Stell, A.J. and Sinnott, R.O. and Jiang, J. (2009) *A clinical grid infrastructure supporting adverse hypotensive event prediction.* In: Cappello, F. and Wang, C.L. and Buyya, R. (eds.) 9th IEEE/ACM International Symposium on Cluster Computing and the Grid CCGrid 2009, Shanghai, China, May 18-21, 2009. IEEE Computer Society, Los Alamitos, USA, pp. 508-513. ISBN 9781424439355

http://eprints.gla.ac.uk/7421/

Deposited on: 4 November 2009

A Clinical Grid Infrastructure Supporting Adverse Hypotensive Event Prediction

Anthony Stell, Richard Sinnott and Jipu Jiang
National e-Science Centre
University of Glasgow
Glasgow, UK
a.stell@nesc.gla.ac.uk

Abstract—The condition of hypotension – where a person's arterial blood pressure drops to an abnormally low level – is a common and potentially fatal occurrence in patients under intensive care. As medical interventions to treat such events are typically reactive and often aggressive, there would be great benefit in having a prediction system that can warn health-care professionals of an impending event and thereby allow them to provide non-invasive, preventative treatments. This paper describes the progress of the EU FP7 funded Avert-IT project, which is developing just such a system using Bayesian neural network learning technology based upon an integrated, real-time data grid infrastructure, which draws together heterogeneous data-sets from six clinical centres across Europe.

Index Terms—Hypotension, Data-Grids, Performance, Security

I. INTRODUCTION

IN many intensive care units (ICUs) the treatment of hypotension is a common but often unpredictable event. Most systems designed to alert clinicians to a patient's hypotensive status simply react to the drop in blood pressure. This in turn, means that the clinician must often administer relatively invasive measures to restore the pressure to a regular value. If a system were to predict such an event in advance over a feasible timescale (say, 15 minutes), this would alert a clinician to the imminent event and allow them to administer less invasive preventative measures. It has been estimated that the cost saving in both human and financial terms would be around 1600 euros per patient per day across the EU25 [1].

There are a variety of commercial systems available which allow the prediction of forthcoming clinical states. For example the BioSign device [2] produces an index predicting cardiovascular instability based on several vital signs such as heart rate, respiration rate, etc. The Philips

Medical Event Surveillance Monitor [3] supports the manual correlation of patient parameters into discrete "events". However, none of these systems provide any probabilistic measure of the causative information of the events occurring and therefore cannot be tied to the context.

Additionally there are a number of medical and clinical research groups that are attempting to enumerate the correlation of patient parameters such as blood pressure and heart rate to the onset of hypotension, however these are typically non-automated [4,5]. Therefore, a gap in the process of patient diagnosis has been identified: the automated prediction of hypotension in a patient. It is this gap that the Avert-IT project is attempting to address, by producing a predictive system, trained through an appropriate decision-support tool, on the data provided by six specialist centres from across Europe.

The technical aspects of this project break down into three distinct components:

- A "hypo-predict" engine trained on a Bayesian neural network to alert clinicians to imminent hypotensive events (to be eventually developed into a standalone "black-box" technology).
- A "hypo-net" application which federates the data from various neurological centres across Europe, and gives a much broader training set than could be achieved from one site – this is the primary motivation for setting up a data grid infrastructure.
- Supplementary web-based interfaces that allow input of neuro-physiological data not routinely input or collected electronically.

How these components fit together in diagrammatic form can be seen in Figure 1. The diagram covers the patient's admittance to the ITU on the left, their details and readings being entered into the Avert-IT database, through to being validated and communicated to the central project, based in NeSC on the right. The ClinicalAPI and ClinicalClient applications near the bottom, form the main components of the data federation process. The centralized uploading of data allows project-wide interpretations of data and prediction validation to be analysed, then fed back into the system later.
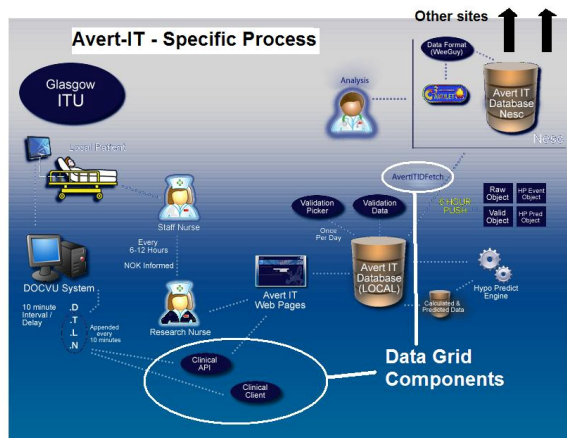
Figure 1: Avert-IT site specific process.

The focus of this paper is the grid-based, Hypo-Net federated data infrastructure. We describe the informational and performance challenges to the project, the security challenges inherent to this domain and how they have been implemented. We include an analysis of initial results integrating the various neuro/physiological applications from the AvertIT sites with a view to validating the performance of the Hypo-Predict engine and supporting individual site data-gathering, and the goal of eventually replacing manual input with an automated system. Once the data from the Hypo-Net application has been validated and analysed, the results are iteratively fed back into the refinement of the Hypo-Predict engine. Eventually the Hypo-Predict engine will be developed into a "black box" instrument to be made commercially available to the wider clinical community.

There are many e-Health projects in other specialties attempting similar integration and analysis of clinical data-sets – examples include aneurIST [17], Health-e-Child [18], CaBIG [19] and CLEF [20]. Projects also exist which provide data repository platforms for these e-Health projects (such as the UK Biobank [21]). The central premise of most grid-based projects is the desire to set up distributed platforms that have no centralized component (which thus removes single, critical points of failure). However, the primary goal of the Avert-IT project is to construct a grid that underpins the ability to infer information that is "greater than the sum of its parts". Given that, and the administrative layout of the consortium – two technical partners (one to federate data, one to develop the Bayesian network) and six clinical partners – and it is apparent that the technical infrastructure of a "hub and spoke" model is the most efficient way to proceed.

II. CLINICAL BACKGROUND

Before the necessary data from neurological centres can be processed to allow the prediction of adverse events, and since there is no prescriptive numerical definition of hypotension, it must be agreed between the consortium what actually constitutes a hypotensive event and subsequently how to identify the cause of that event.

A. Hypotensive event definition

There are various definitions of what constitutes a hypotensive event – largely based around the definition of the event threshold, the event hold-down, and the clear hold-down, in measurements of the patient's blood pressure. The various definitions used by the partner sites involved in Avert-IT are shown in Table 1:

TABLE I
AVERT-IT PARTNER SITE THRESHOLD DEFINITIONS

The first phase of the project has been to identify a

Centre	Measure & Threshold (mmHg)	Event Hold-Down (m)	Clear Hold-Down
Uppsala	BPs < 100	2	BPs > 100;5m
Glasgow	BPm < 70	5	BPm > 70;5m
Kaunas	BPs/BPd < 90/50	5	BPm > 70;5m
Heidelberg	CPP < 50	5	CPP > 60;5m
Monza	BPs < 90	5	BPs > 90;10m
Barcelona	BPs < 90	5	BPs > 90;15m

common definition that all sites can agree to use for the training of the Hypo-Predict engine. The definition agreed upon was that published in the Edinburgh University Secondary Insult Grade (EUSIG) paper (threshold of 90 mmHg for a hold-down of 5 minutes) [6]. The arrival at this definition is described in an earlier project paper which can be found at [7].

B. Identifying hypotensive event causes

There are many parameters that could potentially be used to predict the onset of a hypotensive event. These include: Heart rate (HRT); Respiration rate (RR); Mean blood pressure (BPm); Systolic blood pressure (BPs); Diastolic blood pressure (BPd); Blood oxygen concentration (SpO2); Temperature (T); Mean intracranial pressure (ICPm); Cerebral perfusion pressure (CPP).

The primary aim of the project is to discover how these factors can be incorporated into a prediction system, which has a high probability of accuracy whilst the numbers of false positive results (predictions of adverse events when no hypotensive event actually occurs) are minimized.

In the context of current clinical practice, it is understood that the development of such a predictive system, no matter how accurate, cannot completely replace clinical judgment. Rather, until extensive clinical trials are conducted to prove the accuracy and effectiveness of such a system, it can and should only be regarded as a decision support system for clinicians. Thus there is also scope to incorporate wider contextual and demographic information to ascertain a patient's likelihood of having a hypotensive event, e.g. factors such as their physical condition, associated medical treatments they might have been receiving for existing conditions etc.

509

A secondary finding of [7] was that due to the differing measurement procedures, it was common for significant proportion of shorter duration events to be missed [24, 25]. Also, some sites would only measure systolic, rather than mean, blood pressure (which can be calculated from systolic and diastolic using the formula: 4/3 BPs – 1/3 BPd). Therefore the results of this project have already contributed significantly to the evidence that mean blood pressure should have at least equal weighting in importance as systolic pressure when taking measurements.

C. Clinical Decision Support

To accurately predict adverse events, a decision-support module must be created, which takes account of the data retrieved so far, then makes a decision whether to send an alarm or not. There are a variety of methods of implementing a decision-support system, in addition to simple look-up tables and case-based reasoning, and these are listed below:

- Genetic Algorithms (GA) – this is a search technique that involves generating a random selection of solutions across the domain, then testing this against a fitness function. The fittest solutions are then carried over to the next generation of solutions until a pre-determined level of generations has been completed. Though useful in scheduling and timetabling optimizations, there is some doubt as to whether the GA would be applicable in this system.
- Bayesian Belief Networks (BBN) – this is a probabilistic model that represents a set of variables and their probabilistic dependencies. Inherently capturing probability in its development, training a BBN architecture with data-driven mechanisms has proved difficult.
- Artificial Neural Networks (ANN) – this is a computer-simulated model of the brain, where individual neurons are interconnected from an input layer, through one or more hidden layers, to an output layer.

Balancing the project requirements and the various advantages and disadvantages of the different approaches, it has been agreed by the project consortium that a Bayesian approach to training an Artificial Neural Network (a BANN) is the most effective way of training the system to detect the onset of a hypotensive event based on the input data of associated parameters. The motivation for considering the BANN is its application to the classification and modelling of highly nonlinear relationships whilst also considering probabilistic factors, expected to be a major aspect in the clinical inputs involved. Though the investigation of this area through comparison of the different methods would be of great benefit, given the limited time-scale of the project it was decided to choose one method most likely to yield useful results. It is upon this platform that the Hypo-Predict engine will be developed.

III. GRID APPLICATION AND INFRASTRUCTURE

A. Data Grid

The cost associated with any trial or study is most often made up of employing research staff in data collection and validation work. A significant intended outcome of this project is to reduce, or even eliminate completely, this cost by automating the data collection process.

In order to do this, it is necessary to leverage the Electronic Health Record (EHR) systems or clinical care systems more generally that are currently used by clinical sites. Despite various global efforts (HL7 [8], SNOMED-CT [9]), no single standard with regards to clinical electronic records has been widely adopted. As a result, it is not uncommon for many sites to have unique and idiosyncratic systems.

The data collected as part of the Avert-IT study has been based upon the core data schema developed through two previous EU grants by the BrainIT group [10], which captures a broad class of data that can be associated with brain trauma patients. However, because of the heterogeneity of the partner site EHR systems, the coverage of this core data schema is patchy with the subsets of the data requiring translation from the site IT systems.

After the first manual step (using a simple text editor GUI) of creating an XML file that maps local values to core BrainIT values, the Avert-IT Hypo-Net data grid automatically integrates the data from these sub-systems and captures as much of the local IT data sets that relate to the BrainIT core data-set as possible. Where automated data capture is infeasible, e.g. if only paper-based data sets exist at a given centre, manual input web-based software solutions are used to fill the gaps that remain. A follow-up aim of this project is to influence the development of EHR systems and identify how relevant parameters associated with brain trauma care can be measured, recorded and integrated into wider patient care.

Figure 2 shows the local Hypo-Net application process required to gather the BrainIT core data-set ("ClinicalClient" in figure 1). The program is packaged individually for each partner site, as the structure is largely similar in each case but with certain idiosyncratic interface points including:

- Parsing episodic (lab and physiological) data; Gathering demographic data (also used to identify the occupants of an ICU);
- A supporting ontology to interpret/map local data sets and models onto the BrainIT schema;
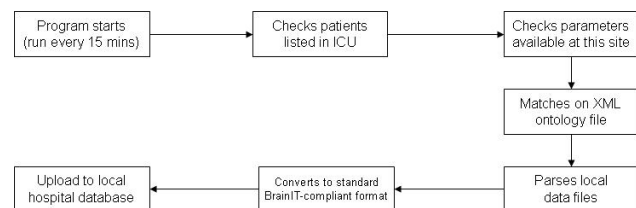
*Figure 2: process diagram for gathering episodic data
(mostly lab and physiological).*

The ClinicalAPI application in figure 1 follows a similar procedure (harvesting demographic information rather than physiological/lab data), but writes the information upon patient entry to a local text file rather than uploading it to the database. This is to allow supplementary demographic information to be manually input using the web-pages before finally committing the information to storage (part of the user interface section on the left of figure 1).

The final grid aspect is the aggregation of the full data-set with sets from other sites. This is done using web services across firewalls between the clinical site and NeSC (depicted in figure 1 as "AvertITIDFetch" – so-called because this is where the fully-qualified project-wide patient ID is assigned), and in regularly scheduled batch processes. The centralized data can then be analyzed and fed back into the Hypo-Predict engine development.

B. Performance

Another major aspect of the Avert-IT project is that of performance. Depending on the format of the output of the individual site EHR systems, and on the requirements of the Hypo-Predict engine, there is the possibility of significant high-volume data processing requirements (see next section).

Throughout the entire Avert-IT process, the following potential performance bottlenecks have been identified:

- The compilation of the hashmap containing the patient's physiological data (within ClinicalClient);
- Converting this hashmap to readable ASCII text;
- Uploading the relevant text to the database;
- The 6-hourly push of data to NeSC (shown in figure 1).

To mitigate against these performance concerns, the local Hypo-Net program has been coded to automatically switch between looped and threaded processing depending on data volume. Currently, for security considerations, the application is coded in Java. However, if as the project transpires, the need for optimal performance coding is required then a switch to C++ will be considered. If the volume of processing demands threaded processing there is also a requirement for the database to be able to support transactional updating over many connections (which it currently does).

Finally, the issue of security over performance is a required consideration. Throughout most of its operation, the program works locally, i.e. it is deployed within the clinical site itself. As such, protection is assumed to be provided by the site security policy. At the final collection to the centralized AvertIT database however, the data stream must cross firewall boundaries using web services and move from the clinical site to the central harnessing site

(depicted in Figure 1 as "NeSC"). To protect this data requires encryption at the message level using public-key cryptography which incurs a significant performance overhead. This data transfer however, can be operated in batch mode - i.e. real-time synchronous responses are not required.

IV. PROGRESS AND RESULTS

A. Data Grid

The six centres involved in the Avert-IT project have the following different EHR systems: Glasgow – Philips DocVu [11]; Monza – Draeger [12]; Uppsala – GE Quantitative Sentinel system [16]; Heidelberg – ADI [13] and CMA Microdialysis [14]; Barcelona – ADI; Kaunas – combination of in-house software systems;
Currently the Avert-IT Hypo-Net application has been targeted to the Glasgow, Monza and Uppsala data centres (with interfaces to other centres under development).

The Philips DocVu system collects a variety of patient information and records it in file-based ASCII text format (the example used in figure 1). The files that are generated include files post-fixed with .D (demographics), .T (physiological trends), .L (lab and treatment data) and .N (extra notes). The format of the .T file provides minute-by-minute interval recording (though this time interval is configurable). The files are updated from the bedside systems in the ICU every night in a batch process. Occupancy of the ICU is denoted by the presence of a .D file in the appropriate folder. The file-naming convention is to prefix the extension with the hospital number assigned to the patient (e.g. SGH0000.D). An important feature of the .T files are the overlapping channels – these data points are recorded in case certain parameters are for some reason not recorded. This has led to the use of hashmaps for individual patient data storage, (c.f. the use of array storage objects, which would significantly increase search time and reduce performance).

The Draeger system similarly records minute-by-minute patient information (this is again configurable), is updated nightly with ICU occupancy details, and outputs the results to ASCII text format using .hl7 files. Overall, the system appears to have more comprehensive coverage of possible BrainIT patient parameters than DocVu. The directory structure is also slightly different – files indexed on time sequences rather than patient ID – and requires a different parsing algorithm.

The systems in use in Uppsala have interfaces to the physiological data, which are immediately available due to previous work conducted locally to harness that data. Access to the wider set of lab and demographic data is currently under negotiation.

These three systems provide the most viable available option to create a data-grid using the core BrainIT data schema, which in turn minimizes the requirements for data entry by project research nurses. Though it is estimated that the coverage of BrainIT core schema by centres is low

(approximately ~33% for each centre), the parameters involved are the most important for development.

B. Performance

The performance of the application has been tested in only one area so far: the rendering of the bedside output data to a readable (and calculable) string format. The other main area, which is due for performance testing, is in the transmission of data across web services, and the associated encryption overhead. For the conversion, the XML example of the output from Heidelberg was used as this had been configured to give second-by-second values over large time intervals, covering several hours.

There are two variables to control the process of threading: a switch to determine whether to loop or thread, and a divisor to calculate the optimum number of threads. For the machine setup used, experimentation of record numbers found the equality point to be around $n = 64$ (where n is the number of records), above which it was advantageous to thread the process rather than loop. The divisor has been calculated to be the nearest integer to $sqrt(n)$, which assumes a regular underlying Gaussian distribution of processing speed versus divisor*thread product. (The author notes that this is an intuitive result based on Gaussian theory – quantitative results to support the claim are currently being investigated). Similar test results are expected to follow when processing the WS performance.

C. Security

The main security issue in the Avert-IT project is one of protecting patient identity, and keeping all data anonymised as much as possible. As the bulk of the application is to be run within the firewalled domain of the various hospitals, this is not a major issue during development and is subject to the individual hospital site policies. The international nature of the project means that each site must have ethical approval granted to recruit patients to the study, and have their data transmitted to the central site. The central site holding this data is bound by the UK Data Protection and Freedom of Information acts, as required by the legal entity housing the database (in this case the University of Glasgow), which has equivalent legislation in all the participant countries.

In order to keep the patient identity anonymised once the data has left the hospital domain, there is a strict policy enforcement of not having any identifier used within the hospital within the centralized database. This has lead to the implementation of three identifiers of a patient, representing their status in the local hospital, the hospital Avert-IT database, and the centralized Avert-IT database. This latter identifier is assigned as a foreign key during the batch processing and uploading of site results to the centralized database, and can therefore never be traced to an individual patient (unless the hospital security is also breached).

Authentication occurs between the local clinical database and the NeSC service using digital certificates. Advanced authentication processes such as single sign-on and credential delegation are unnecessary as the grid infrastructure is a back-end batch process with a low number of well-defined communication channels (authentication of the hospital interface is managed with a standard database access control policy). This largely negates the need for grid security stacks (such as PERMIS [23] or Shibboleth [22]), which also come with a heavy overhead of supporting requirements.

Encryption of the data in transit over https between sites is performed using public/private keypairs (for performance reasons this uses symmetric algorithms (e.g. DES) to protect the cipher-text, and asymmetric algorithms (e.g. RSA) to protect the symmetric key), issued from a certificate authority set up ahead of run-time, and communicated between parties using secure, out-of-band channels. Communications are signed using these key-pairs as well (see previous paragraph on client authentication).

Implementation of this encrypted security using Java 6 is relatively straight-forward and configurable, using the Java WS environment Glassfish [15]. The server used to house the individual web services is the Sun Application Server (v9.1).

As with other software and server packages, all reasonable administrative precautions are taken within the server and machine environments, including but not limited to: Easily-configurable firewalls implemented using IPtables; SSH protected to specific IP addresses, users and groups; Administrative machine accounts (and others) protected by strong passwords; Stealthing capabilities of all machines enabled; Regular virus sweeps; Real-time port monitoring; All other software and servers locked down or switched off; Databases protected with strong passwords; All third-party software is regularly patched and updated.

In terms of programming security, the principles of default-deny exception-handling and input bounds checking are adhered to throughout the code.

In the event of any component failure of the Hypo-Net application, a redundant failover is present to seamlessly take over the duties of the first. The application is therefore robust to failure, and protects the operation with all the same security measures outlined in this section.

IP addresses are stealthed using IPtables, which allows close configuration of what machines are allowed to access the services, locks down all other ports on the machines, and enables logging of any attempted intrusions. The identification of individual IP addresses also allows the Denial of Service (DOS attack) vector to be greatly minimised. The "defence in depth" philosophy is manifest in the fact that clients also need to be authenticated to the individual databases used by the application, as well as through the encrypted Web service connection. It should be noted here that the application never requires any participating clinical sites to open their firewalls to incoming connections. Because the clinical site initiates the conversation, they are able to converse with the server without having to leave any ports open – in exactly the

512

same way as users surfing the Internet from behind a firewall.

V. CONCLUSION

This paper has shown the development of a data grid infrastructure to support adverse hypotensive event prediction. The ultimate aim of the Avert-IT project is to produce a "black box" that can be integrated with various healthcare systems and support the prediction of hypotensive events. In order to do this, several challenges of security, performance and data formatting must be overcome. The authors have presented possible solutions to these issues.

The major results pertaining to the clinical context of the project have provided evidence that measurement of the mean blood pressure is of much greater significance than had previously been thought in the clinical community (more information on these results can be found in [7]).

Work in the immediate future is to provide much greater performance results data and use this to optimize the application as much as possible. In parallel with this, the development of the data grid for the rest of the participating centres will be completed in time for the clinical trial launch of the software. As this work continues, the scope of the Hypo-Predict engine – investigating how many parameters and of what significance they are – will be made more concrete, which will feed directly into performance considerations. The project is on-going with large scale clinical trials of the developed technologies scheduled for June 2009.

REFERENCES

[1] Avert-IT project proposal Advanced Arterial Hypotension Adverse Event prediction through a Novel Bayesian Neural Network (available at http://wiki.avert-it.org/wordpress)

[2] Tarassenko et al., BioSign: multi-parameter monitoring for early warning of patient deterioration. 3rd IEE International Seminar on Medical Applications of Signal Processing (2005/11199), p71-76

[3] Philips Medical Event Surveillance Monitor: http://www.medical.philips.com/main/\\products/patient-monitoring/assets/docs/event-surv4522-982-91371.pdf

[4] Geoffrey Manley, M. Margaret Knudson, Diane Morabito, Susan Damron, Vanessa Erickson, and Lawrence Pitts. Hypotension, hypoxia, and head injury. ARCH SURG, 136:1118–1123, 2001.

[5] A. Marmarou, R.L. Anderson, J.D. Ward, et al. Impact of ICP instability and hypotension on outcome in patients with severe head trauma. Journal of Neurosurgery, 75:S59–S66, 1991.

[6] Patricia Jones, Peter Andrews, Susan Midgley, Shirley Anderson, Ian Piper, Janis Tocher, Alma Housley, Jane Corrie, James Slattery, Mark Dearden, and Douglas Miller. Measuring the burden of secondary insults in head-injured patients during intensive care. Journal of Neurosurgical Anesthesiology, 6:4-14, 1994.

[7] Stell, Sinnott, Jiang et al, Federating Distributed Clinical Data for the prediction of adverse hypotensive events, e-Science All-Hands conference 2008

[8] HL7: http://www.hl7.org

[9] SNOMED-CT: http://www.connectingforhealth.nhs.uk/systemsandservices/data/snomed

[10] BrainIT group: http://www.brainit.org

[11] Philips DocVu: http://www.medical.philips.com/main/products/patient-monitoring/products/doc-center/

[12] Draeger Medical: http://www.draeger.com

[13] ADI: http://www.adinstruments.com

[14] CMA Microdialysis, ICU Pilot: http://www.icupilot.com

[15] Java Glassfish: http://glassfish.dev.java.net

[16] GE Healthcare – http://www.gehealthcare.com

[17] Aneurist – http://www.aneurist.org

[18] Health-e-Child – http://www.health-e-child.org

[19] Kakazu, Cheung, Lynne - The Cancer Biomedical Informatics Grid (caBIG): pioneering an expansive network of information and tools for collaborative cancer research, Hawaii Med Journal, 2004 Sep;63(9):273-5

[20] Taweel, Rector, Kalra, Rogers, et al - CLEF – joining up healthcare with clinical and post-genomic research, Current Perspectives in Healthcare Computing, 2004

[21] UK Biobank – http://www.ukbiobank.ac.uk

[22] Shibboleth – http://www.internet2.edu/shibboleth

[23] Chadwick, Otenko - The PERMIS X.509 role based privilege management infrastructure, Future Generation Computer Systems, Vol. 19, issue 2, Feb 2003, p277-289

[24] Zanier E, Ortolano F, Ghisoni L, Colombo A, Losappo S, Stochetti N. Intracranial pressure monitoring in intensive care: clinical advantages of computerized monitoring over manual recording. Crit. Care Med. 2007; 11:117

[25] Corrie J, Piper I, Housely A, Tocher J, Anderson S, Midgley S, Slattery J, Dearden N, Miller J. Microcomputer based data recording improves identification of secondary insults in head injured patients. British Journal of Intensive Care, June 1993. 226-233.