

# BIOSIGN™ : MULTI-PARAMETER MONITORING FOR EARLY WARNING OF PATIENT DETERIORATION

L. Tarassenko, A. Hann, A. Patterson<sup>+</sup>, E. Braithwaite<sup>+</sup>, K. Davidson\*, V. Barber\* & D. Young\*

Signal Processing & Neural Networks Group, University of Oxford

<sup>+</sup>Oxford BioSignals Ltd, Magdalen Centre, Oxford Science Park

\*Nuffield Department of Anaesthetics, John Radcliffe Hospital, Oxford

**Keywords:** Patient monitoring, vital signs, early warning systems, data fusion, models of normality.

## Abstract

Early recognition of abnormalities in the physiological parameters of hospital patients followed by rapid intervention should result in an improvement in functional outcome or mortality rate. We have developed a real-time system, BioSign™, capable of analysing physiological parameters in order to identify adverse trends in multi-parameter space (departure from “normality”) and prompt clinical staff to intervene. The model of normality is based on five vital signs, which can all be recorded non-invasively: the heart rate, blood pressure, arterial oxygen saturation, respiration rate and temperature. We have evaluated the trained model of normality in a two-year Randomised Controlled Trial of 405 patients monitored on general medical and surgical wards. Preliminary results show that BioSign™ can provide early warning of changes in clinical status and deterioration in patient condition.

## 1 Introduction

Catastrophic deterioration of patients in hospital is often preceded by a change in physiological parameters that may pass unnoticed for a considerable period of time. It has been shown, for example, that at least 50% of the patients who suffer a cardiac arrest have one or more physiological disturbances recorded on admission or within the 8 hours preceding the arrest [1]. Failure to recognise patterns of physiological disturbance may contribute to death in patients with head injuries [2] and to unscheduled admission to the ICU [3].

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 [4]. 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 [5]. Even if patients are continuously monitored on the ward, it is

often the case that important changes signalling a catastrophic decline go undetected, for the following reasons:

- the pattern of decline is not recognised or is missed by the ward staff (the recording of data, “observations”, only occurs at 4-hourly intervals [6]);
- existing bedside monitors only alarm when a single variable falls outside a wide band of values defining “normality” for that variable;
- simple limit alarms are highly sensitive and not at all specific;
- artefactual values, often caused by patient motion, cause large numbers of false positives, which may lead to the disabling of alarms by medical personnel [7].

The effect of the last of these can be minimised through careful signal (pre)-processing. The first three in the above list require a more fundamental shift in approach, however: the one described in this paper is based on the *data fusion* of continuously-monitored physiological variables using *trained models*.

Medical Emergency Teams rely either on alarm limits on single parameters or on a scoring system (Early Warning Scores – EWS) to trigger the call to the patient’s bedside. The scoring system is usually based on five physiological variables: heart rate, respiratory rate, blood pressure, temperature and a measure of alertness or coma [8]. The scores increase as the patient diverges from normality, and values above a set threshold mandate a call to the Medical Emergency Team. The scoring systems have all been derived empirically and rely on the linear summation of the scores assigned to each parameter individually, with no knowledge of the values of the other parameters. We have developed a real-time system, BioSign™, capable of identifying adverse trends in vital signs from prior knowledge captured in a trained model of normality. The model of normality is based on the following five vital signs: heart rate, respiratory rate, blood pressure, temperature and arterial oxygen saturation ( $S_pO_2$ ). We have substituted the latter for the assessment of alertness or coma, because it can be recorded both non-invasively and continuously using pulse oximetry; it provides

valuable clinical information; it does not require the intervention of a nurse (as with alertness and coma) and it is now routinely monitored in high-risk patients.

The vital signs are fused in a model of normality which is trained on a data set acquired from a representative sample of patients. An approximation to the unconditional probability density function (pdf) of the normal vital sign data in the training set is used as the model of normality. This model is stored in BioSign™ and used to evaluate the probability that the set of vital signs acquired second-by-second from the patient being monitored can be considered to be normal. Whenever this probability falls below a threshold previously determined (see section 7), an alert is generated, which could be used to trigger the intervention of a Medical Emergency Team.

BioSign™ has just completed a two-year evaluation in a Randomised Controlled Trial (RCT) of 405 patients monitored on the general medical and surgical wards at the John Radcliffe Hospital in Oxford, UK. In the sections that follow, we describe the acquisition of the vital sign data, the pre-processing of the five parameters, the training of the model of normality and, finally, the preliminary findings from the RCT.

## 2 Vital sign data acquisition

The five parameters used in the BioSign™ model of normality are collected using multi-parameter patient monitors, which record the following signals:

- One, two or three leads of continuous ECG;
- Systolic and diastolic blood pressure readings recorded at periodic intervals using an oscillometric method with a blood pressure cuff placed over the medial artery;
- A single, continuous, electrical impedance pneumograph signal from which respiration rate can be calculated;
- Arterial oxygen saturation (SpO<sub>2</sub>) measurements from pulse oximetry;
- Continuous skin temperature measured from a thermistor attached to the arm underneath the blood pressure cuff and insulated from the surrounding environment (note that, even under these conditions, skin temperature is between 1.0 and 1.5 °C below core temperature).

The heart rate and respiration rate are derived according to industry-standard algorithms. The systolic-diastolic average (SDA) is calculated from the blood pressure readings as the arithmetic mean of the two pressure readings; i.e.  $SDA = \frac{1}{2}(\text{diastolic} + \text{systolic})$ . The SpO<sub>2</sub> and temperature readings are sampled once per second prior to further pre-processing (see below).

## 3 Training data set

The model of normality was constructed from vital sign data collected from patients at the John Radcliffe Hospital in Oxford, between 2001 and 2003. Prior to the commencement of the data collection period, an international advisory board consisting of members from Europe and North America was asked to define a suitable patient population to model. This population was used to generate the BioSign™ model. It consisted of 150 patients who were connected to a multi-parameter monitor for, on average, twenty-four hours per patient. The data set therefore included approximately 3,500 hours of patient data. Patients were studied on the general wards of the John Radcliffe Hospital. Ethics approval was granted before the data collection exercise started and informed consent was obtained in every case. Patients included, but were not limited to, the following “high-risk” patient groups:

- patients monitored for at least 24 hours after a myocardial infarct and again for a few hours five 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 pre- and post-operatively (this is a patient group for which the mortality rate one year after the operation is of the order of 20%).

## 4 Parameter normalisation

In order to construct the model of normality, the unconditional probability density  $\hat{p}$  of  $\mathbf{x}$ , the vector of vital sign parameters, must be estimated. As the five parameters have different dynamic ranges, they need to be normalised before they can form the feature vector  $\mathbf{x}$ . Such normalisation could take place on the basis of clinical knowledge (for example, is an increase in body temperature of 0.5 °C equivalent to a change in heart rate of 30 bpm?) and indeed the EWS scoring system does incorporate some clinical judgement. However, the scores have never been validated and we have opted instead for data-driven normalisation.

The distributions of the five parameters in the training set are near-Gaussian, except for SpO<sub>2</sub> which has a one-sided distribution (it cannot go above 100%). Hence a zero-mean, unit-variance transformation can be applied to the parameters:

$$x_n = \frac{x - \mu}{(s.d.)}, \quad (1)$$

where  $x_n$  is the normalised parameter,  $\mu$  is the mean value of the parameter in the training database and  $(s.d.)$  is the standard deviation of that parameter in the same database. The mean values for each of the five parameters in the training set are shown in the second column of Table 1.

It is important to check the validity of the values of  $\mu$  and (*s.d*) for each parameter, using an *independent* validation data set acquired from a different set of patients but drawn from a similar population. The 150 patients in the validation data set were monitored on average for forty hours each, providing nearly 6,000 hours of validation data.

The “tolerances” (third column of Table 1) around each mean value in the training data set were chosen *a priori*, on the basis of clinical experience. If the tolerances were exceeded on the validation data set, then the latter could be said to have normalisation characteristics different from those of the training data set.

Parameter	Mean	Tolerance
Heart rate (BPM)	83.8	$\pm 5$ beats per minute
SDA blood pressure (mmHg)	94.7	$\pm 5$ mmHg
Oxygen saturation (%)	95.2	$\pm 1\%$
Skin temperature ( $^{\circ}\text{C}$ )	36.0	$\pm 1^{\circ}\text{C}$
Respiration rate (RPM)	18.3	$\pm 3$ breaths per minute

**Table 1** – Mean values (with associated tolerances) for each of the five BioSign<sup>TM</sup> parameters in the training data set

Table 2 below shows the means of the five parameters for the validation set. It can be seen from the third column (the differences in the mean values between the two data sets) that none of the mean values in the validation set are outside the pre-set tolerances.

Parameter	Mean	Difference between means
Heart rate (BPM)	84.5	0.7
SDA blood pressure (mmHg)	91.1	3.6
Oxygen saturation (%)	96.0	0.8
Skin temperature ( $^{\circ}\text{C}$ )	35.5	0.5
Respiration rate (RPM)	20.2	1.9

**Table 2** – Mean values for each of the five BioSign<sup>TM</sup> parameters in the validation data set (with differences with respect to training set)

## 5 Parameter pre-filtering

A further step of filtering is then applied to the normalised parameters. This step involves short-term median filtering for noise removal and longer-term, or historic, median filtering for coping with a missing parameter stream (see below). For example, mechanical or movement artefact can cause outliers

within the vital-sign parameters. Short-term filtering removes these so that they do not affect the calculation of  $\hat{p}(\mathbf{x})$ .

The following short-term median filters are applied to each parameter:

Parameter	Short-term Median Filter Length
Heart Rate	5
SDA Blood Pressure	No filtering
Oxygen Saturation (SpO <sub>2</sub> )	5
Skin Temperature	5
Respiration Rate	9

**Table 3** – Short-term median filter lengths for each of the five BioSign<sup>TM</sup> parameters

For the first and fifth parameters in Table 3, the short-term median filter lengths were chosen in relation to typical numbers of heart beats/breaths per minute. The third and fourth parameters are both sampled at a rate of one sample per second and hence the filter length for both of these is equivalent to a period of 5 seconds.

If no valid measurement of a parameter, with the exception of blood pressure (which is only measured every 30 minutes, at most), has been acquired for one minute, BioSign<sup>TM</sup> uses the value from a historic median filter. This value is the median value of the last five minutes of valid data for that parameter. If a new measurement is not received for thirty minutes, the mean value in the training database 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  $\hat{p}(\mathbf{x})$ .

## 6 Computing the unconditional probability

There are a number of methods available for estimating the unconditional pdf,  $\hat{p}$ , of the training data, the BioSign<sup>TM</sup> model of normality. For example, a Gaussian Mixture Model (GMM) representation has the flexibility required to model general distributions, and there is certainly no shortage of training patterns (with 3,500 hours of patient data) with which to set the free parameters of a full-covariance GMM. Instead, we chose to use a combination of k-means clustering, Parzen Windows and data visualisation for enhanced data understanding.

The k-means clustering algorithm is used to select 500 cluster centres, or prototype patterns, from the tens of thousands of feature vectors in the training database. Each of the prototype patterns  $\mathbf{x}_{n,j}$  can then be a kernel in the Parzen Windows estimator of the pdf, given by the equation shown at the top of the next page.

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

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

The location of the 500 kernels in five-dimensional space can be visualised using a multi-dimensional scaling algorithm such as Sammon's mapping [9]. The image points in Sammon's map are distributed so that the Euclidean distances between pairs of these points are as close as possible to the Euclidean distances in  $d$ -dimensional space between the corresponding pairs of feature vectors. Alternatively, it is possible to use the Neuroscale algorithm [10], in which the mapping from the  $d$ -dimensional space to 2-D is parameterised using a Radial Basis Function (RBF) neural network. This allows interpolation and new patterns can be displayed using the map constructed from the training data set. The 2-D map for the 500 prototype patterns extracted from the training set is shown in Figure 1 below.

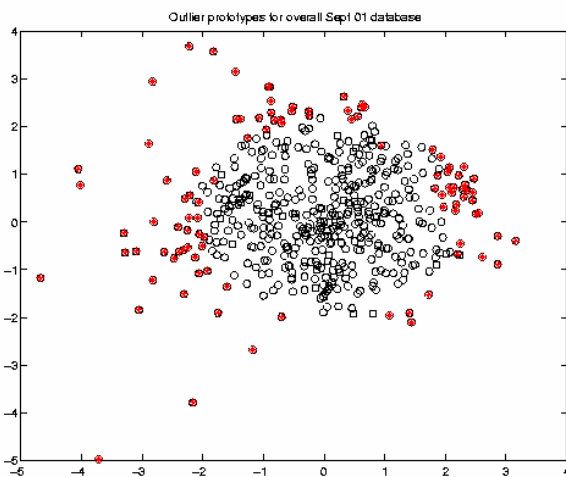


Figure 1 – 2-D visualisation of prototype patterns

It is clear from the 2-D map that there are outlier patterns, which is to be expected when the data is recorded from high-risk patients. The label of “normal” applied by the clinicians to the patients whose data was included in the training database is only a global label, mostly influenced by the eventual patient outcome. With many of these patients, there were often transient departures from normality during the 24 hours in which they were monitored. The visualisation shows these patterns (shown in darker colour) lying outside the “region of normality” at the centre of the map. These patterns are excluded from the Parzen Windows estimate of  $p(\mathbf{x})$ , so that  $N$  is closer to 400 than 500 in the model of normality.

## 7 Patient Status Index and generation of alerts

We define a “Patient Status Index” to quantify departures from normality (i.e. low values of  $p(\mathbf{x})$ ) so that alerts can be

generated when this index rises above a threshold value. The Patient Status Index is calculated by transforming the probability so that abnormality increases along the vertical axis (the horizontal axis being time in a trend plot):

$$\text{Patient Status Index} = \log\left(\frac{1}{\hat{p}(\mathbf{x})}\right) \quad (3)$$

The Patient Status alert threshold is set to indicate major changes in the vital signs, and trigger the possible intervention of a Medical Emergency Team. In order to choose an appropriate threshold, the variation of the Patient Status Index was investigated as a function of each of the five normalised parameters in turn. A BioSign™ parameter has no effect on the model (and hence on the Patient Status Index) when it is replaced by the mean value which it has in the training data set. Because of the zero-mean, unit variance transformation which is applied to each parameter, this is equivalent to setting the normalised parameter to a value of zero.


In order to study the effect of parameter 1 (heart rate) on the Patient Status Index, *normalised* parameters 2 to 5 (SDA blood pressure, arterial oxygen saturation, temperature and respiration rate) were set to zero, while the normalised heart rate was swept from  $-4$  to  $+4$ . The latter is equivalent to the heart rate being increased from a value of 13.9 bpm to a value of 173.7 bpm. The heart rate (HR) ramp is shown in the second plot (from the top) in Figure 2. The corresponding variation in Patient Status Index is the left-hand most parabola (data points 1 to 800) in the top plot.

In order to study the effect of parameter 2 (SDA blood pressure) on the Patient Status Index, *normalised* parameters 1 and 3 to 5 (heart rate, arterial oxygen saturation, temperature and respiration rate) were set to zero, while the normalised SDA BP was swept from  $-4$  to  $+4$ . The SDA blood pressure (SDA BP) ramp is shown in the third plot (from the top) in Figure 2. The corresponding variation in Patient Status Index is the second parabola from the left (data points 801 to 1600) in the top plot.

The effects of parameters 3 (arterial oxygen saturation), 4 (temperature) and 5 (respiration rate) were investigated in the same manner and the ramps are shown in the 4<sup>th</sup>, 5<sup>th</sup> and 6<sup>th</sup> plots from the top, respectively. The corresponding variations in Patient Status Index are represented by the third parabola from the left (data points 1601 to 2400), the fourth parabola from the left (data points 2401 to 3200) and the fifth parabola from the left (data points 3201 to 4000), all in the top plot. It should be noted that the ramp for arterial oxygen saturation (SpO<sub>2</sub>) is not increased beyond  $+1.52$  s.d. from the mean, as this corresponds to 100% oxygen saturation.

Deviations away from the mean cause the Patient Status Index to *increase* (as it is an index of patient *abnormality*), hence the parabolic shape of the Patient Status Index plots in Figure 2. The rate of increase away from the minimum is not

the same for each parameter, nor are the plots necessarily symmetrical. For example, a high temperature (pyrexia) has the largest effect – fourth parabola from the left – and hypotension has a greater effect than hypertension – second parabola from the left. Heart rate and respiration rate, on the other hand, give rise to symmetrical parabolae. This is the information learnt by the Parzen Windows model from the training data set.



**Figure 2** – Variation in Patient Status Index (top plot) as each of the five BioSign™ normalised parameters are ramped up from -4 to +4

The Patient Status Index alert threshold was chosen such that the threshold is crossed whenever an individual BioSign™ parameter is close to  $\pm 3$  standard deviations away from its mean value in the training data set (the other four being assumed to be normal). The nearest integer value for which this is the case is a value of 3 for the Patient Status Index, as shown by the dotted line in the top plot of Figure 2.

## 8 Clinical Trial of BioSign™ – design

The study design was a Randomised Controlled Trial (RCT) 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. Nested within the “monitoring” arm of the RCT was a within-patient assessment of the sensitivity and specificity of BioSign™ in detecting clinically significant events. The study was designed to answer two questions, “Does mandated

continuous physiological monitoring alter the number of clinically significant events detected?” and “Can BioSign™ detect these events with suitable sensitivity and specificity?” This trial design was chosen as it separates the effects of monitoring from the assessment of BioSign™.

Patients in the trial were randomly assigned to receive mandated continuous 5-parameter physiological monitoring for up to 72 hours post-surgery or post acute medical admission (Group A) or to receive usual ward care involving intermittent, single channel or multi-channel physiological monitoring at the attending staff’s discretion (Group B).

In Group A patients the multi-parameter patient monitor was connected to BioSign™ which continuously analysed the vital sign data. The ward staff were blinded to the Patient Status Index and the clinical event log. All occurrences of two types of clinically significant events were recorded: type 1 and type 2 events. Type 1 events were deemed to occur when a change in the patient’s condition triggered an action recorded in the patient record by clinical staff. Typical type 1 events are: unscheduled visit by medical staff requested by ward staff, arrest call or commencement of oxygen therapy. Type 2 events are those for which a clinically significant change in the patient’s physiological status occurs which is detectable from the recorded vital sign data. Type 1 events were recorded in both groups; by design only group A patients have records of type 2 events.

## 9 Clinical Trial of BioSign™ – preliminary results

A total of 405 patients undergoing high-risk surgery or following emergency admission for acute non-surgical conditions were studied at the John Radcliffe Hospital. Three patients were withdrawn after randomisation, and so there are 201 Patients in both Group A (mandated multi-parameter monitoring) and Group B (usual care). The outcome data so far collected on these patients is summarised in Table 3.

The data is currently being analysed retrospectively to determine:

- 1) the sensitivity and specificity of BioSign™ in detecting clinically recognised deteriorations (type 1 events) in the patient’s condition;
- 2) the time differences between detection of events by BioSign™ (type 2 events) and detection of the same events by the clinical team (type 1 events);
- 3) if the Patient Status Index has any ability to predict clinically significant changes in single or multiple physiological variables (type 2 events) before they happen;
- 4) if mandated standard 5-parameter continuous physiological monitoring alters the patient’s hospital mortality or length of stay.

We have carried out detailed analysis of one particular type 1 event (unscheduled admission to ICU) for Group A patients.

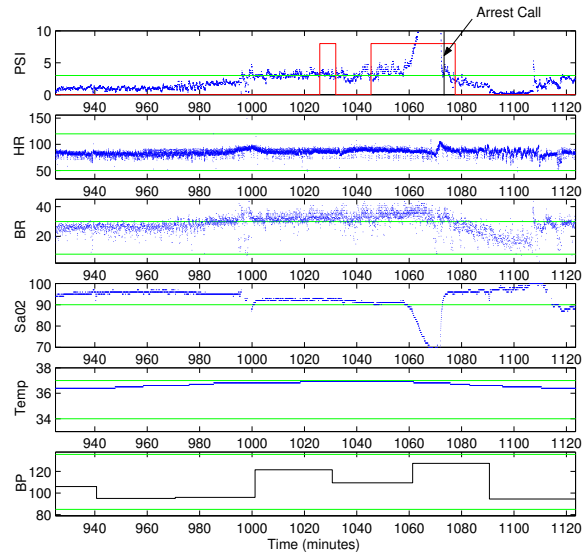
	All patients (n=402)	Group A (n=201)	Group B (n=201)
<b>Patients alive at 30 days</b>	333	169	164
<b>Acute hospital Length Of Stay</b>			
Mean		16	17
Median		10	11
Range		0-203	0-135
<b>Patients alive at hospital discharge</b>	315/396	156/197	159/199
<b>Total hospital Length Of Stay</b>			
Mean		21	19
Median		10	11
Range		0-299	1-158

**Table 3** – Data on outcomes of patients in RCT

There were ten occurrences of such an event in this patient group. Of these ten, three patients were admitted to ICU more than 24 hours after the end of the mandatory 72-hour monitoring period and they therefore cannot be included in the early warning analysis. Two of the remaining studies yield no Patient Status Index alerts, owing to missing data during the critical period prior to admission. With another study, the recording only started 75 minutes before ICU admission and the Patient Status Index was very close to, or above, the alert threshold throughout this period. The remaining four studies have a number of Patient Status Index alerts which appear to be predictive of ICU admission within the next 24 hours.

A 200-minute section (which includes an Arrest Call) from one of these studies is shown in Figure 3. At the beginning of the record, the value of the Patient Status Index is well below the alert threshold of 3, as most of the parameters are within their normal range. The heart rate (HR) and blood pressure (BP) remain normal throughout the period shown in Figure 3. The temperature gradually increases, however, until the patient becomes pyrexemic (skin temperature of 37 °C). More significantly, the breathing rate (BR) increases at a rapid rate, usually an indication of shallow breathing. This increase, together with the dip in SpO<sub>2</sub> at t = 1,000 min, causes the alert threshold to be crossed at this time, eventually causing an alert to occur just after t = 1,020 min. Beyond this, the breathing rate continues to rise and a further alert is generated soon after t = 1,040 min. This precedes the catastrophic decrease in arterial oxygen saturation starting at t = 1,060 min and continuing for ten minutes before the arrest call is made and clinical intervention occurs.

This sequence of events demonstrates that the BioSign™ alert would have caused the nursing staff to intervene at t = 1,025 min, approximately 50 minutes before they did so in this case.



**Figure 3** – Patient Status Index (top plot) and the five vital sign parameters for a Group A patient

## References

- [1] M.D. Buist, G.E. Moore, S.A. Bernard et al. “Effects of a medical emergency team on reduction of incidence of and mortality from unexpected cardiac arrest in hospital: preliminary study”, *BMJ*, **324**, pp. 387-390, (2002).
- [2] R.V. Jeffreys, J.J. Jones. “Avoidable factors contributing to the death of head injury patients in general hospitals in Mersey region”, *Lancet*, pp. 459-461, (1981)
- [3] P. McQuillan, S. Pilkington, A. Allan et al. “Confidential inquiry into quality of care before admission to intensive care”, *BMJ*, **316**, pp. 1853-8, (1998).
- [4] M.A. De Vita, R.S. Braithwaite, R. Mahidhara et al. “Use of medical emergency team responses to reduce cardiopulmonary arrests”, *Qual Saf Health Care*, **13**, pp. 251-4, (2004).
- [5] G. Garcea, S. Thomasset, L. McClelland et al. “Impact of a critical care outreach team on critical care readmissions and mortality”, *Acta Anaesthesiol Scand*, **48**, pp. 1096-1100 (2004).
- [6] J.T. Sharpley, J.C. Holden. “Introducing an early warning scoring system in a district general hospital”, *Nursing in Critical Care*, **9**, pp. 98-103, (2004)
- [7] C.L. Tsien, J.C. Fackler. “Poor prognosis for existing monitors in the intensive care unit”, *Crit Care Med*, **25**, pp. 614-619, (1997).
- [8] C.P. Subbe, M. Kruger, P. Rutherford, L. Gemmel. “Validation of a modified Early Warning Score in medical admissions”, *QJM*, **94**, pp. 521-526, (2001).
- [9] J.W. Sammon. “A nonlinear mapping for data structure analysis”, *IEEE Transactions on Computers*, **18**, pp. 401-409, (1969).
- [10] M. Tipping, D. Lowe. “Shadow Targets: A Novel Algorithm for Topographic Projections by Radial Basis Functions”, *Proc 5<sup>th</sup> IEE Int Conf Artif Neural Networks*, Cambridge, pp. 7-12 (1997).