

A Distributed Clinical Data Platform for Physiological Studies in the Brain Trauma Domain

Anthony Stell¹, Richard Sinnott², Rob Donald³, Iain Chambers⁵, Giuseppe Citerio⁷, Per Enblad⁸, Barbara Gregson⁹, Tim Howells⁸, Karl Kiening¹⁰, Pelle Nilsson⁸, Arminas Ragauskas¹¹, Juan Sahuquillo⁶ and Ian Piper⁴ (on behalf of the Brain-IT Group)

¹National e-Science Centre
University of Glasgow
Glasgow, UK
a.stell@nesc.gla.ac.uk

²Melbourne eResearch Group
University of Melbourne
Melbourne, Australia

³Dept of Statistics
University of Glasgow
Glasgow, UK

⁴Dept of Clinical Physics
Southern General Hospital
Glasgow, UK

⁵Dept of Medical Physics
The James Cook University Hospital
Middlesbrough, UK

⁶Neurosurgery
Vall d'Hebron University Hospital
Barcelona, Spain

⁷Neuroranimazione
Hospital San Gerardo
Monza, Italy

⁸Dept of Neurosurgery
Uppsala University Hospital
Uppsala, Sweden

⁹Dept of Neurosurgery
Newcastle General Hospital
Newcastle, UK

¹⁰Neurosurgery
Ruprecht-Karls-Universität Hospital
Heidelberg, Germany

¹¹Telematics Sc. Lab
Kaunas University of Technology
Kaunas, Lithuania

Abstract— There are many serious and acute physiological conditions about which we have incomplete medical knowledge that can support optimal healthcare intervention. To develop effective treatments a wealth of clinical data is required for collection, analysis and feedback. Such data often does exist but is typically held in a variety of different formats and locations. This paper describes the EU FP7-funded Avert-IT project (www.avert-it.org), which has developed an integrated, real-time physiological data grid infrastructure (HypoNet) to address the specific issue of prediction of hypotensive events in the brain trauma domain and is currently being used as part of a large multi-centre clinical trial. The implementation and application of the HypoNet system is described here.

Keywords- Hypotension, Clinical Data-Grids, Security

I. INTRODUCTION

Hypotension is abnormally low blood pressure (BP). If left untreated it can be particularly harmful in unconscious head injured patients who have impaired ability to regulate cerebral blood flow and can impact on long-term outcome. The treatment of hypotension (a currently unpredictable event) in many intensive care units is common. Most physiological monitoring systems are designed to provide notification

“once” a hypotensive status has been detected – through measurement of the blood pressure – i.e. once the event is already underway. This in turn, means that the clinician must intervene clinically often through administration of vasoactive or cardiogenic drugs. If a system were to predict such an event in advance over a feasible timescale (e.g. 15 - 30 minutes), this would alert a clinician to the imminent event and allow them to commence a more moderate intervention. 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].

Hypotension is an example of an abnormal physiological condition, but there are several others that can be identified using similar bedside monitored parameters (BP, heart rate (HRT), etc), such as hypertension, raised intracranial pressure or Tachycardia [2].

As such, a clinical data platform would thus have great use as a research tool for a wide range of studies. In order to collect, analyse and feed back the results of this data in a meaningful timescale, the AvertIT HypoNet platform was developed. This system interfaced with the local clinical data gathering systems, standardised the data into a uniform format

then streamed this data into a central repository. As well as the data-gathering component, a key challenge has been to use these data sets for prediction of upcoming adverse hypotensive events. The AvertIT project was structured accordingly and has three distinct components:

- A “hypo-predict” engine trained on a Bayesian neural network to alert clinicians to imminent hypotensive events;
- A “hypo-net” data grid application that federates the data from various neurosurgical centres across Europe, and gives a much broader training data than could be achieved from one site;
- Supplementary web-based software solutions that allow input of patient management data not routinely input or collected electronically, e.g. procedures periodically administered by nurses.

How these components fit together in diagrammatic form can be seen in Figure 1.

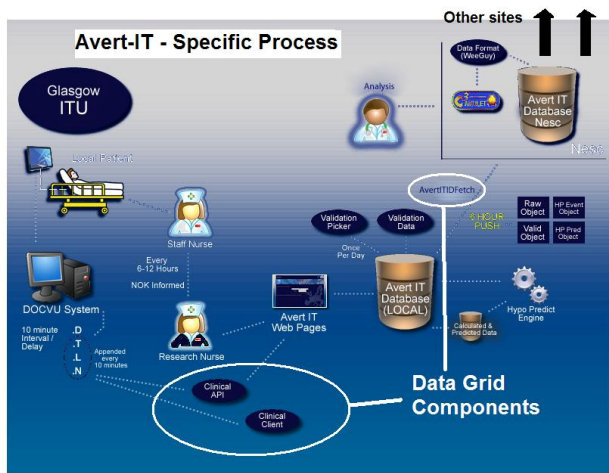


Figure 1: AvertIT Workflow Schematic

The main focus of this paper is the Hypo-Net federated data infrastructure. We provide a full description of its operation, implementation, security, support and maintenance. We describe the results provided by the system in terms of patient data, events detected and quality control feedback to the centres involved. We also describe some of the operational features of the system as a whole – such as test metrics available, email support and automated report generation. Full details of the Bayesian Neural Network results and testing will be described in a later paper.

II. CLINICAL BACKGROUND

Before the necessary data from neurological centres can be processed to allow the prediction of adverse events, it had to be agreed between the consortia what actually constitutes a hypotensive event, as there is no universally accepted definition of hypotension in the medical literature.

A. Hypotensive event definition

There are various definitions of what constitutes a hypotensive event – largely based around the definition of the event threshold (when the event starts), the event hold-down (how long the event lasts), and the clear hold-down (when the event is over), 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 (BPm = mean blood pressure, BPs = systolic blood pressure, CPP = cerebral perfusion pressure).

TABLE I
AVERT-IT PARTNER SITE THRESHOLD DEFINITIONS

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

The first phase of the project was to identify a common definition that all sites agree upon 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 systolic or 70 mmHg mean arterial pressure, for a hold-down of 5 minutes) [2]. The justification for this definition is described in an earlier project paper [3].

B. Identifying hypotensive event causes

There are many parameters that could potentially be used to interpret the onset of a hypotensive event that are commonly monitored in an intensive care unit. These include:

- Heart rate (HRT)
- End tidal CO2 (ETCO2)
- Mean blood pressure (BPm)
- Systolic blood pressure (BPs)
- Diastolic blood pressure (BPD)
- Blood oxygen concentration (SaO2)
- Temperature (TC)

The primary aim of the project was to discover how these parameters could be incorporated into a prediction system that would have a high probability of accuracy (and a clinically relevant degree of specificity) whilst minimising the numbers

of false positive results (false positives are predictions of adverse events when no hypotensive event actually occurs).

It is noted that there is currently no solution to this hypotensive prediction problem on the market and the commercialization aspects of the project are currently being considered. That is, a variety of commercial systems do allow the prediction of forthcoming clinical states, but none are able to provide any probabilistic measure of the causative information that can be tied to context. For example the BioSign device [5] produces an index predicting cardiovascular instability based on several vital signs such as heart rate, respiration rate, etc. The Philips Medical Event Surveillance Monitor [6] also supports the manual correlation of patient parameters into discrete “events”.

Additionally, there are a number of medical and clinical research groups that are attempting to enumerate the relationship between patient parameters such as blood pressure and heart rate to the onset of hypotension, however these are typically non-automated [7,8]. 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 neurosurgical centres from across Europe.

In the context of current clinical practice, it is accepted that the development of such a predictive system, no matter how accurate, cannot completely replace clinical judgment. Rather 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.

A completed analysis has shown that, due to the differing measurement procedures, it was common for hypotensive events to be missed [19, 20]. Our further analyses [4] have shown that this is due in part to when hypotensive events are detected based solely upon systolic BP without considering mean BP mean. 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 measuring.

C. Clinical Decision Support

Though not the main focus of this paper, it is instructive to have a section describing the theory behind the ultimate aim of the project: namely, to develop a predictive engine.

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 raise 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:

- Genetic Algorithms (GA) – provide a search technique that involves generating a random selection of solutions across the domain, then testing each of them 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) – offer 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 proven difficult.
- Artificial Neural Networks (ANN) – comprise a computer-simulated approach derived from the brain model, 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 an effective way of training the system to detect the onset of a hypotensive event based on the input data from associated parameters. The motivation for considering the BANN is its efficiency for the classification and modelling of highly nonlinear relationships whilst also considering probabilistic factors, which themselves are expected to be a major aspect in the clinical inputs involved. Though the comparison of the different methods would be of great benefit, given the limited time-scale of the project it was decided to focus upon one method (BANN) most likely to yield useful results. It is upon this platform that the Hypo-Predict engine has been developed.

D. Physiological, treatment and demographic data

The ability to harness the wealth of clinical data available in the neurointensive care domain is a pre-requisite for the development of a probabilistic Bayesian neural network solution. However, there are wider ramifications for the construction of such a research platform. This project focuses on the traumatic brain injury (TBI) domain, but the data collected and methods applied could be applicable in other pathophysiological populations. The different types of data collected broadly fall into three categories:

- Physiological – this provides the most detailed minute-by-minute information on a patient’s status. Blood pressure, heart rate, intracranial pressure can be measured and directly related to events. The step size of such feeds must be varied to allow the discretization

of individual hypotensive “events”. However, in general the more detailed the resolution of the physiological data that can be retrieved, the better the analysis. This detail naturally depends on the granularity of the hospital system with which the platform is integrated (currently Philips DocVu, Draeger Medical, Datex-Ohmeda, ADI Instruments and the Odin Browser). In figure 3, the need to distinguish between an “event” and an “episode” will be outlined and discussed.

- Treatment (or Episodic) data – this is data that is not available at regular intervals but still has an impact on the status of the patient at any given time. The most obvious example is the administration of any kind of drug, perhaps in a bid to restore a patient’s vital signs to normal after the onset of a hypotensive event. Whilst irregular, this is information that has a direct impact on the explanation of the presence (or absence) of an event in a given timeline.
- Demographic data – this is data only entered once per patients (eg: age, sex, centre, type of injury) and is important in providing wider contextual information about the status of a patient. The most likely use of this data would be in secondary studies on the incidence of hypotensive events connected to information such as, say, lifestyle habits or geographic area. The information is still valuable in being held in the same repository, though as the nature of this data can be identifying, the motivation for secure storage and linkage is paramount.

Whilst these definitions are not yet fully standardized, much progress has been made by the Brain-IT consortium (www.brainit.org) to settle on best practice standards in the neurointensive care domain. The Avert-IT project has adopted this data model directly. It comprises a standard ontological description that can be shown in properly marked up XML files with textual descriptions of the terms and their usage in traumatic brain injury management.

III. HYPO-NET

The HypoNet system has three main components: ClinicalAPI, ClinicalClient, and AvertITIDFetch. They interact and operate as outlined in Figure 1. The typical scenario for their interoperation is as follows:

- The patient enters the hospital system as per normal procedure for that institution. A local identifier is assigned to the patient, and their next-of-kin are informed. Monitoring data is now in principle available to be parsed from the local hospital system (should the patient be entered into the study).
- The patient is identified as being a potential candidate for the Avert-IT study. A consent form is generated and signature obtained from family.
- Upon successful receipt of consent, the local Avert-IT research nurse enters the patient ID into the study using either the web application for manual entry, or the

ClinicalAPI module of the HypoNet application for point and click selection entry of patients currently in the ICU.

- This ID is assigned a study number and an entry is created in a specific list file (*avertitpatients.txt*).
- The data from the hospital system (e.g. Philips DocVu, Draeger, etc) is now parsed by the HypoNet system, rendered into a standard SQL format and input to a local Avert-IT database. The developed application runs standard classes for each centre, the only difference being the parsing of output. Standard properties can be input from *avertit.properties* (such as file paths, centre numbers and data volume thresholds) and the BrainIT description in *ontology.xml* is used to map local parameter labels to the BrainIT standard.
- The HypoPredict system runs as a service in the background and does real-time analysis on the data in the local database. The results of this analysis are fed back into separate tables within the local database.
- Every hour, the AvertITIDFetch module contacts the central web service located at the National e-Science Centre (NeSC) at the University of Glasgow and uploads the latest anonymised data to the main repository database. Along with the data, logfiles and throttling mechanisms, system status availability features are also implemented to facilitate support, maintenance and system reliability.

A. Implementation

The system is implemented with the following specifications. In each clinical centre hardware has been provisioned including:

- Sun Fire X4140 server to house the local database, HypoNet and HypoPredict applications (64Gb memory, 2x 2.7GHz quad-core AMD Opteron processors, 8x 2.5” SAS drives)
- Dell E4500 Laptop for bedside data input
- Wireless router for connectivity within the ICU
- Operating systems of Windows Vista and Windows Server 2008. (Most advanced Windows OS’s at time of development allowing use of advanced security features and in line with prevailing hospital IT policies)

The central repository, housed at the University of Glasgow, is written using Java v6, Glassfish WS v3 and WSIT encryption [9]. It is a similar setup of a Sun Fire X4140 running Windows Server 2008, with the important addition of being run on a virtual machine that allows on-demand access to much larger storage facilities provided by Sun Thumper machines also housed locally (providing 24 – 48 Tb of storage).

B. Individual Clinical System Features

Mainly covering the monitoring systems and network infrastructures, each centre had individual specifications and issues with regard to their clinical IT infrastructure, of which it is pertinent to list here.

- Glasgow – the Southern General Hospital uses Philips DocVu [10] monitoring systems, which report output to a central server, then export this to ASCII text files. The main complication in parsing was the presence of overlapping channels in the text output, requiring the use of hashmap rather than array storage of values. The only network complication was the need to contact the web service in outgoing traffic over port 80, as restricted by the UK NHS Net.
- Uppsala – the monitoring system used here was the Odin browser [11], created by one of the Avert-IT lead developers, Tim Howells. Output was again to pre-formatted ASCII text (previously used by the BrainIT consortium). There were no complications to the network infrastructure here.
- Monza – HL7 Infinity Gateway [12] is used at the San Gerardo Hospital, which broadcasts to a specific IP address messages in HL7 v3. The parsing of this output was made possible using the HAPI [13] interface, available in Java. Through threaded blocking queues, it was possible to get a stream of minute-by-minute data archived then parsed by ClinicalClient. The network infrastructure was complicated by the presence of a sophisticated proxy serving the hospital network. The solution to this was to set the IP address outwith the proxied network (or to “proxify” it) using a proxy application, Proxifier [14].
- Heidelberg – similar to Monza, Heidelberg used an HL7 gateway, broadcasting to a specific IP address and parsed using the HAPI interface. In this centre, the HL7 broadcast had to be split due to the presence of another project that also required the same stream. Using applications developed in Java, the stream was repeated then broadcast to the other machine, then archived for parsing by ClinicalClient. To send the information from the network to the central server required using NTLM authentication [15] for a pre-set account within the hospital.
- Barcelona – ADI Instruments LabChart [16] was the monitoring system used in this centre. The application runs on Windows and outputs an Excel worksheet with a stream of patient variables. Again, due to the presence of a separate project, two sets of configuration files were required for ADI input. The latest version (v7) required a Visual Basic wrapper to allow the conversion to versions compatible with all Office applications. Once converted, the output comma-separated value file is parsed by ClinicalClient. The network

was proxied using a transparent proxy, therefore only requiring the setting of “http.proxy” variables in the Java code.

- Vilnius – this hospital centre used individual bedside monitors running Datex-Ohmeda [17] systems. Using a wrapper C program to interpret the broadcast signal, output was sent to ASCII text, transferred to the central server then parsed by ClinicalClient. The main challenge in network infrastructure here was the lack of VPN, requiring remote support for testing and maintenance.

It was found that having virtual support greatly increased the efficiency of work and decreased the time spent debugging issues. To this end, four of the six centres implemented VPNs and remote accounts to allow the developers to login and provide solutions. The two that didn’t have this facility (in both cases because of administrative rather than technical obstacles), on-site support was provided by the project employee, guided remotely by the main project developers.

C. Support and maintenance

The system is supported with a variety of features:

- Logfile uploads – alongside the data upload from each centre on the hour, a set of anonymised logfiles monitoring the output and status of each system component is uploaded to the central server in Glasgow.
- Regular automated email updates – at 8am every morning an email is sent to the main Avert-IT developers, giving a summary status of the main components from each centre. Regular expression parsing of the logfiles picks up time gaps and error messages.
- Trac server – short- and medium-term issues, bug-fixes and issues requiring investigation are monitored and updated using a Trac service housed at Glasgow University [18].
- Master-to-master database synchronization – two databases are run to mirror each and allow a robust service to be provided. To prevent inconsistent data updating at the central server in the event of one failing, both databases are synchronized at regular intervals.
- WS redundant failover – in a similar vein, the web service gateway to the central repository has two redundant failovers to provide uninterrupted service.
- Data volume throttling – the nature of the data involved means that security takes priority over performance. However, this means that the volume of data over the web services (WS) is greatly increased due to encryption and XML padding of data, and also due to the data being obfuscated away from direct SQL ASCII text. To avoid a bottleneck of data blocking the WS, a throttling mechanism has been added to the AvertITIDFetch module, which allows the volume of

data being uploaded to the central repository to be controlled by the individual centre. The control lies on the hospital side as one of the main challenges is to deal with the backlog of data records accumulated from the hospital system in the time window of a patient’s arrival at the ICU to inclusion in the AvertIT study. Once this backlog has caught up the real-time data streaming is scheduled accordingly (on the hour with predictable data sizes). We note that automatic load-balancing is a feature that would be desirable to add here if the project timescale would allow.

D. Security

As mentioned previously the nature of the data involved means that security must take precedence over performance in situations where this contention arises. The following security features are present within the Avert-IT data grid:

- Encryption of data streams using WSIT technology and authentication allowing the assertion of user authority.
- Advanced firewall configuration to allow selective streaming of data (with default settings of “closed” only opened if deemed absolutely necessary – e.g. to read an HL7 broadcast within a hospital network)
- WS access through port 80 – the only port available to all centres for communication from their individual networks (standard port for http web traffic).
- Administrative security – all applications are locked down according to typical system administration procedures (Windows, SQL Server, etc).
- Programmatic security – during the development of code used, typical security measures were adhered to, such as the isolation of classes, and protection of dropping byte-code or SQL injection attacks within the web service itself.

E. Reporting

Several tools have also been developed that support a combination of SQL to the central repository database, R statistics application, Latex report rendering and Adobe PDF reader to provide automatic report generation. These have allowed the rapid creation and dissemination of results – shown in the next section – to be analysed and returned by the clinical community in collaboration with the project.

IV. RESULTS

The scientific results relating to the development of the HypoPredict engine are to follow in a separate paper. The results detailed here are to provide supporting evidence for the technical efficacy of the clinical data platform that has been developed throughout the project.

A. Patient Information

After 2 months of full operation the following numbers of patients have been collected from the participating centres

(note that this is a full collection – patients in the Phase I clinical trial are a subset of these numbers):

Centre	No. of patients
Glasgow	15
Uppsala	9
Monza	14
Heidelberg	9
Barcelona	10
Vilnius	1

Table 2: No of patients in central repository

As seen, all sites are now using the HypoNet system for collection of clinical patient information used in the AvertIT clinical trial itself. Ad hoc investigation of the anonymised aggregation of the stored patient data can be displayed graphically using the Amulet tool [19]. However, for dissemination amongst the clinical partners a project specific report was developed showing vital sign traces of individual patients along with the HypoPredict trace of predicted events. Figures 2 and 3 show patients within the data-set that have different representative features.

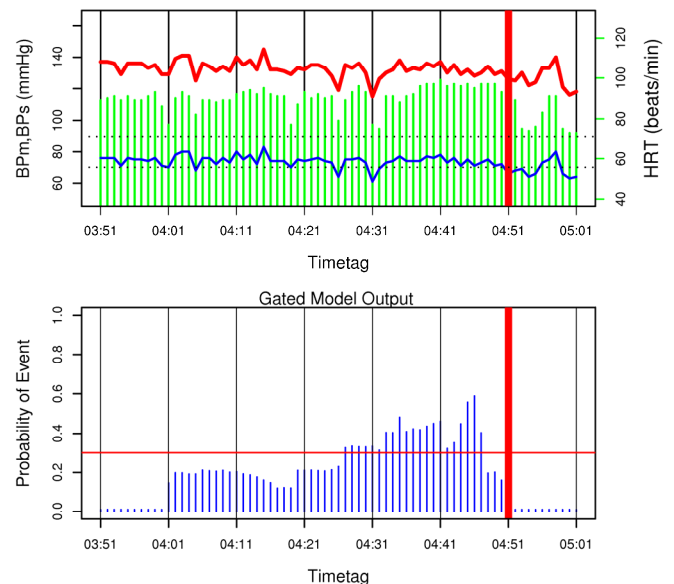


Figure 2: the top graph shows vital signs (BPs – red, Bpm – blue) ahead of an hypotensive event (vertical red bar). The bottom graph shows the increasing probability of an event occurring before the actual event.

Figure 2 shows a patient with a clearly defined hypotensive “event”. Shown in three bands in the top graph – mean BP (blue), systolic BP (red) and heart rate (green bars). The thick vertical red line indicates the time of an event. The bottom graph shows - with blue bars denoting the probability of an event – that this probability steadily increases crossing the 0.3 threshold within 20 minutes of the event occurring. In this case, the patient showed no sign in any of the vital signs charted of

an upcoming event. However, as a result of the BANN-trained probability calculation, the gated model output shows a distinct increase in the 20 minutes before the event. The ramifications for the HypoPredict engine is that with the usage of the HP application, this is an event that would be picked up and would alert clinicians to the approach of a hypotensive event ahead of time. Figure 3 shows a patient having multiple events triggered by a single event. This type of output has been regularly observed and has led to the requirement of a definition of a “patient episode” rather than a single “event”, as to attempt to extract single events reduces the clinically practical information of what is observed.

The other type of event of note is a comparison between two exactly equal lead conditions, which resulted in different outcomes. Upon correlation with the treatment data also collected, it was discovered that the bedside nurse registered the possibility of occurrence of the second hypotensive event and administered a preventative treatment. Therefore, the prediction output of the HypoPredict engine is accurate, but because of the successful treatment, a subsequent hypotensive event did not occur.




Figure 3: the three graphs show many events clustered together. To chart every event would reduce the clinical practicality of these readouts, therefore a new definition of “episode” is used.

It should be noted that these graphs merely illustrate the type of measurement ongoing and the possible practical applications. The scientific output and the more philosophical debates of making causal links with previously observed data are beyond the scope of this paper.

B. Availability

This is real data obtained from patients undergoing treatment for brain trauma injuries in an ICU in participating centres. The data covers the patient stay from the start of monitoring in ICU to discharge or death. Within one hour of entry to the study, the fully parsed, anonymised physiological data is stored in the central repository and can be made available for clinical study with a variety of tools.

To participate and/or contribute to the centralized repository, a simple requirement and procedure document has been put together. It is summarized below:

- Purchase of hardware and software required for local gathering of data. Typically, this includes a server, laptop, wireless router and operating system (the final two may not be necessary depending on local hospital resources).
- Outline of which patient monitoring system is being used at the participating centre. Currently, interfaces have been developed for Philips DocVu, Draeger Medical, Datex-Ohmeda, the Odin Browser and ADI LabChart. These applications were specifically chosen as they cover a large percentage of the monitoring system market. However, the interfaces are standardized to a degree where it is possible to incorporate a new type of system if the consortium deemed it worthwhile.
- For support and maintenance, it is ideal to have some kind of remote connection available – usually through a Virtual Private Network (VPN) – to allow administration from the central developers. This often requires administrative “buy-in” from the local hospital, which is not always easy to obtain. However, the system is constantly being developed to allow for this. The main requirement in the instance of no remote access is to have someone with a detailed knowledge of the IT systems and a direct line of communication to the central developers. This is often non-trivial to establish when dealing with highly federated and security-oriented organisations.

V. CONCLUSIONS

In this paper we have outlined the AvertIT HypoNet data infrastructure that allows the real-time admission and uploading of physiological, episodic and demographic data from neurosurgical centres across Europe. Through this platform it is now possible to support in-depth analysis and development of a range of clinical studies with the example shown here being prediction of arterial hypotension. Furthermore, whilst our approach for prediction is based upon a BANN, trained on existing brain trauma data sets, other

approaches as outlined in section 2.C can equally be applied. The extended ramifications and application of such a data platform for wider physiological research are considerable.

As mentioned in section 2.B, an important scientific output of the Avert-IT project has been analyses showing that hypotensive event detection is dependent upon whether mean blood pressure (BPm) is considered in the event definition. Further information on this result can be found at [3].

The platform is now live and in active use to support the admission of patients to the clinical trial designed to test the accuracy of the HypoPredict application, running until March 2011. A key feature of this study was undertaking the ethics approval process for each site. This has been successfully completed and approval granted.

In terms of the trial itself, we aim to establish the accuracy and sensitivity of the Hypo-Predict engine and its use in a clinical setting. In feedback from the clinicians, their expectations of accuracy vary, but a consensus that the HypoPredict system should provide a prediction accuracy of at least 30% to be of clinical value. It is also important to note that the HypoNet data platform as described here is not part of the trial itself, i.e. it is the clinical use that is the focus. In short, the HypoNet platform works and is used for real, live clinical research.

The AvertIT core server and associated infrastructure itself resides at the University of Glasgow and is available for follow-up studies, e.g. development of Virtual Physiological Human studies. The results of the HypoPredict development and outcome of the final clinical trial will be outlined in a later paper.

ACKNOWLEDGMENTS

The authors would like to acknowledge the EU FP7 program (grant number 217049), kindly funding the Avert-IT project.

We would also like to acknowledge the work of the BrainIT group investigators and participating centres to the BrainIT dataset without whom this work would not have been possible. These include: Barcelona, Spain: Prof. Sahuquillo; Cambridge, UK: Prof. Pickard; Edinburgh, UK: Prof. Whittle; Glasgow, UK: Mr. Dunn; Gothenburg, Sweden: Dr. Rydenhag; Heidelberg, Germany: Dr. Kiening; Iasi, Romania: Dr. Iencean; Kaunas, Lithuania: Prof. Pavalkis; Leipzig, Germany: Prof. Meixensberger; Leuven, Belgium: Prof. Goffin; Mannheim, Germany: Prof. Vajkoczy; Milano, Italy: Prof. Stocchetti; Monza, Italy: Dr. Citerio; Newcastle upon Tyne, UK: Dr. Chambers; Novara, Italy: Prof. Della Corte; Southampton, UK: Dr. Hell; Uppsala, Sweden: Prof. Enblad;

Torino, Italy: Dr. Mascia; Vilnius, Lithuania: Prof. Jarzemaskas; Zürich, Switzerland: Prof. Stocker.

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] 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.
- [3] Stell, Sinnott, Jiang et al, Federating Distributed Clinical Data for the prediction of adverse hypotensive events, e-Science All-Hands conference 2008 [need to update ref to Royal Society]
- [4] 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
- [5] Philips Medical Event Surveillance Monitor: <http://www.medical.philips.com/main/\products/patient-monitoring/assets/docs/event-surv4522-982-91371.pdf>
- [6] 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.
- [7] 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.
- [8] WSIT standard: http://en.wikipedia.org/wiki/Web_Services_Interoperability_Technology
- [9] Philips DocVu: <http://www.medical.philips.com/main/products/patient-monitoring/products/doc-center/>
- [10] Odin Browser Program – T.P. Howells, I.R. Piper, P.A. Jones, M. Souter, and J.D. Miller. Design of a research database for the study of secondary insults following head injury. *Journal of Neurotrauma*, 12:471-472, 1995.
- [11] HL7: <http://www.hl7.org>
- [12] HAPI: <http://hl7api.sourceforge.net/>
- [13] Proxifier: <http://www.proxifier.com>
- [14] NTLM: <http://en.wikipedia.org/wiki/NTLM>
- [15] ADI: <http://www.adinstruments.com>
- [16] GE Healthcare monitors: http://www.gehealthcare.com/euen/patient_monitoring/products/imm-monitoring/datex-ohmeda/index.html
- [17] Trac: <http://nescforge.nesc.gla.ac.uk/trac/avertit>
- [18] C3 Global: <http://www.c3global.com>
- [19] Zanier E, Ortolano F, Ghisoni L, Colombo A, Losappo S, Stocchetti N. Intracranial pressure monitoring in intensive care: clinical advantages of computerized monitoring over manual recording. *Crit. Care Med.* 2007; 11:117
- [20] 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.