored for at least 24 h after a myocardial infarct and again for a few hours 5 days later;

- patients with severe heart failure;
- patients with acute respiratory problems (for example, acute asthma or pneumonia);
- elderly patients with hip fracture, who were monitored both before and after operation.

The BioSign model of normality is the unconditional probability density function (pdf), $\hat{p}(\mathbf{x})$, of the training data, where $\mathbf{x}=\{x_1, x_2, \dots, x_5\}$ is the vector of vital sign parameters, with x_1 =heart rate, x_2 =breathing rate, etc. As the five parameters have different dynamic ranges (an increase of 0.5°C in temperature is much more significant than an increase of 0.5 mm Hg in blood pressure), they need to be normalized before they can form the vector \mathbf{x} . The distributions of the five parameters in the training data set were found to be near-Gaussian, except for Sa_{O_2} , which has a one-sided distribution (it cannot go above 100%). Hence a standard zero-mean, unit-variance transformation is applied to the parameters to normalize them. The mean values of the five parameters for the training data set are shown in Table 1.

The unconditional pdf, $\hat{p}(\mathbf{x})$, of the training data is estimated using a combination of k -means clustering and Parzen Windows.⁶ Firstly, the k -means clustering algorithm is used to select 500 cluster centres, or prototype patterns, from the tens of thousands of normalized vectors in the training data set. Each of the prototype patterns \mathbf{x}_j is then a kernel in the Parzen Windows estimator of the pdf, given by the equation below:

$$\hat{p}(\mathbf{x}) = \frac{1}{N(2\pi)^{d/2}\sigma^d} \sum_{j=1}^N \exp\left(-\frac{\|\mathbf{x}-\mathbf{x}_j\|^2}{2\sigma^2}\right)$$

where each spherical kernel has the same global width σ and d is equal to 5.

We define the ‘PSI’ to quantify departures from normality [i.e. low values of $p(\mathbf{x})$] so that alerts can be generated when this index increases above a threshold value. The PSI is calculated by transforming the probability so that abnormality increases along the vertical axis (the horizontal axis being time in a trend plot):

$$\text{Patient status index} = \log_e \left[\frac{1}{\hat{p}(\mathbf{x})} \right]$$

A PSI of 3.0, corresponding to a value of $p=0.05$, was chosen for the alerting threshold.²⁸ Alerts are generated

Table 1 Mean values for each of the five BioSign parameters in the training data set. The systolic–diastolic average (SDA) is calculated from the blood pressure readings as the arithmetic mean of the two pressure readings; that is, $SDA = \frac{1}{2}(\text{diastolic} + \text{systolic})$

Parameter	Mean
Heart rate (beats min ⁻¹)	83.8
SDA blood pressure (mm Hg)	94.7
Oxygen saturation (%)	95.2
Skin temperature (°C)	36.0
Breathing rate (bpm)	18.3

when the PSI is above the threshold of 3.0 for 80% of the time in a time window of fixed length (usually 5 min). The threshold was chosen to correspond approximately to the PSI generated when one vital sign is ± 3 SDs away from its mean value in the training data set (the other four being assumed to be normal).

BioSign—results from clinical studies

When BioSign is used to monitor a patient on the ward, the vital sign parameters are pre-processed before normalization and estimation of $p(\mathbf{x})$. Short-term median filtering is applied to remove artifactual values caused by patient movement, for example, and longer-term, or historic, median filtering is used to cope with a missing parameter stream. If no valid measurement of a parameter, with the exception of blood pressure (which is only measured every 30 min, at most), has been acquired for 1 min, BioSign uses the value from the historic median filter. This value is the median value of the last 5 min of valid data for that parameter. If a new measurement is not received for 30 min, the mean value in the training data set is used instead. (This may occur if a probe becomes disconnected from the patient or a signal degrades for a prolonged period of time.) The effect of this is that the missing parameter no longer influences the calculation of $\hat{p}(\mathbf{x})$, as it is replaced by its ‘most normal’ value in the parameter vector \mathbf{x} .

BioSign has recently been used in a number of clinical studies in Europe and the US to analyse the vital sign data from patients undergoing high-risk surgery or after emergency admission for acute non-surgical conditions. The main study design was a randomized controlled trial (RCT) in Oxford of the effect of mandated continuous physiological monitoring on the clinically significant event rates in patients with a high risk of death from medical or surgical conditions.²⁹ Patients in the trial were randomly assigned to receive mandated continuous five-parameter physiological monitoring for up to 72 h post-surgery or post acute medical admission (Intervention Group) or to receive usual ward care involving intermittent, single channel or multi-channel physiological monitoring at the attending staff’s discretion, usually an ECG monitor, a pulse oximeter, or both (Control Group). Nested within the ‘monitoring’ arm of the RCT was a within-patient assessment of BioSign alerts (PSI>3.0), with the PSI retrospectively calculated from the five-parameter vital sign data continuously recorded and logged for the Intervention Group patients. This trial design was chosen as it separates the effects of monitoring from the assessment of BioSign.

In the Intervention Group, there were 168 patients with valid data. Of these patients, 105 had at least one BioSign alert, whereas 63 had none. The mean alert rate was 0.13 h⁻¹, that is 1 every 7.8 h on average. There were 690 BioSign alerts in total, which were subsequently analysed by two senior clinicians. After reviewing trend plots of the five vital signs in the 5 min leading up to each

BioSign alert, they deemed 652 of the alerts to be true episodes of severe physiological abnormalities, that is the BioSign alert was classified as ‘True’ in 95% of cases (positive predictive value of 0.95). In 30 cases, shortcomings in the rejection of motion artifact on the pulse oximeter or electrical impedance pneumograph signals gave rise to a spurious alert as a result of artifactual values of Sa_{O_2} or breathing rate, and in eight cases the clinicians could not determine whether the recordings were valid data or artifact.

We also investigated whether the data fusion provided by BioSign is able to give warning of patient deterioration *before* the single-channel alerts which would be generated by one of the individual parameters going out of range. The PSI alerting threshold of 3.0 can be crossed for two reasons: a single parameter deviating by close to ± 3 SDs from its normal value (i.e. the mean value Sa_{O_2} for that parameter in the training data set—see Table 1), an event which we call a ‘Type A alert’; or, alternatively, two or more parameters moving away from normality at the same time, but by a smaller amount (a ‘Type B alert’).

We therefore asked the following question: ‘How often does a BioSign alert (PSI>3.0) precede an alert generated by any one of the vital signs recorded with a single-channel monitor?’ The single-channel alert limits used for comparison were those published in the literature⁵ by one of the leading advocates of METs/RRTs. All types of single-channel alerts from the Oxford RCT were selected for the early-warning analysis although some of these events occurred too rarely for the resulting statistics to be meaningful. There were 49 instances of early warning before low Sa_{O_2} alerts (<85%), 11 instances before high diastolic blood pressure alerts (>110 mm Hg) and 10 instances before low systolic blood pressure alerts (<80 mm Hg). When the low Sa_{O_2} alerts were analysed more closely, it was found that

there were between 6 s and 5 min of ‘early warning’ for 31 of 49 of these alerts (see Fig. 1).

In these 31 cases, the decreasing Sa_{O_2} value combined with other, mostly normal, vital signs to cause the BioSign alert to trigger *just before* the single-channel alert, that is the latter was triggered by the same event as the impending single-channel alert (Type A alert). For 18 cases of 49, however, the BioSign alert occurred between 5 and 126 min earlier than the low Sa_{O_2} alert (on average, 60.1 min). This is because the less pronounced deterioration in Sa_{O_2} at the time of the BioSign alert was accompanied by other evidence of physiological derangement (Type B alert); at least one of the other four parameters also had an abnormal value. This provides evidence of the early-warning capability of BioSign as a result of data fusion.

Table 2 provides an analysis of the vital sign profile at the time of the 18 Type B (early warning) BioSign alerts. In 16 of 18 cases, the breathing rate was elevated and in 15 of 18 cases the oxygen saturation was below the mean value in the training set (95.2%). In 12 cases, the patient’s skin temperature was higher than normal. For 13 of the 18 cases, the breathing rate or oxygen saturation was the most ‘abnormal’ parameter at the time of the BioSign alert. Thus a pattern begins to emerge: sudden deteriorations in oxygen saturation are often preceded by elevated breathing rates (shallow breathing) and a gradual decrease in the oxygen saturation itself.

Conclusions

BioSign can be used to trigger calls to a MET/RRT (but the patient’s mental status and urine output also need to be assessed on a regular basis by nursing staff). BioSign




Fig 1 Cumulative frequency plot of the number of minutes of early warning with respect to single-channel alerts triggered by low Sa_{O_2} (85%). Early warning alerts greater than 120 min are plotted as occurring at 120 min.

Table 2 Vital signs at the time of BioSign alerts that gave >5 min early warning of low Sa_o₂. Note that skin temperature is on average 1.2°C lower than core temperature (in the training data set)

	Heart rate (beats min ⁻¹)	Breathing rate (bpm)	Sa _o ₂ (%)	Skin temp. (°C)	Systolic/diastolic av. (mm Hg)
Average value at the time of BioSign alert	80.8	26.0	90.4	34.8	96.9
Above/below normal value for that parameter	8/10	16/2	3/15	6/12	9/9
Most abnormal vital sign at the time of the BioSign alert	0	6	7	3	2

alerts occur either when a single parameter deviates by close to ± 3 SDs from its normal value or when two or more parameters depart from normality, but by a smaller amount. In a trial with high-risk elective/emergency surgery or medical patients, BioSign alerts were generated, on average, every 8 h; 95% of these were classified as 'True' by clinical experts. Retrospective analysis has also shown that the data fusion algorithm in BioSign is capable of detecting critical events in advance of single-channel alerts.

Acknowledgement

A.H. is funded by an Engineering & Physical Sciences Research Council (EPSRC) studentship.

References

- Bellomo R, Goldsmith D, Uchino S, et al. A prospective before-and-after trial of a medical emergency team. *Med J Aust* 2003; **179**: 283–7
- Buist MD, Moore GE, Bernard SA, et al. Effects of a medical emergency team on reduction of incidence of and mortality from unexpected cardiac arrest in hospital: preliminary study. *Br Med J* 2002; **324**: 387–90
- Chellel A, Fraser J, Fender V, et al. Nursing observations on ward patients at risk of critical illness. *Nurs Times* 2002; **98**: 36–9
- Cuthbertson BH. Outreach critical care—cash for no questions? *Br J Anaesth* 2003; **90**: 4–6
- De Vita MA, Braithwaite RS, Mahidhara R, et al. Use of medical emergency team responses to reduce cardiopulmonary arrests. *Qual Saf Health Care* 2004; **13**: 251–4
- Duda RO, Hart PE, Stork DG. *Pattern Classification*, 2nd Edn. New York: John Wiley and Sons,
- Fiesemann JF, Hendryx MS, Hlems CM, et al. Respiratory rate predicts cardiopulmonary arrest for internal medicine inpatients. *J Gen Intern Med* 1993; **8**: 354–60
- Garcea G, Thomasset S, McClelland L, et al. Impact of a critical care outreach team on critical care readmissions and mortality. *Acta Anaesthesiol Scand* 2004; **48**: 1096–100
- Goldhill DR. The critically ill: following your MEWS. *QJM* 2001; **94**: 507–10
- Goldhill DR, Worthington L, Mulcahy A, et al. The patient-at-risk team: identifying and managing seriously ill ward patients. *Anaesthesia* 1999; **54**: 853–60
- Goldhill DR, White SA, Sumner A. Physiological values and procedures in the 24 hours before ICU admission from the ward. *Anaesthesia* 1999; **54**: 529–34
- Hillman K, Chen J, Cretikos M, et al. Introduction of the medical emergency team (MET) system: a cluster-randomised controlled trial. *Lancet* 2005; **365**: 2091–7
- Hillman KM, Bristow PJ, Chey T, et al. Duration of life-threatening antecedents prior to intensive care admission. *Intensive Care Med* 2002; **28**: 1629–34
- Jeffreys RV, Jones JJ. Avoidable factors contributing to the death of head injury patients in general hospitals in Mersey region. *Lancet* 1981; **2**: 459–61
- Institute for Healthcare Improvement. The 'MERIT' trial of medical emergency teams in Australia: an analysis of findings and implications for the 100,000 lives campaign. [Available from <http://www.ihl.org/NR/rdonlyres/F3401FEF-2179-4403-8F67-B9255C57E207/0/LancetAnalysis81505.pdf> (accessed 23 March 2006)]
- McArthur-Rouse F. Critical care outreach services and early warning scoring systems: a review of the literature. *J Adv Nurs* 2001; **36**: 696–704
- McGloin H, Adam SK, Singer M. Unexpected deaths and referrals to intensive care of patients on general wards. Are some cases potentially avoidable? *J R Coll Physicians Lond* 1999; **33**: 255–9
- McQuillan P, Pilkington S, Allan A, et al. Confidential inquiry into quality of care before admission to intensive care. *Br Med J* 1998; **316**: 1853–8
- Morgan RJM, Williams F, Wright MM. An early warning score for the early detection of patients with impending illness. *Clin Intensive Care* 1997; **8**: 100
- Nurmi J, Harjola V, Rosenberg P, Castrén M. Pre-arrest signs in different types of hospitals. *Crit Care* 2004; **8** (Suppl. 1): P290
- Pittard AJ. Out of reach? Assessing the impact of introducing a critical care outreach service. *Anaesthesia* 2003; **58**: 882–5
- Schein RMH, Hazday N, Nolan J, et al. Clinical antecedents to in-hospital cardiopulmonary arrest. *Chest* 1990; **98**: 1388–92
- Schoenberg R, Sands DZ, Safran C. Making ICU alarms meaningful: a comparison of traditional vs. trend-based algorithms. *Proc AMIA Symp* 1999; 379–83
- Sharpley JT, Holden JC. Introducing an early warning scoring system in a district general hospital. *Nurs Crit Care* 2004; **9**: 98–103
- Smith AF, Wood J. Can some hospitals cardio-respiratory arrests be prevented? A prospective study. *Resuscitation* 1998; **37**: 133–7
- Subbe CP, Kruger M, Rutherford P, Gemmel L. Validation of a modified Early Warning Score in medical admissions. *QJM* 2001; **94**: 521–6
- Subbe CP, Davies RG, Williams E, et al. Effect of introducing the modified Early Warning Score on clinical outcomes, cardio-pulmonary arrests and intensive care utilisation in acute medical admissions. *Anaesthesia* 2003; **58**: 797–802
- Tarassenko L, Hann A, Patterson A and et al. BioSign™: Multi-parameter monitoring for early warning of patient deterioration. *Proc 3rd IEE International Seminar on Medical Applications of Signal Processing* 2005; 71–6
- Young JD, Barber VS, Watkinson PJ, et al. Effect of continuous electronic physiological monitoring on the adverse event rate in high risk medical and surgical patients: randomised controlled trial. Submitted for publication