# Life-Signs Determination Model for Warfighter Physiological Status Monitoring

**Mr. Maurizio Borsotto[a], Dr. C. Thomas Savell[a], Dr. Jaques Reifman[b], Dr. Reed W. Hoyt[c], Mr. Gavin Nunns[a], and Mr. Christopher J. Crick[a]**

[a]GCAS Incorporated
1531 Grand Avenue, Suite A, San Marcos, CA 92069-2463, USA
[b]U.S. Army Medical Research and Materiel Command
Telemedicine and Advanced Technology Research Center, MCMR-ZB-T
504 Scott Street, Fort Detrick, MD 21702-5012, USA
[c]U.S. Army Research Institute of Environmental Medicine
Building 42, Kansas Street, Natick, MA 01760-5007, USA

Email: mborsotto@gcas.net / ctsavell@gcas.net / reifman@tatrc.org / reed.hoyt@na.amedd.army.mil / gnunns@gcas.net / ccrick@gcas.net

## *ABSTRACT*

*The U.S. Army is leading an effort for a Warfighter Physiological Status Monitoring (WPSM) system that interprets data from a suite of wearable physiological sensors to infer a soldier's current health status on the battlefield. The future WPSM system will consist of a body-worn network of biosensors with a central processing control unit whose firmware contains a probabilistic Bayesian Network for assessing the soldier's physiological status. The Bayesian Network will assess the status of the soldier in terms of Life-Signs Presence, Absence or Unknown (PAU) state. Together with this health status assessment, another goal of the Bayesian Network will be to assess the related level of confidence in the diagnosis as resulting from clinical uncertainty, sensory information patterns and reliability of the hardware. This information will be made available to field medics and others over separate communication channels, in order to help prioritize the urgency of medical assistance and evacuation. This paper describes the current development of the PAU Determination Model, which demonstrates the various techniques that will be adopted in the final version of the Bayesian Network to fulfill the health status assessment goals and highlights the robustness of the approach.*

## 1.0 INTRODUCTION

Perennial objectives of battlefield force health protection include:

- The reduction in mortality and morbidity rates,
- The enhancement of force effectiveness by reducing the likelihood of non-battle injuries (such as heat stroke and acute mountain sickness), and
- The improvement of casualty management in remote situations.

Many technological advances in body-worn sensory devices have made these goals realizable. Since the first hour after injury is crucial [1], the ability to rapidly locate, triage, diagnose, and render appropriate initial treatments are vital to improving the outcomes of battlefield injuries. In the near future, it will be possible to identify in real time the wounded soldiers on the battlefield who require immediate assistance, even when they are greatly dispersed and out-of-site, which could result in a reduction in battlefield morbidity and mortality.

In response to this challenge, the U.S. Army is developing a set of computerized devices as part of future combat systems, comprised in the Warfighter Physiological Status Monitoring (WPSM) system [2]. The system will feature a configurable array of miniaturized and computationally capable sensors. Among them, physiological sensors will monitor heart and breathing rates, metabolic energy expended while working or marching, skin and core temperatures, activity patterns (body positioning: upright lying face-up or face-down), and several other parameters. These sensors will transmit physiologic information to a small central processing control unit carried by the soldier, where appropriate Bayesian Network model will perform higher-level data analysis. The resulting assessment of the soldier's status will be then made available to the field medic and upper echelons of care, as well as small unit leaders and commanders, as required. To allow quick localization, the uniform will also feature global positioning system capabilities.

Among the assessment capabilities featured by this system are a series of Life Sign Decision Support (LSDS) algorithms. The LSDS algorithms process the sensory data streams and produce meaningful information to help combat medics assess, triage, and manage life-threatening injuries. Specifically, as a primary indicator, a Life-Signs (PAU) status is estimated by these algorithms and transmitted to the field medic or other desired locations as part of the output provided by the WPSM system.

Due to the critical nature of life-signs determination, a key requirement is to reach a high level of confidence in the reliability of the PAU assessment. This means that the system must perform a statistical evaluation of the accuracy of the incoming signals as well as a probabilistic interpretation of the soldier's physiological state. As appropriate for the PAU status determination, these specific algorithms are not concerned with the future state of a wounded soldier but rather are concerned with the use of the present signal information and the associated uncertainties to determine the most likely current state of the soldier. To achieve this result, the LSDS algorithms perform a temporal analysis of the sensory data, including the arbitration of contradictory information, processing of multiple sensors, and performing non-monotonic reasoning on the collected data. The algorithms take into consideration the various elements of data imprecision, as derived from possible sensor/device faults and data transmission failure. They also assess the reliability of the integrated array of sensors and devices by taking into account the probability of failure of each component as well as the probability of failure of the entire sensor array and data transmission system as a whole. All this information on the data imprecision and system reliability is finally merged as part of a diagnostic model that represents both data imprecision and clinical uncertainty pertinent to remotely determine the life-signs status of a soldier. When the system cannot reach a definite determination, the algorithms will report an Unknown condition.

These diverse results can be achieved using the Bayesian Network (BN) probabilistic modeling method [3]. BNs are used to develop knowledge-based applications in domains that are characterized by inherent uncertainty. BNs provide an organized representation of knowledge resulting from the combination of human expertise and statistical analysis. BNs also accept real-time information that they use with the stored knowledge in order to formulate diagnostic or predictive conclusions.

In the case of LSDS, a set of BN models has been developed to satisfy our diagnostic goals. One set of BNs model the behavior of the sensory system, the influence on the sensors of several external factors (e.g., temperature, vibration) and sensor reliability. These sensor models appropriately analyze the incoming sensory data streams, identify possible inconsistency patterns and evaluate the "health" of each sensor in the

WPSM system.

These processed data streams are then forwarded to the PAU Determination Model (PAU-DM). This model incorporated through another set of BN modules the heart of the LSDS algorithms and reproduces the human inference process for PAU determination. A detailed description of the architecture of the PAU-DM subsystem is presented in [4]. In this paper, we provide an overall description of the PAU Determination Model and future research directions.

## 2.0 METHODOLOGY

The PAU-DM will be continually evolved by cycling through stages in an iterative fashion. The first iteration, called "Phase I", which is featured in this article, was intended to highlight the robustness of our approach and demonstrate various techniques that will be adopted in the second version of the system, called "Phase II". In Phase I no claim is made about the correctness of the health status assessment performed by the model. In fact, it is understood that the Phase II model, currently under development, will extend the Phase I model by incorporating sound medical expert knowledge of the human physiology.

The Phase I model can generate a series of ancillary assessment conclusions. These are provided to illustrate how the approach can be used to deliver "amber" outcome warnings indicating that critical situations exist that require immediate intervention. We show that, as a by-product of exploiting appropriate steps in the PAU determination process, a set of indicators can be used as alerts for triggering the medic's attention/intervention. Future releases of the model will include indicators that recognize meaningful situations for use in first-level triage, thus helping the medic establish a weighted priority for providing assistance of the injured soldiers.

The information used for the Phase I PAU-DM was gathered from literature, legislation and consultation with a medical doctor experienced in emergency room trauma situations. A meticulous search was conducted in Phase I for existing procedures or algorithms for determination of death [5] from a remote location. No appropriate models were found. For example, to establish a legal final determination of death, a visual inspection by a physician over a period of time is always required [6]. When determination of brain death is involved, there are also standard procedures that require a medical facility several hours to complete [7]. Since the brain death determination procedure is used to establish death for subjects whose cardio-pulmonary activity has been artificially maintained, it is inappropriate to determining the death of a soldier who has been injured out in the battlefield and not directly helped by a medic. The intent of PAU-DM is not to establish a clinical assessment of death, but rather to be used as a tool in estimating if a wounded soldier has life signs present and thereby helping triage prioritization in operational settings.

The above considerations suggest that it is not possible to establish determination of death with absolute certainty by remote sensor measurements alone, at least from a legal standpoint. In fact, while the lack of both heartbeat and breathing are excellent clues of possible death, they are not conclusive for timely determining the incipience of death. Many other considerations need to be taken into account, including time elapsed, ambient temperature and drug intake among others. Nevertheless, a sensory system can provide reliable information on possible death (absence of life signs) or extreme physical distress that requires immediate attention/medical intervention. In the latter case, it would raise a warning before an irreversible condition is reached.

## 2.1 Modeling Framework

To capture medical assessment expertise for PAU determination, we made use of the Bayesian Networks probabilistic framework [3]. BNs provide a method to represent interdependencies between variables that represent elementary chunks of knowledge, even if the relationships involve uncertainty, unpredictability or imprecision. The relationships may be learned automatically from data files, constructed from experiments or other data, created by an expert, or developed by a combination of these approaches. BNs are used to develop knowledge-based applications in domains that are characterized by inherent uncertainty. A BN allows us to combine prior knowledge and incoming data with the likelihood of a hypothesis of interest, such as a soldier being dead given his/her physiology and the sensory data-time series.

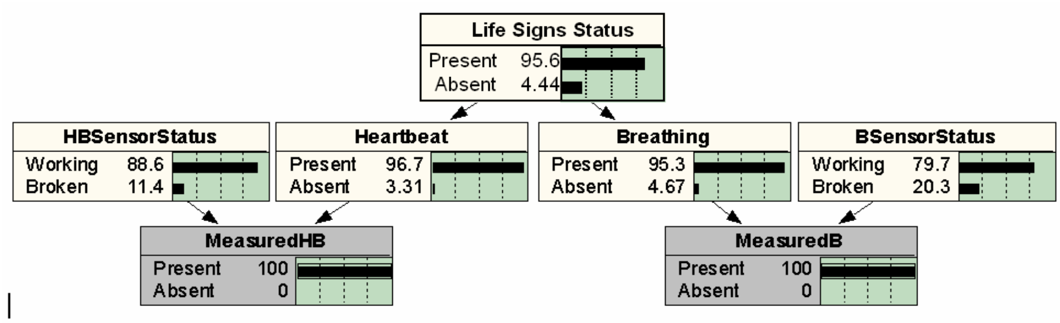


**MeasuredHB**

| HBSensor... | Heartbeat | Present | Absent |
|---|---|---|---|
| Working | Present | 94.000 | 6.000 |
| Working | Absent | 6.000 | 94.000 |
| Broken | Present | 10.000 | 90.000 |
| Broken | Absent | 10.000 | 90.000 |

**MeasuredB**

| BSensorSt... | Breathing | Present | Absent |
|---|---|---|---|
| Working | Present | 70.000 | 30.000 |
| Working | Absent | 5.000 | 95.000 |
| Broken | Present | 15.000 | 85.000 |
| Broken | Absent | 15.000 | 85.000 |

**Figure 1: Bayesian Network and Conditional Probability Tables.**

In a BN, the problem domain is modeled as a set of nodes interconnected with arcs to form a directed acyclic graph. Each node represents a probabilistic variable that can take two or more possible values. The arcs signify the existence of direct influences between the linked variables, and the strength of each influence is quantified by a forward conditional probability. Bayesian Networks do not use "algorithms" in the conventional procedural sense. Rather they give a probabilistic association between an assembly of input variables, as established by experts in the field, as to the relative influence of these variables on the outcome of any given node in the network. The key to this association is the conditional probabilities assigned by the experts to each incoming variable on the state of the receiving node. These conditional probabilities relating the variables to the output state of the node are defined in Conditional Probability Tables (CPTs). Each node in the network has a CPT describing the relative dependence of that node on its parents. Therefore, a CPT is used to define the conditional probability (or likelihood) of a specific value of a variable based on combination of states of the parent variables. A typical BN and some of its CPTs are shown in Figure 1 above.

The time dimension can be handled in a BN by resorting to two strategies. The first strategy is to introduce explicit time variables, starting from the moment a specific condition or set of conditions occur. Examples of this approach in our PAU-DM network are a variable representing time elapsed since the heartbeat stopped and one indicating when breathing has stopped.

The second strategy handles time in an indirect way. This leads to “Dynamic BNs”. In this case, time is represented by discrete values and a model is created to represent the status of the world at a given time-slice. Time is implicitly represented by specifying how one set of variables at a given time-slice is affected by another set of variables from the preceding time-slice. While this is a more elegant approach, it has the drawback of slowing down the inference process. Care must be taken in order to balance cost and benefits. An example of this second approach in our PAU-DM network is the modeling of the dynamics of oxygen saturation. Since we use both strategies to represent time in our model, the mix between the two methods has been selected so to minimize the interdependence between subsequent time slices in order to improve the speed of inference.

### 2.2.1 Why use a BN Framework?

The PAU Determination Model utilizes a BN used to infer the PAU status of a soldier using the information provided by the various sensory outcomes. The reason for using the BN framework comes from the need for representing and handling the uncertainty, or level of confidence, in the data streams. The BN framework naturally handles this issue and allows us to merge data uncertainty with clinical uncertainty in order to derive a final confidence level in the derived soldier’s health status.

The clinical uncertainty stems from the need for representing medical expertise. Since no procedural algorithm exists for remote death determination, our goal is to reproduce the reasoning process of physicians that are expert in trauma and emergency procedures. This kind of knowledge is best expressed through a probabilistic framework, since the experts themselves are not able to conclude a definite diagnosis for a combination of sensory data. That is, they are not able to easily classify all the possible combinations of data streams into three precise classes, i.e. PAU.

To illustrate, suppose that a subject has been experiencing a lack of circulatory activity for two minutes. The experts we consulted were unable to say whether the subject is either definitely alive or definitely dead, since several contextual, and often intangible factors are not detected with the remote sensory system that would influence this conclusion.

Using this approach, the experts in our team agreed on the fact that a subject is 50% likely to be dead after two minutes of lack of circulatory activity. After another minute, this likelihood may rise to 95%. Finally, they felt comfortable in stating that the subject is definitely dead (100% likely) after a total of six minutes of lack of circulatory activity. All these rules are considered valid unless the subject is experiencing hypothermia. The BN framework has explicitly been devised to capture this kind of non-deterministic reasoning. As discussed in the previous section, it also provides a good level of flexibility in handling the time dimension.

Another issue related to clinical uncertainty is that the set of possible combinations of data streams can be quite large. This makes impractical (and possibly unreliable) to map all the possible combinations into only three categories such as PAU. On the other hand, the BN approach allows us to break down this classification into a combination of much simpler processes that reproduce the human reasoning activity in the specific domain. This has two beneficial effects. The first one is that we need not explicitly encode the full mapping of all the possible combinations of data streams. In fact, this mapping will automatically emerge as a combination of the various simpler reasoning processes. In this way, we can greatly simplify the modeling activity and obtain significant savings in computational resources. The second benefit is that, by reproducing the human reasoning process, every step of reasoning will have a clear meaning and adjustments can be made by focusing on narrow aspects of the problem at a time.

The BN framework also offers a natural approach for taking into account the effects of data uncertainty in the diagnostic process. For example, suppose that we are not sure whether the heartbeat actually stopped beating two minutes or three minutes ago, given the received data streams and the reliability of the hardware. In fact, although the heart rate data stream indicates that the heart stopped two minutes ago, we may have clues that immediately prior data were unreliable. Even if in that minute the heart rate was reported as present, it is possible that it was actually absent. We will therefore provide the PAU-DM network with a likelihood distribution for the heart rate being absent. This distribution may indicate that the two-minute absence is 80% likely to be correct, and the three-minute absence is 20% likely. This kind of information, called evidence, will propagate through the PAU-DM network and merge with the domain knowledge that is encoded in the model. This propagation, called inference, will result in an overall likelihood of the subject being alive. This overall likelihood will then take into consideration the various possibilities regarding the absence of heart rate, and substantially weigh the effect of each one on the final assessment.

The advantage of the BN framework is that it provides a principled way for handling the non-deterministic knowledge described above. Furthermore, the inference process is mathematically exact, in the sense that it obeys the laws of probability theory. This ensures that the outcomes of inference will be always coherent, that is, they do not violate common sense logic. Finally, the graphical approach at the basis of the BN framework is also intuitive. This feature greatly helps the task of knowledge elicitation, since the experts can easily understand the basics of the formalism and contribute directly to the development of the networks.

## 3.0 THE PAU DETERMINATION MODEL

For its health assessment, the PAU-DM network receives measurements from a human subject. Each measurement follows the path shown in Figure 2 before reaching the network. A Human Subject utilizes a physical sensor represented by the Physiological Sensor box. The sensor performs a specific measurement such as heart rate. The resulting measurement is fed into a Pre-Processing module that performs a series of operations such as computing the average heart rate occurring in the last minute, or counting how much time has elapsed since the last heartbeat was present. The Pre-Processor also translates the resulting quantity into a format that can be used by the next step in the chain. Finally, the pre-processed data enter into an appropriate Sensor Model, encoded with a BN. Here the data are analyzed and conditioned by the sensor reliability, and possible sensor failures are detected. The resulting information is then delivered to an appropriate node in the PAU-DM network.

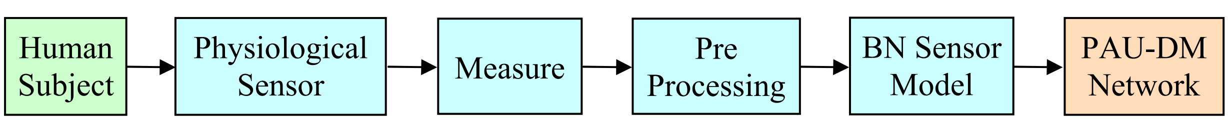


**Figure 2: Information Flow towards the PAU-DM Network.**

The PAU-DM Bayesian Network is shown Figure 3. It encompasses 4,343 probabilistic rules relating 45 variables. This initial model may not be clinically accurate. It is rather intended to represent a reasonable behavior with the goal of illustrating how a BN can be used to solve our diagnostic problem. A more realistic health state assessment model is now under development using the techniques that we are about to illustrate.

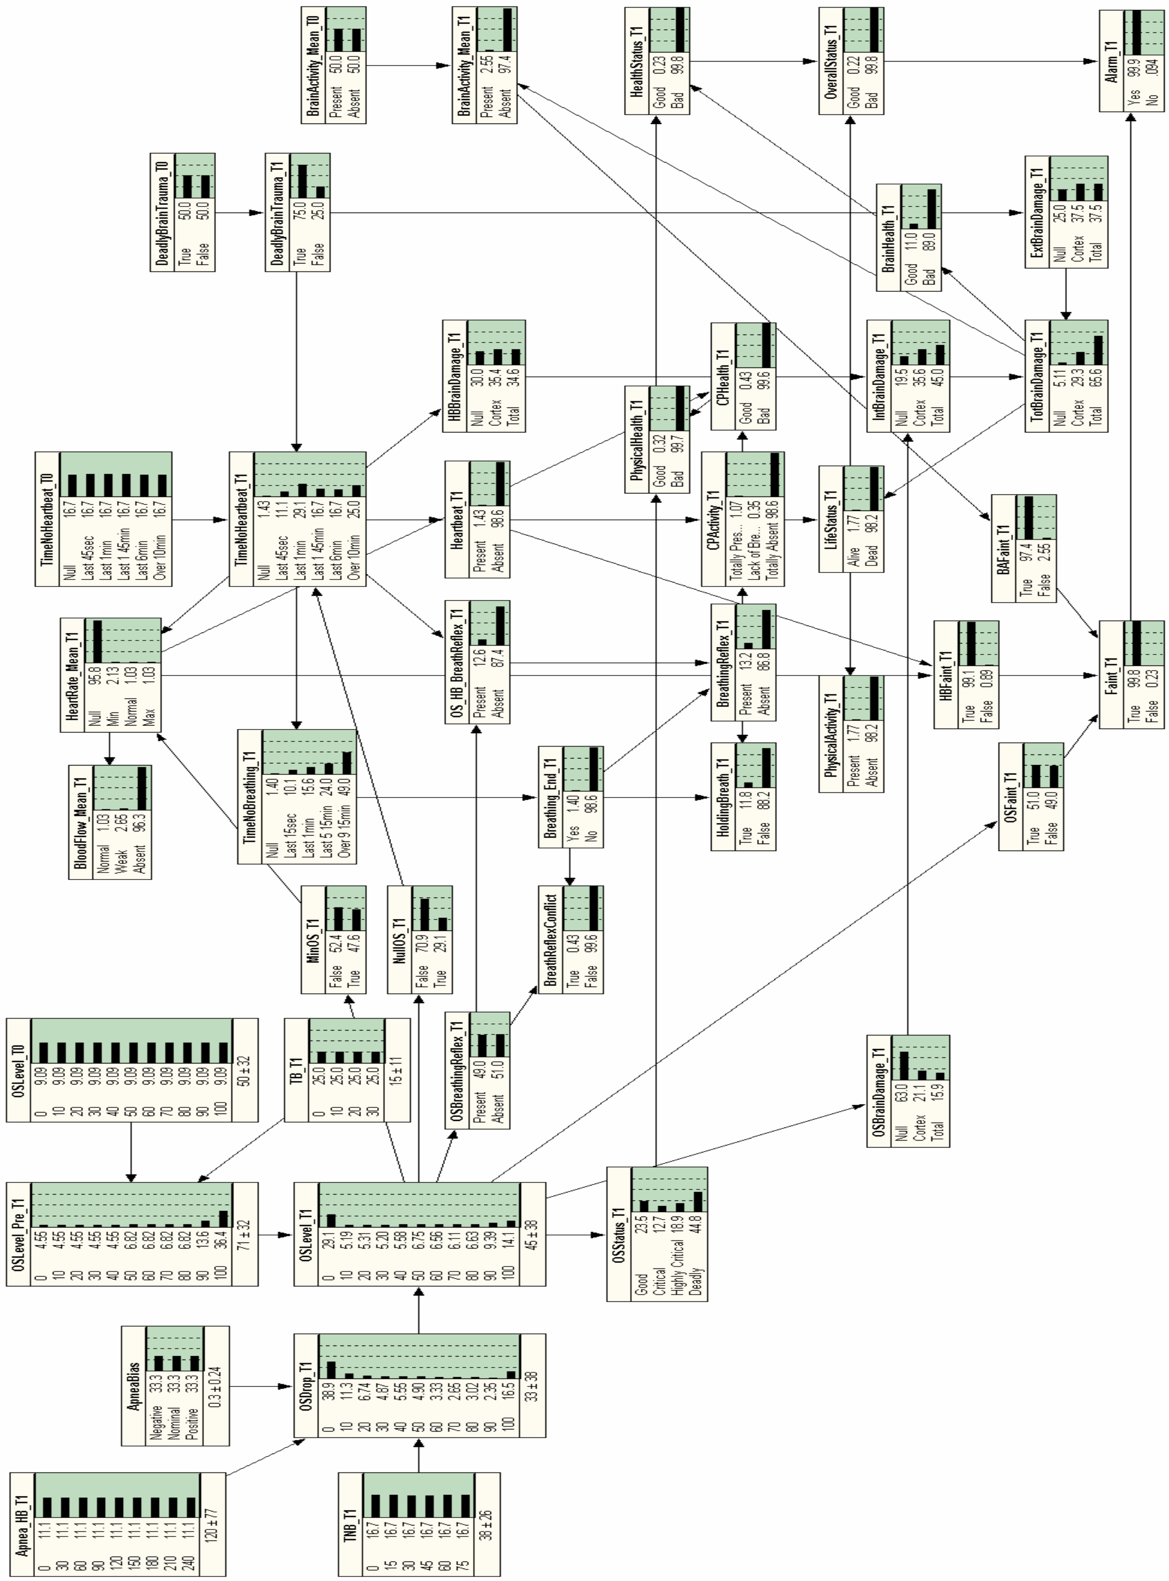


Figure 3: PAU Determination Model.

A general remark on the model is that it represents the subject's status in a one-minute time window. Several variables in the model refer to a one-minute average. For example, the *HeartRate_Mean_T1* variable at the center top of the figure stands for the average value of heart rate in the last one-minute interval. Other variables refer instead to persistence of conditions that can stretch far beyond the last minute, such as *TimeNoHeartbeat_T1*. This variable in fact represents how long the subject has shown no sign of heartbeat.

With this model, a non-deterministic PAU diagnosis can be performed and simple "amber" outcomes can be synthesized as a by-product. An "amber" outcome is an alert that can be used to warn about a critical situation that requires a medic's immediate attention and possible intervention. For example, an oxygen saturation level below 50% sustained over a period is deemed no compatible with life. This oxygen saturation level will therefore raise an alert.

## 4.0 TESTING THE MODEL

In order to test the PAU-DM network with sensory data, we simulated a subject's physical condition over a period of time. This task was carried on by a Human Physiology Simulator that generates realistic data streams for the input variables in the model. The Human Physiology Simulator attempts to provide realistic physiological behavior during the dying process. It represents the assumed "true" physiology model interrelating the time evolution of a set of physiological parameters. A central simulation manager contains a number of rules that link the various physiological states to provide reasonable cause and effect patterns.

Each simulated parameter features random fluctuations to provide realism and takes into account different possible physiological behaviors as emerging from different subject profiles. Three basic profiles are encoded in the simulator representing individuals with different physical fitness. This, for example, affects the levels at which the heart rate is considered too fast or too slow, or the individual capability of coping with apnea.

Using the Human Physiology Simulator, we generated a set of data streams representing different possible dying processes. We then fed the PAU-DM network with these data streams and analyzed the appropriateness of the assessment conclusion, with the help of a medical doctor specialized in intensive care. We also studied the behavior of the model when the network receives only some of the above data streams, to verify what information is more relevant and how the assessment model degrades in performance. Finally, we corrupted the data streams with several levels of noise in order to analyze the robustness of our approach.

Given that we did not use real-world data and because our Human Physiology Simulator was quite simple, we were not able to precisely quantify the PAU-DM performance in terms of sensitivity and specificity. Instead, our medical consultant analyzed the results in order to identify assessments that were clearly inappropriate and to understand the limits of the model. With this qualitative testing, we proved that we are able to obtain an accurate life-signs PAU determination. This holds also for situations in which the data streams present a consistent level of noise, thanks to the choice of averaging the input quantities over a one-minute interval.

There were no simulations in which the system gave a gross misdiagnosis of the subject's condition. The main difference between the physician and the PAU-DM assessment of death was normally related to the onset of the condition. Our model tends to slightly delay this determination mainly because the oxygen saturation model was developed for a full-lung voluntary apnea. This is rarely the case for a subject that has suffered a traumatic injury.

We also want to underline that we did not perform any fine-tuning of the set of parameters present in the PAU-DM network before our testing. The results are therefore even more encouraging because the model is not optimized. The best approach for tuning would use a set of the simulations to calibrate the parameters, and the rest of the simulations as verification baselines. It is expected that by proceeding in this way the assessment performance will greatly improve. We plan to use this approach in the development of the Phase II model.

## 5.0 CONCLUSIONS AND FUTURE WORK

The PAU-DM network we have discussed here represents the starting point for a new, more accurate life-signs PAU determination model. The techniques introduced in this demonstrative network have proved that with an appropriate mix of dynamic behavior and timers, it is possible to properly handle the uncertainty connected to the time dimension. The BN formalism allowed us to translate medical expertise into simple and intuitive elementary models that can be combined together in order to perform the desired assessments. This separation of a complex model into subparts that can be handled independently and then interrelated, provided a powerful yet simple enough tool to compose a coherent set of over 4,000 probabilistic rules.

The resulting PAU-DM network contains several parameters that can be fine-tuned in order to enhance the health state assessment performance. Other parameters allow us to take into consideration the different individual physiological responses expressed by different individuals.

As a byproduct of the PAU assessment process, we also synthesize information that may be effectively used to signal a serious problem requiring immediate medical assistance, before a complete absence (i.e. death) condition is reached. This result goes beyond the original life-signs presence/absence goal and represents an additional benefit that is worth exploiting in future versions of the model.

Building on the success of this proof of concepts, we are now developing a new PAU-DM that makes use of a larger set of sensory data to infer the subject's status. This model will reuse the techniques presented here and perfect the medical knowledge to reproduce a more accurate assessment process.

To generate the appropriate physiological data streams, the user will be able to exploit state-of-the-art physiological simulators, data collected on the field and stored in a file, and even real-time data. In fact, in order to validate the final product we intend to use real measurements of humans and/or animals. This will allow the PAU-DM network to assess the outcome that was experienced by a subject in real life. The resulting assessment from the model will then be compared to the actual clinical outcome.

We are also generalizing the system architecture. We are creating a user-friendly development platform using a flexible framework, whose qualifying features will be easy scalability and modularity. A physiologist will be able to select and compose the assessment system through a simple interface, by selecting a pool of sensors, placing them at appropriate locations on the human body and performing a set of simulations. The software will assemble the appropriate simulation and diagnostic algorithms in the background, given the description of the used sensors, their characteristics and possible redundancy, the behavior of their components, and their performance.

The next generation software under development will allow the user to perform sensitivity analysis. The user will be able to select a sensor suite to measure a set of parameters that are meaningful for PAU determination. For each one of the sensors, the user will be able to specify accuracy levels, failure modes and the likelihood of their occurrence. This information will be taken into account to establish the expected level of confidence in PAU determination over a specified mission time. By changing parameters that describe the quality of the sensors, the user will be able to investigate the change in the PAU determination confidence as the sensors change in performance. Moreover, the user will be able to compare the impacts of different sensor suites that incorporate different kinds of sensors. Further issues such as the effects of environment, aging of the sensors and their self-diagnostic capabilities will also be taken into account.

Several of the sensors that are appropriate for PAU determination are quickly evolving. Our software will be flexible enough to accommodate changes in the architecture and behavior of those sensors. The software architecture has been designed to accept sensors and features that will likely be available in the foreseeable future. Indeed, the software will be used to provide indications of the desired performance of the sensors in order to achieve a reliable health status assessment. Thanks to these features, our software will constitute a valuable tool during the design of the sensor suite appropriate for the PAU goal.

### 5.1 Acknowledgements

The work described in this paper was funded by OSD/Army SBIR contract DAMD17-03-C-0046.

### 5.2 Disclaimer

The opinions and assertions expressed in this paper are those of the authors and do not necessarily express the official views of the U.S. Department of the Army or the U.S. Department of Defense.

"This paper has been approved for public release; distribution is unlimited."

## 6.0 REFERENCES

[1] CCCRP, Combat Casualty Care Research Program, U.S. Army Medical Research and Materiel Command, http://www.usaccc.org/aboutccc.jsp.

[2] R.W. Hoyt, J. Reifman, T.S. Coster, M.J. Buller, "Combat Medical Informatics: Present and Future", Proceedings of the AMIA 2002 Annual Symposium, San Antonio, TX, USA, November 9-13, 2002, pp. 335-339.

[3] J. Pearl, Probabilistic Reasoning in Intelligent Systems, San Francisco, CA: Morgan Kaufmann Publishers, 1988.

[4] C.T. Savell, M. Borsotto, J. Reifman, R.W. Hoyt, "Life Signs Decision Support Algorithms", to be presented at the MedInfo 2004, San Francisco, CA, USA, September 7-11, 2004.

[5] Uniform Determination of Death Act, National Conference on Commissioners of Uniform State Law, the American Medical Association and the American Bar Association, approved October 19, 1980 and February 10, 1981, http://www.law.upenn.edu/bll/ulc/fnact99/1980s/udda80.pdf, 1980.

[6] Voluntary Consensus Guidelines for Determination of Death, New York State Department of Health, http://www.health.state.ny.us/nysdoh/bsd/guidelin.htm, 2003.

[7] Guidelines for Determination of Death by Irreversible Cessation of All Functions of the Entire Brain, Including the Brain Stem - Age Greater Than One Year, New York State Department of Health, http://www.health.state.ny.us/nysdoh/bsd/attacha.htm, 2003.