

Probabilistic Unit Life Status Estimation (PULSE)

Subrata Das, Josh Introne, and Dave Lawless

Charles River Analytics, Inc.
Cambridge, MA 02138
USA
sdas@cra.com

Reed Hoyt and Steve Muza

US Army Research Institute of Environmental Medicine
Natick, MA 01760
USA
reed.hoyt@na.amedd.army.mil

Abstract – We present here the Probabilistic Unit Life Status Estimation (PULSE) methodology and software architecture. The PULSE approach leverages Temporal Belief Networks (TBNs) to model the subject's fundamental physiological dynamics (e.g. shock response) over time, to perform life status estimation in a robust manner that accounts for sensor hardware malfunctions. We developed a limited-scope prototype using our in-house TBN engine to estimate clinical status, which served as a framework to demonstrate the PULSE approach for two separate but related clinical assessment tasks. In the first task, the PULSE prototype provides a life status assessment of a simulated unit, allowing for the principled introduction of noise to demonstrate the system's ability to detect sensor failure. In the second task, the PULSE prototype provides a clinical assessment of a unit's degree of acclimatization, using pre-recorded data from studies of soldiers at altitude.

Keywords: Life Status Assessment, Acclimatization Assessment, Temporal Belief Networks.

1 Introduction

The growing digitization of the battlefield will produce a quantum leap in the quality of medical care afforded to thousands of soldiers in the field. This important advance will be enabled by near real-time physiological status monitoring of warfighters, incorporating an array of personal biomedical sensors worn by each soldier.

To realize such enormous potential requires sophisticated Sensor Processing and Assessment (SP&A). In this paper we describe the PULSE (Probabilistic Unit Life Status Estimation) methodological framework and software architecture for the development of custom solutions to remote (telemedical) SP&A problems. Our approach meets four criteria we have identified as crucial for SP&A software solutions; it is flexible, extensible, robust, and dynamic.

An SP&A component must be very flexible, both in terms of its sensor inputs and its derived outputs. This is because the inputs from the sensor array may be custom configured for different situations; various combinations of sensors may be employed to provide such varied data as heart rate (pulse), blood pressure, respiration rate, body position and/or motion, core and skin temperatures, blood oxygen saturation, electroencephalogram (EEG), and electrocardiogram (EKG). Moreover, the active sensor configuration may change dynamically, depending on the inputs and outputs currently needed; for example, simple life status monitoring may suffice most of the time, but a sudden drop in pulse may necessitate activating an intensive monitoring mode, possibly adding more sensors to the active configuration.

An SP&A component must be very extensible, because new and improved sensors are frequently developed, and it will be necessary to easily incorporate them into the SP&A component. Similarly, support hardware such as sensor harnesses, along with communication, logic, and power modules can all be expected to undergo regular upgrades, which the SP&A component must readily cope with. The same is true of the clinical estimation algorithms used within the SP&A component: these must be readily upgradeable.

An SP&A component must be robust, i.e. it must utilize reasoning techniques that incorporate uncertainty. This is because it must cope with inherently uncertain data, characterized by a range of sensing and measurement artifacts, such as noise, burst errors, bias faults, or stuck sensor conditions. Moreover, given the difficult battlefield conditions anticipated, it is a given that any sensor may fail at any time, e.g. due to sensor dislocation or malfunction, or failure of the sensor array's supporting components (e.g. power supply, communication links). This makes data collection still more uncertain, but the SP&A component will be expected to cope with such failures.

An SP&A component must be dynamic in nature; it must incorporate temporal reasoning and dependencies, due to the temporal nature of the sensor data. Without exception, the sensor data to be analyzed will have temporal dependencies which must be considered to properly deduce the clinical state. As a simple example, a single abnormally low blood pressure reading may not be especially significant, but several such readings over several minutes of monitoring would probably be cause for concern: they may indicate shock and impending death (although vaso-vagal syncope might also be suspected, i.e. fainting).

The PULSE methodology and architecture have been designed with these criteria in mind. To ensure both flexibility and extensibility, we have applied the PULSE approach to two separate domains with significantly different features. In one domain, unit mortality detection, we used data from two simulated sensors to make diagnoses within a relatively small time window (several minutes). In the other domain, acclimatization assessment, we utilize real (very noisy) data to make diagnoses within a relatively long time window (12 days).

To ensure robustness and temporal sensitivity, we leverage Temporal Belief Network (TBNs) technology (Kanazawa, Koller, and Russell, 1994; Nicholson and Brady, 1994; Murphy, 2002), which explicitly provides the temporal reasoning needed. TBNs provide great flexibility in modeling real world problems, including

sensor processing and medical diagnosis. TBNs are also robust in the sense of providing probabilistic reasoning for dealing with uncertain data. We also expect to benefit from the modularized BN structures, which allow sensor behaviors to be modeled individually, providing the required extensibility needed. We expect that we can apply the powerful sensitivity analysis feature (Das et al., 2001, 2002) to fault diagnosis of the modeled sensor array, thereby improving the output clinical state estimates by adjusting for sensor failures. We prefer TBN technology because it integrates the aforementioned features, especially the temporal aspect, which cannot readily be obtained from other traditional AI techniques such as Neural Networks, Rule-based Expert System, and Fuzzy Logic.

The rest of the paper is organized as follows. Section 2 provides a brief introduction to TBN technology. Section 3 presents our methodology, and section 4 introduces our software framework. Our approaches to life status estimation and acclimatization assessment are detailed in sections 5 and 6 respectively.

2 Temporal Belief Networks (TBNs)

A Bayesian Belief Network (BN) (Pearl, 1988; Jensen, 1996) is a graphical, probabilistic knowledge representation of a collection of variables describing some domain. The nodes of the BN denote the variables and the links denote causal relationships between the variables. The topology encodes the qualitative knowledge about the domain. Conditional probability tables (CPTs) encode the quantitative details (strengths) of the causal relationships. There are two key ideas in extending a BN to a TBN:

1. All nodes of the BN are associated with particular time steps, simply by indexing the nodes with a time step value. The time steps may be irregular, especially if an event-driven scenario is modeled, or they may be regular, in all possible fixed increments.
2. Some BN nodes for a given time step may have causal dependencies on nodes from earlier time steps (in addition to the usual causal dependencies on nodes from their own time step); such dependencies are called temporal dependencies.

The result is a TBN (Dean and Kanazawa, 1989; Nicholson and Brady, 1994, Ghahramani, 2001; Murphy, 2002). Figure 1 illustrates the general case of time indexing and temporal dependency.




Figure 1: General Temporal Dependency

Variable $Y(t)$ from time step t has causal dependencies on some variables $X(i,t)$ from its own time step t ; it also

has causal dependencies on some $X(i,t-j)$ from earlier time steps. When temporal dependencies between nodes are only from the immediately prior time step, we say we have a Markovian model. Refinements to the basic TBN extension paradigm are discussed in the following section, in particular iteration and estimation techniques.

We introduce the formalism of a *primary iterant*, which is a template network consisting of all the nodes and causal links for a particular time step t of the TBN, together with any nodes and *temporal links* from prior time steps to step t . In this section we provide a simple example in the weather domain for the sake of explanation; in the following sections, we will provide specific examples for our target domains. In Figure 2, the example indicates that, for any given time t ($t=3$ is shown), there are two time-indexed variables in the TBN, namely $Clouds(t)$ and $Rain(t)$, and the $Rain(t)$ variable for any time step t is causally dependent upon the variables $Clouds(t)$, $Clouds(t-1)$, and $Clouds(t-2)$. Semantically, the diagram models the belief that the chance of rain occurring on any given day depends not only whether it's cloudy that day, but also upon the cloud cover for the past two days.




Figure 2: Sample TBN Primary Iterant

The primary iterant is the initial version of the TBN; it is also a template used to "iterate" the network, i.e. to enlarge the network to include the next time step. Iteration is a two-step process (illustrated in Figure 3):

1. A copy of the primary iterant is produced with a high index 1 greater than the current network's high index.
2. The iterated network is then created by equating nodes of the current network and the primary iterant copy having the same indexed names.

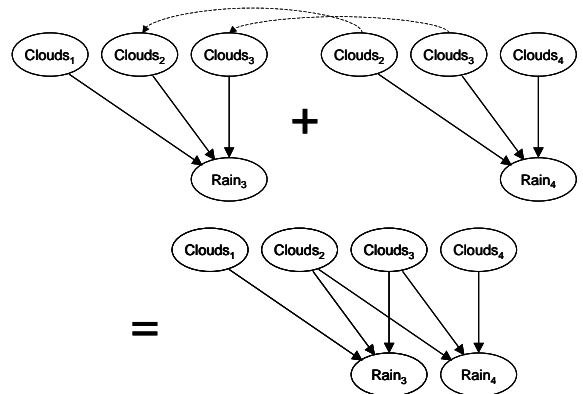


Figure 3: TBN Iteration Process

After much iteration, the TBN will inevitably grow so large that it becomes difficult to work with. To deal with this problem, we devised a technique of estimating TBNs using special nodes we call *history nodes*, which allow truncation of a TBN prior to a given iteration, with the

truncated portion summarized in the history nodes; these are created as follows (illustrated in Figure 4):

1. The index N of the latest desired complete iteration for the estimated TBN is chosen ($N=26$ in Figure 4).
2. The history nodes will be those nodes with index $<N$ having direct causal influence on (i.e. are parents of) nodes with index $\geq N$.
3. The belief values of the history nodes are noted.
4. All non-history nodes with index $<N$ are deleted, together with their incident causal links; any causal links between history nodes are also deleted, so that the history nodes are causally independent. However, causal links from history nodes to non-history nodes (their children) are preserved.
5. The history nodes' *a priori* belief values are then set as noted above.

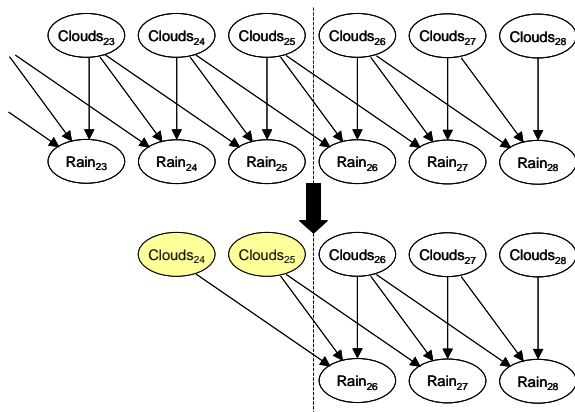


Figure 4: TBN Estimation with History Nodes

The causal independence of history nodes (which are always root nodes) is key to the estimation technique, as it allows us to set their belief values to our choice of historical information; we use this flexibility to help “preserve” belief values throughout the network. We found that this technique generally results in the estimated TBN matching the belief metrics of the full TBN to within purposes. The technique facilitates repetitive estimation as the TBN is iterated, keeping the working copy of the TBN reasonably small. It has the additional significant benefit of allowing the use of the same CPTs for all iterations of non-history nodes; this is because the parents of each non-history node are always kept the same, by using the history nodes. Finally, it helps to reduce graphical clutter, although such clutter is generally unavoidable in practice (e.g. a large, fully specified spacecraft model), as each time iteration could require dozens of nodes. For various other exact and approximate inference algorithms for TBNs, see (Boyen and Koller, 1998; Kanazawa et al., 1995; Murphy, 2002).

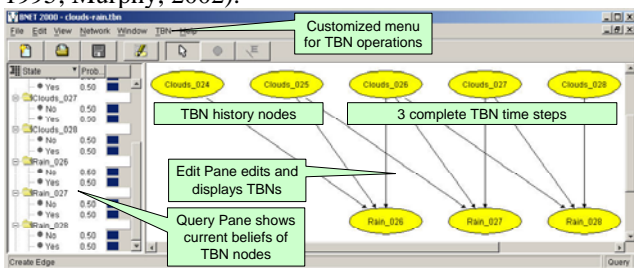


Figure 5: TBN software running Clouds/Rain example

Using standard object-oriented programming techniques, an in-house software package for BNs was extended with new functionality for TBNs, including the ability to create, edit, save, load, query, iterate, estimate, retract, change the number of iterations kept, and check and change evidence in the context of PULSE.

TBN iteration and estimation are invoked automatically, driven by user specified parameters. TBNs are saved in an XML format, including any posted evidence. All work was done in platform-independent Java. Figure 5 shows the software running the estimated TBN of Figure 4.

3 A TBN Methodology for SP&A

We have developed a methodology for the application of Temporal Belief Networks to the task of SP&A. Our methodological approach captures the human expert’s approach to sensor failure detection. Simply stated, given a set of metrics used to monitor any situation, if some subset of those metrics do not match the dynamic behavior that are known (whatever these behaviors are indicative of) it is likely that those metrics are being reported inaccurately.

To illustrate the approach, consider the following (simplistic) example. When driving a car, several metrics are used to monitor the speed and acceleration of the car – there is the sensation of acceleration, the sound of the engine, the appearance of the external world, and the speedometer. As the car accelerates, certain observables are expected from each of these metrics – the speedometer needle moves clockwise, there is a sense of pressure against the seat back, etc. If all metrics except the speedometer needle indicated acceleration was taking place, one would naturally assume (and likely be accurate in that assumption) that something was faulty with the speedometer itself.

Different sensors have different operating characteristics, and an expert will have different levels of confidence in each sensor. Furthermore, when observed metrics do not correlate with known possible behaviors, suspicion of sensor failure will be influenced by this confidence. Continuing with the car example, if the sense of pressure against the seat back was felt, yet the view outside the car appeared to be static, one might assume that acceleration was in fact occurring, but that something was wrong with the reporting of visual information (maybe the windshield was painted?) and not in fact the pressure information (corresponding to a greater confidence in the tactile sense of motion than the visual appearance of motion).

Figure 6 illustrates this approach to status estimation and sensor failure for life status estimation. Knowledge Elicitation (KE) work with medical experts will elicit information about physiological trends that would occur in biometric data corresponding to various physical conditions, e.g. the SME would likely indicate that serious arterial bleeding would result in increasing pulse and decreasing systolic pressure (due to severe shock). From this we can programmatically create a TBN (the correlative physiological model in the figure) that recognizes the biometric data waveform corresponding to

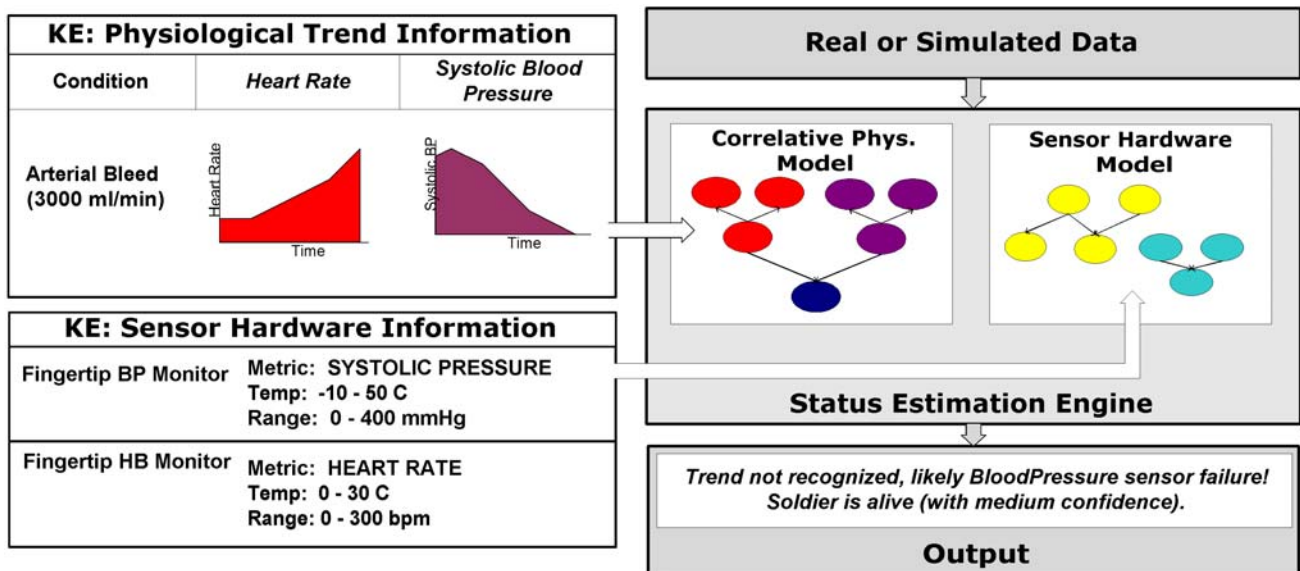


Figure 6 Overview of Methodology for SP & A

the increased pulse expected during arterial bleeding, and save the TBN as a subnet in the library. A comprehensive library of such subnets would be created covering all known physical conditions, for all biometric data types expected.

Additionally, via more KE work with hardware engineers, the BN library would receive subnets (the sensor hardware model in the figure) that model the behavior of all available biometric sensors that could be used. This provides confidence estimates for the various sensors employed, so that a set of behaviors that are not known will allow us to determine which hardware is most likely failing.

The Temporal BN subsequently assembled and used for assessment would load pertinent subnets from the BN library as building blocks, including: 1) subnets for the sensor hardware used; and 2) subnets for all known physical conditions derivable from the sensors chosen; this is shown in the center right portion of the figure. Finally, when the assembled Temporal BN is loaded into the belief network engine and used for assessments, incoming sensor data will produce waveforms (shown in the upper right of the figure) that will be recognized and correlated to a particular known physical condition. Failing that, we expect there must be some sensor error since we have covered all human possibilities, so the Temporal BN will trigger an error indication (shown in the lower right of the figure).

4 A Software Framework for SP&A

To demonstrate the utility and feasibility of our proposed temporal BN approach to SP&A, we developed the Probabilistic Unit Life States Estimation (PULSE) architecture. The PULSE system leverages the Temporal Belief Network tools described to produce dynamic clinical assessments. As a byproduct of the clinical assessment, the PULSE system also produces an assessment of reliability of the underlying hardware, and supports fault isolation analysis techniques.

We have developed a modular architecture for the development and testing of TBNs for SP&A. A schematic of this architecture is shown in **Error! Reference source not found.**

As seen in the diagram, there are three major functional components to the PULSE system:

- 1 **Simulation Layer** - creates and/or provides the various biometric measurements to be used.
- 2 **Sensor Layer** - retrieves biometric measurements from the simulation layer, and adds noise depending on the type of sensor failure to be simulated. These parameters can be modified at run time.
- 3 **Status Estimation Layer** takes sensor readings and performs assessments using Temporal BNs to deliver clinical status estimates along with confidence estimates and sensor error alerts.

In the following, we document the application of our software framework to two state-of-the-art problems within telemedicine: life-status estimation and acclimatization assessment.

5 Life Status Estimation

The life status estimation task involves making a simple determination as to whether or not a unit is alive on the basis of remote sensor data. In this task, overcoming possible sensor failures and determining where such failures lie is far more difficult than the clinical assessment itself.

5.1 Data

There is a dearth of available data for fielded units with life-threatening medical conditions. As a result, it was necessary to utilize a third-party simulation environment to generate a plausible data-stream.

Following a survey and review, we selected Bodysim™ (Advanced Simulation Corporation, <http://www.advsim.com>) as the best candidate for our needs. Bodysim™ provides a realistic sensor data stream at 165 Hz (165 measurements per second, not necessarily

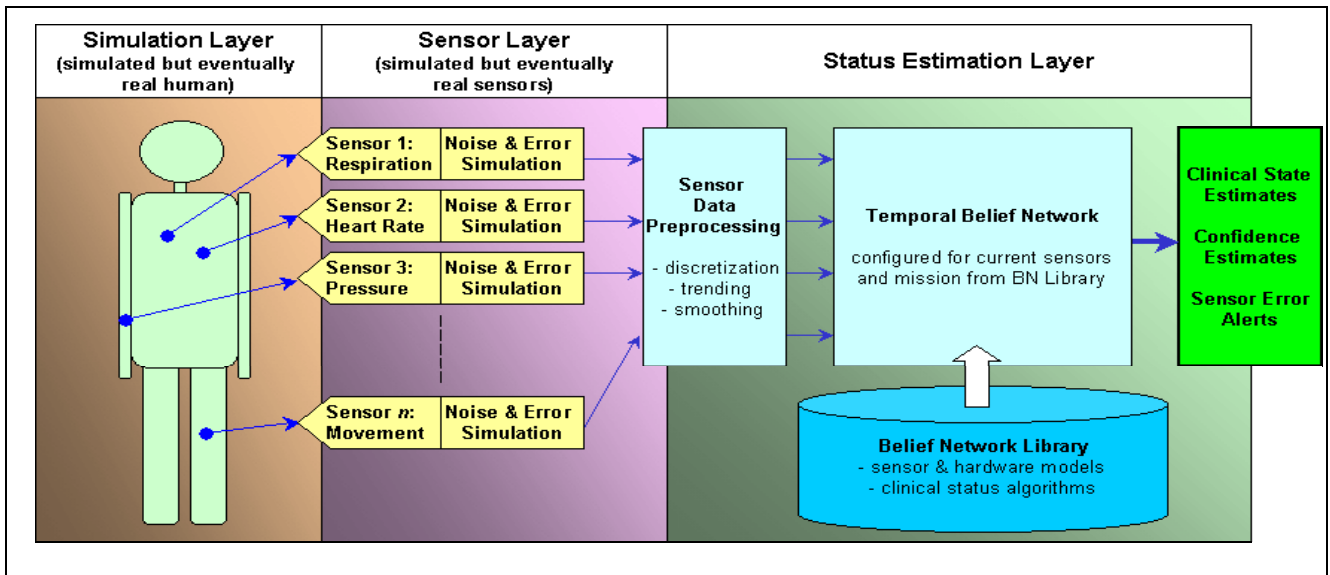


Figure 7 Software Framework for SP & A

in real time), although physiological signals provided by the package are not necessary those that would be expected from a deployed sensor array. Furthermore, trauma simulation within Bodysim is not well supported. For example, severe bleeding could be simulated, but the wound site cannot be specified. Other stock trauma simulations provided with the Bodysim™ package could be interpreted as modeling some sort of battlefield injury, although not usually with the specificity we would have preferred.

Despite its drawbacks, Bodysim™ was the best available simulation package for software developers at the time of this study.

5.2 The Assessment Network

In order to develop an assessment network for mortality assessment, knowledge elicitation sessions that were held with medical experts. In these elicitation sessions, we sought to accumulate two types of information:

- Physiological patterns indicating likely trauma, thus providing information for assessment.
- Physiological patterns that were unlikely to appear, providing information for sensor validation.

The elicited information was transcribed into a TBN which uses two sensors: Systolic blood pressure, and pulse. Figure 8 is a screenshot of the Mortality Assessment Network in its entirety. The network covers three distinct time slices, and has two history that accumulate state in the network, and four input nodes that

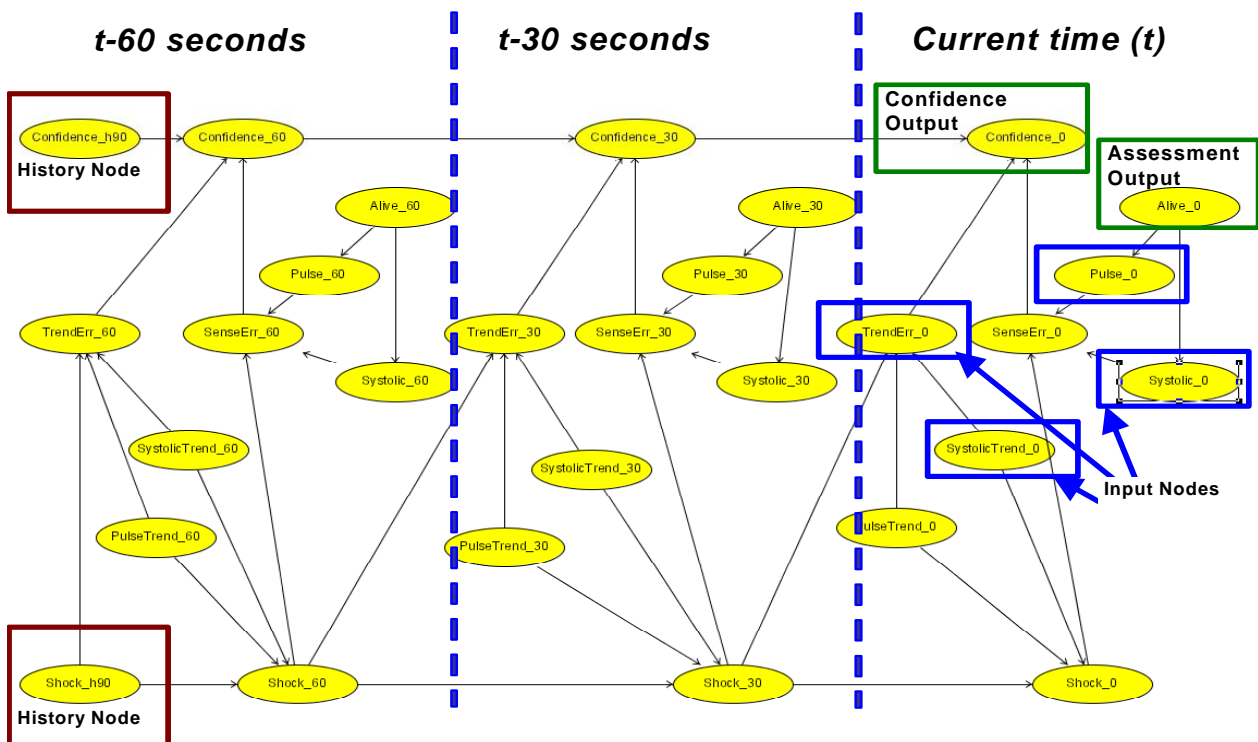


Figure 8 The Assessment Network

are posted with trend and absolute evidence from the sensors. The network spans a ninety seconds including history nodes, as this is a useful time-scale when trying to detect the physiological patterns that accompany serious trauma (discussed further below). Output is read from the Confidence and Alive nodes.

Figure 9 provides a conceptual breakdown of a single slice of the above network. There are three “modules” within a time-slice. These modules are : 1) the Physiological Module, which provides a confidence component according to how well observed metrics match known physiological (in this case, shock or normal physiological states) trends; 2) the Hardware Module, which provides a confidence given the operating parameters of the sensor hardware (the was maximum and minimum values allowed for each sensor); and 3) the Clinical Module, which provides an assessment of the subjects mortality on the basis of current readings from the sensor hardware. Note that there is no direct connection between the physiological module and the clinical module in this network; thus, the goodness of fit between the known physiological patterns and the observed sensor readings contribute to the confidence in the hardware, but do not influence the assessment directly. This separation was useful in the development of this network, but is not a necessary feature.

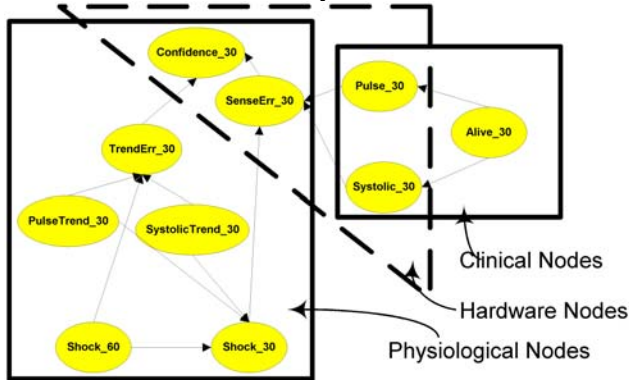


Figure 9: Conceptual Breakdown of Mortality Network

The mortality assessment is made by detecting whether or not the monitored unit has both a pulse and viable blood pressure reading. Illegal states (impossible values on either heart rate or blood pressure sensors) or states that are unlikely given the current stage of shock register an error at the “SenseErr” node, and this also reduces confidence in the hardware.

As shown in Figure 8, the network has two outputs, a clinical assessment (the “Alive node”), and a hardware confidence (the “Confidence”). These are to be interpreted as follows:

Alive₀ – The current best estimate as to whether or not the unit is still alive. In general, this value should be either “yes” or “no” (with some very small deviance from certainty). Occasionally, when the clinical state is not recognizable (e.g. no blood pressure yet continued heart beat, which is possible in Bodysim™) each state will have a non-negligible probability. This may be interpreted as an “indeterminate” state, in which an assessment is not really possible.

Confidence₀ – The current estimate as to the reliability of the underlying hardware. The interpretation of this node is fairly straightforward

Ultimately, the output from these nodes would be used to drive a “user-friendly” layer that interprets the output so as to provide a well-tuned decision support tool. However, such a layer was not implemented in this prototype.

Validation was performed with an SME, and although it was difficult to find subjects who were willing to provide data, our network performed very well for simulated data over a variety of sensor conditions. More formal validation experiments are slated.

6 Acclimatization Assessment

The acclimatization assessment problem involves assessing a subject’s degree of adaptation to high-altitude conditions over time. Unlike the life status estimation problem, the clinical assessment in the acclimatization problem is not well defined (more medical professionals will agree about a person’s life status than about degree of acclimatization).

6.1 Data

The data used in this scenario was “real” data collected from a sensor harness that was originally designed to be used for monitoring humans at rest, and was consequently very noisy. Finally, because the adaptation process spans several days (as opposed to the much shorter significant intervals in the trauma case detected by the life-status estimation network), two level of historical summarization were required in the assessment network to summarize data over longer periods of time.

The data that served as the source of simulation for the acclimatization scenario was collected from actual military field tests at Pike’s peak in 2001. Each data set provided roughly eight hours of data collected on days one, six, and twelve via a sensor harness. The data contained two separate ECG traces from two separate leads (sampled every 4 ms), one respiration rate (RR) trace (sampled every 40 ms), and one SpO2 trace (sampled every 400ms). All classes of data were subject to significant noise in the form of artifacts and drop-outs, although a proprietary algorithm was available from the manufacturer to extract heart-rate from ECG data which performed some error-correction.

However, data from the RR sensor was too noisy to process directly, but there were no available algorithms for smoothing the data. Figure 10a shows a representative sample of the respiration data (two minutes of data are shown). People generally breathe at somewhere between ten and fourteen breaths per minute, thus the sample shown should have on the order of twenty to thirty complete waves.

As it was difficult to distinguish actual peaks from artifactual data, we implemented a simple low-pass filter in an effort to screen out some of the noise. The low-pass




Figure 10: Respiration data a) Before smoothing; b) After smoothing c) After regression

filter converts the signal to the time-domain (via FFT), discards all but N low-frequency components (where N is determined at compile time), and then converts these

components back to the time domain (via iFFT) to obtain the smoothed waveform. Figure 10b shows the result of applying this filter to the raw data in Figure 10a.

After initial smoothing, a linear regression was run (with a window size chosen at compile time) to identify significant inflection points. This, in effect, screens out peaks that are “small” where small is a function of amplitude and wavelength, as well as the size of the window. Half the number of zero crossings (the number of complete waves) in the minute preceding the current time is thus the respiration rate in breaths per minute. This procedure was applied at runtime, and thus no preprocessing of the respiration data was necessary. Figure 10c shows the result of this procedure (note, the inverse of the slope is shown here to make the zero crossings more obvious).

6.2 The Assessment network

The acclimatization network (shown in Figure 11) is designed to produce assessments about the monitored unit’s degree of adaptation to high altitude conditions over a series of days, as opposed to the mortality assessment, which was possible to do over a series of minutes. This required modifications to be made the underlying temporal belief network software framework, as caching 12 days (or even a single day’s worth) of data from a data stream that was sampled at the sub-second time scale creates tractability problems.

To address this difficulty, the temporal framework was augmented with summary nodes, which are similar to history nodes in that they capture the beliefs of the network at given time, but are updated with coarser temporal granularity than the other nodes. In our case, the summary nodes are only updated when the next available day’s worth of data is observed. Thus, there are two levels of historical summarization present within the Acclimatization Network.

Figure 12 provides a conceptual breakdown of the “Current Time” portion of the Acclimatization network. The breakdown is somewhat different than in the case of the mortality estimation network, as the clinical module in this case is entirely subsumed by the physiological

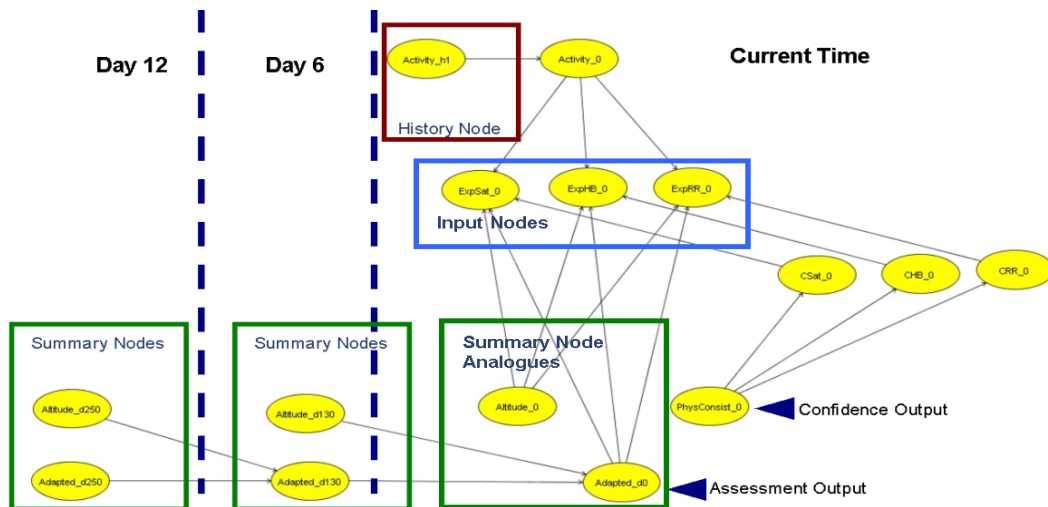


Figure 11 The Acclimatization Network

module. In this case, the physiological characteristics given a degree of adaptation, the altitude the monitored unit is stationed at, and the units' degree of activity are captured in the probability tables of these nodes. The clinical module is actually implicit in these nodes, as the propagation of evidence through the network will cause the "Adapted_0" node to be updated according to evidence that is posted at the input nodes, and the degree of adaptation during previous days.




Figure 12: Conceptual Breakdown of the Acclimatization Network

The hardware module models both unlikely states in the individual sensors (extreme values that are not possible) and unlikely states given the assumed physiological state. Note the causal direction of the links is from the consistency nodes (CSat_0, CHB_0, CRR_0), to the three input nodes (ExpSat_0, ExpHB_0, ExpRR_0). Thus, the probability tables at the input nodes are conditioned upon the state of the of the consistency nodes. Because the reasoning in belief networks is essentially directionless, sensor values which cause states at the input nodes that fall outside the values that could be expected if the sensors are operating correctly, the consistency nodes will register the potential sensor failure after evidence is propagated.

As in the case of the mortality network, validation was performed with the assistance of medical subject matter experts. In this informal validation, the network was observed to provide clinically plausible assessments, and was able to identify locations in the data logs which were very likely to be erroneous. More formal validation studies are planned.

7 Conclusions

We have presented a general modeling methodology for SP&A, with application to two domains in the field of telemedicine. The assessment algorithms provided probabilistic status estimates, along with probabilistic hardware status estimates (i.e. probability of a sensor reading being correct). These probabilities were computed using temporal TBN models.

Our methodology is modeled after common sense human reasoning. This methodology involves establishing a plausible probability space and encoding this space into a temporal belief network. In summary, our methodology is as follows: 1) establish expected correlation between individual physiological parameters and environmental conditions; 2) establish confidences and operating parameters for a variety of sensors; 3) combine these in modular fashion into a single temporal belief network.

We plan to extend and refine the PULSE prototype with additional and more refined sensor models, and more complete physiological models. Ultimately, we hope to accumulate a library of data models for various physiological variables, which can be rapidly combined using a framework like the one discussed here into custom built SP&A solutions for a variety of telemedicine applications.

Acknowledgements

This work was supported by the U. S. Army Medical Research and Materiel Command under Contract No. DAMD17-03-C-0060.

References

- [1] Boyen, X. and Koller, D. (1998). Tractable inference for complex stochastic processes. Proceedings of UAI.
- [2] Das, S., Lawless, D., Hanson, M., & Zacharias, G. (2001). "Model-based Reasoning with Temporal Belief Networks", NASA Contract No. NAS2-01048, Final Report. Cambridge, MA, Charles River Analytics.
- [3] Das, S. and Lawless, D. (2002). Network-based Truth Maintenance System. Proceedings of the 15th European Conference on Artificial Intelligence (ECAI), Lyon, France.
- [4] Dean, T. and Kanazawa, K. (1989). A model for reasoning about persistence and causation. Computational Intelligence, Vol 5.
- [5] Ghahramani, Z. (2001). An Introduction to Hidden Markov Models and Bayesian Networks. International Journal of Pattern Recognition and Artificial Intelligence, Vol 15(1).
- [6] Jensen, F.V. (1996). An Introduction to Bayesian Networks. Springer-Verlag.
- [7] Kanazawa, K., Koller, D., and Russell, S. (1995). Stochastic simulation algorithms for dynamic probabilistic networks. Proceedings of the 11th UAI.
- [8] Nicholson, A. E., & Brady, J. M. (1994). "Dynamic Belief Networks for Discrete Monitoring". IEEE Transactions on Systems, Man and Cybernetics, 24(11), 1593-1610.
- [9] Murphy, K. (2002). Dynamic Bayesian Networks: Representation, Inference, and Learning. PhD Thesis, Department of Computer Science, UC Berkeley.
- [10] Pearl, J. (1988). Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. San Mateo, CA: Morgan Kaufmann.

Disclaimer for Government Authors: The opinions and assertions expressed in this paper are those of the authors and do not necessarily express the official views of the Department of the Army. The study data presented is from volunteers who gave their free and informed consent. Investigators adhered to AR 70-25 and USAMRDC Regulation 70-25 on Use of Volunteers in Research.