Predictive Health Monitoring Systems for the Management of Hypertension

O. Diegel*, H. Ryu*, T. Norris*, R. Stockdale*, D. Parsons*, K. Hawick*

*** Centre for Mobile Computing, Massey University, Auckland, New Zealand.**

O.Diegel@massey.ac.nz

ABSTRACT

This paper describes how a predictive health monitoring system was implemented in which the computer learns from its user's habits and then starts to make intelligent decisions based on this learning. It describes how pattern recognition is used to detect anomalies and trends in a patient's blood pressure, and how the knowledge that the system acquires of the patient's habits can be used to fine tune the system.

Keywords: predictive health systems; pattern recognition

1. Introduction

The need for Automated Health Monitoring Systems (AHMS) is growing, especially in developed nations, as demographic ratios are changing and elderly population percentages are increasing. Most advances in patient health monitoring have been in the development of new clinical measurements, or improvements in the processing of existing ones. Health monitoring systems can also be improved by automating the interpretation of the data they produce [1, 2]. Rather than simply displaying measurements for doctors to interpret, monitoring devices can assist both doctors and patients in the complex task of the interpretation itself. Monitoring devices can then be effectively used to help patients improve the management of their health. These advances are made possible through developments in the fields of signal processing, pattern recognition and artificial intelligence (AI).

There are many motivations for automating signal interpretation. The main motivation comes from difficulties doctors and patients face when they continuously monitor patient data. These human factors include the problems of information overload, varying expertise, and human error [3].

Doctors may currently have difficulty in interpreting information presented to them on current monitoring systems [4]. Not only is the amount of information available greater than can be easily dealt with by the mind, but the clinical environment often provides distractions, reducing the effort that can be devoted to data interpretation. Another problem is that current monitors often have a tendency to flood clinicians with false alarms and false readings, providing unnecessary distraction from the task at hand [5].

The level of expertise that individual doctors bring to a task like the interpretation of signals can vary greatly, and it may not always be possible for such deficits to be remedied by consultation with more skilled colleagues. This occasionally leads to errors in diagnosis or selection of treatment. This is even more so with the recent prevalence of home health monitoring systems in which the users may have minimal to non-existent medical backgrounds.

There are several ways in which computer systems can assist in the difficulties described above. One is to automate the process of data validation. At present it is up to the doctor to ascertain whether a measurement accurately reflects a patient's status, or is in error. While in many situations signal error is clear from the clinical context, it can also manifest itself as subtle changes in the shape of the data measured. Without quite specialized expertise, doctors and patients may misinterpret measured data as being clinically significant when it, in fact, reflects an error in the measurement system. For example, the position in which a blood pressure monitoring cuff is worn can significantly alter the measurements it produces.

The interpretations produced by a computerized monitoring system can be much more complex than just an assessment of signal validity. Much of the research in medical AI over the last two decades has been devoted to this area, and good diagnostic performances have been demonstrated in many medical domains [6].

In the case of home health monitoring, however, the main pattern that needs to be identified is any one which deviates from a preset norm, as this often indicates a change in health status. In many cases, it is a sudden change in blood pressure rather than just normally slightly high blood pressure reading, for example, that heralds a change in health. It is also useful to be able to detect gradual trends in measurements, as a gradual increase in blood pressure, for example, may be an indication of arterial disease. For a useful AHMS we are therefore concerned firstly with identifying these 'out of norm' patterns, and then in using what other available data there is, such as pulse, weight, blood sugar levels, arterial stiffness index, etc., to try and make intelligent decisions about them.

This paper describes such an automated home health monitoring system in which patients are able to monitor their vital signs, and the computer system makes decisions about abnormalities in these readings and communicates with the user about any necessary preventative steps to take.

2. Interpreting Patterns

Intelligent pattern recognition and interpretation can be divided into two tasks. Firstly, distinct events within a data set are identified using pattern recognition methods, e.g., detecting missed or abnormal blood pressure readings within a prescribed schedule. Secondly a meaningful label is assigned to the detected events using pattern interpretation methods, e.g., picking up an abnormally high blood pressure reading, and interpreting its clinical significance. There have been significant advances with techniques for performing both these tasks, and new methodologies have emerged, some specifically from research in AI. Some of the more significant methods include blackboard systems [7], Markov models, and neural networks.

The properties of neural networks make them useful both for pattern recognition, and data interpretation. The neural networks not only recognize a pattern, but are then able to associate it with a predetermined diagnostic class. While the interpretive facility of networks has found numerous applications, it can be limited by its inability to explain its conclusions. The reasoning by which a network selects a class is hidden within the distributed weights, and is therefore unintelligible as an explanation. Networks are therefore limited to interpreting patterns where no explanation or justification for selecting a conclusion is necessary.

Once patterns have been identified within a data set, they need to be interpreted. It is with this task that techniques from AI have made major contributions over the last two decades, especially through the introduction of expert systems.

An expert system is a program that captures elements of human expertise, usually in the form of situation or context recognition rules, and performs tasks that rely on specialist knowledge.

Expert systems perform best in straightforward tasks, which have a predefined and relatively narrow scope, and perform poorly with more general tasks that rely on general or common-sense knowledge [8]. This level of performance is well suited to home health monitoring systems as all that is required of it is to select one of a series of cases based on a combination of available data such as high blood pressure, weight, lung capacity and blood sugar levels, for example.

An expert system consists of a knowledge base which contains the rules necessary for the completion of its task, a working memory in which data and conclusions can be stored, and an inference engine which matches rules to data to derive its conclusions.

Clinically deployed expert systems already perform a variety of tasks from the interpretation of ECGs to analysis of laboratory results. Experimental expert systems have been developed with more ambitious goals in mind, including systems that can interpret respiratory parameters and automatically adjust ventilator settings during the process of weaning a patient off a ventilator.

One of the contributions of AI has been an increase in the understanding of the ways in which knowledge can be represented and manipulated. Rule-based representations of knowledge are only appropriate for narrowly defined problems like diagnosing high blood pressure. In real life, humans deal with a broader type of problem by invoking other types of knowledge than the rules-of-thumb stored in an expert system. Especially with difficult problems, humans may attempt to reason from first principles, using models of pathophysiology or biochemistry to explain a set of clinical symptoms. This contrasts with the simple structure of rules which record commonly seen patterns of disease, and which can only deal with interactions by explicitly enumerating them. The vast number of such interactions can make such an enumeration impractical [9].

Model-based systems are designed to utilize models of disease in order to cover a broader set of clinical problems than is possible with rules [10].

Model-based systems are generally perceived as being better at explanation than rule-based systems, and better at dealing with novel or complex problems. They are also however, more computationally expensive to run as it takes longer to reason a problem out from first principles than it does to simply recognize it from a previous case. A more practical system therefore combines the two systems, thus having the facility to invoke deep medical models should they be needed, but also being able to rely on simpler, more efficient rules when they are applicable. These combined models form the basis of knowledge based systems.

3. Smart Health Monitoring Systems

The smart health monitoring system described in this paper learns through pattern recognition. As it learns about the patient, it gradually classifies events into hourly, daily, weekly and monthly events. Each time classification has a user set threshold outside of which the event no longer falls into that category. Over time, as the system learns about the users' habits, the system gradually tightens these thresholds, thus allowing the user to increase their level of compliance.

Typical events that are monitored by the system are blood pressure readings, asthma monitor lung capacity readings, blood sugar level readings, weight, temperature and body fat levels. Each day the user goes through their daily routine, and the computer system logs all the measurements taken and, based on their pattern, classifies them into the relevant classification. Most of the above mentioned measurements, for example can be classified into hourly, daily or weekly even within a time span of between one week and one month.

The system described in this paper was tested with elderly patients since such patients are generally more susceptible to follow a daily routine, and this becomes an advantage when trying to establish regular usage patterns. They will, for example, have a tendency to monitor their blood pressure at the same time each day. This therefore allows the computer to relatively easily detect a late or missed reading, and if other data is available, can make an educated guess as to the possible cause for the missed or late reading.

Once events have been put into their classifications, the computer can make decisions about whether a measurement falls outside its threshold and can then make decisions based on this data. Over the course of time, the computer can also reduce the thresholds as more user data is collected.

The test smart blood pressure monitoring system developed as part of this project employed a pattern recognition algorithm, developed by Cios [11] based on the splitting method [12] for region extraction.

The pattern recognition algorithm and the rule based decision making system were implemented through Microsoft C# using a series of multiple level variables for each blood pressure readings. Each variable contained the following properties: Systolic blood pressure, Diastolic blood pressure, Pulse pressure, Mean arterial pressure, Pulse arrhythmia detection, Time and date of reading, and thresholds for each of the above

In addition to these multi-level variables, a set of time category variables was used. These time variables were initially setup as weekly, daily, and hourly. Each variable was also given an adjustable threshold to determine whether a reading would fall within that category or not.

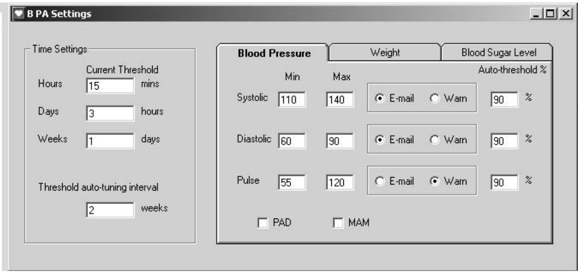

Figure 1: GUI screen allowing thresholds to be set and percentages of values that need to be within threshold before they automatically adjust based on the threshold auto-adjust time.

Thus, as a blood pressure was taken for the first time it was, by default put into the weekly category. If another reading was taken the next day, a comparison between the times of the two readings was performed to determine whether it remained in the week category or, if it was within the threshold for a daily reading, it was recategorized as a daily reading. In the same manner if readings were to be taken four times a day, each subsequent reading would be compared to the previous readings to see if it fell within the threshold of the hourly category and, if so, was categorized as such.

Intelligence was implemented into the system whereby, with the hourly readings, for example, the readings were looked at on a weekly basis (another user adjustable variable) and if the majority of the readings (what constituted a majority was determined by a second threshold variable) were below the hourly threshold, the threshold was reduced to the highest of these values, thus causing the system to become more precise. The few readings that were outside of the threshold, but had already been categorized as hourly readings, were flagged as exceptions to the rule and displayed as such on the graphical user interface (GUI).

Once blood pressure readings had been categorized into time categories, their values were compared both against each other and against the threshold set for each reading.

If, for example, only a systolic reading was outside of its acceptable threshold once, the user would be asked, through the GUI, whether they wanted to treat it as an exception. If, however, the systolic reading was high and the same reading had detected a pulse arrhythmia, the rule for that case might have the result of e-mailing the user's doctor. Ultimately, in an advanced system, the system might in such a case be able to prescribe an increase in medication dosage, for example.

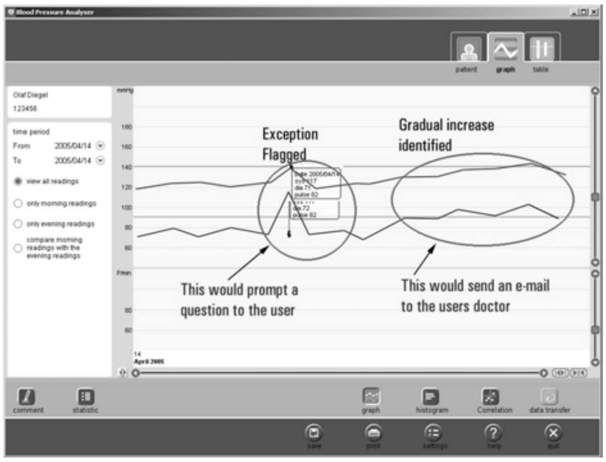

Figure 2: GUI screen showing identification of high readings to be treated as exceptions and identification of a rise in blood pressure over time.

If, on the other hand, the blood pressure readings were consistently high, and at the same time, the weight readings and blood sugar level readings were high, a rule would detect these abnormalities and might produce a result such as prescribing medication or advising the user to exercise.

In the test system, only the blood pressure readings were automatically uploaded from the blood pressure monitor. The other readings, such as weight and blood sugar levels, were entered manually.

4. Conclusions and Future Work

The intelligent health system described in this paper uses the generally regular habits of elderly users to gradually learn about their health measurement habits and then tightens the values of thresholds over time as it better learns their habits. This is of benefit to the user to improve their medication or health compliance rate.

The system then looks for patterns in the data gathered (currently based on blood pressure readings) to make decisions about their health. It currently prompts the user about what they want to do about any abnormal readings and allows for the option of abnormal results, as specified by the user thresholds, to be e-mailed to a specified e-mail address.

The system is also capable of accepting weight and glucose level measurements but these are not yet automated. This is however currently being implemented.

The system, though implemented at a relatively low level demonstrates how a health monitoring system can be used to gradually learn about the habits of its users and thus make better decisions

Most of its intelligence is rule based because of the scarcity of data it currently accepts. The blood pressure readings are fully automated and downloaded directly from the blood pressure monitor, but all other data is still entered manually.

The system is currently being improved to accept readings directly from an asthma monitor, a glucose monitor and a scale in order to automate much of the data gathering.

This extra data will also allow the system to begin using a hybrid decision making system which will be both rule and model based. Evaluation of the hybrid systems will then provide valuable information on the design and application of predictive health monitoring systems

5. References

[1] E. Coiera, Editorial: *Intelligent Monitoring and Control of Dynamic Physiological Systems*, Artificial Intelligence in Medicine, 5, 1-8, 1993

[2] E. Coiera, *Automated Signal Interpretation*, in P. Hutton, C. Prys-Roberts (eds) Monitoring in Anaesthesia and Intensive Care, Bailliere Tindall Ltd., London, 32-42, 1994

[3] C. D. Wickens, *Engineering psychology and human performance*, Harper Collins, New York, 1992

[4] M. B. Weigner, C. E. Englund, *Ergonomic and human factors affecting anesthetic vigilance and monitoring performance in the operating room environment*, Anesthesiology, 73, 5, 995-1021, 1990

[5] E. Koski, A. Makivirta, T. Sukuvaara, A. Kari, *Frequency and reliability of alarms in the monitoring of cardiac postoperative patients*, International Journal of Clinical Monitoring and Computing, 7, 129-133, 1990

[6] W.J. Clancey, E.H. Shortliffe, (eds.), *Readings in Medical Artificial Intelligence - The First Decade*, Addison-Wesley, Reading MA, 1984

[7] H. P. Nii, *Blackboard Systems*, in A. Barr, P. Cohen, E. A. Feigenbaum (eds), The Handbook of Artificial Intelligence, Vol. IV, Addison-Wesley, Reading, Ma., 1-82, 1989

[8] E. Coiera, *The Qualitative Representation of Physical Systems*, Knowledge Engineering Review, 7,1, 55-77, 1992

[9] E. Coiera, *Monitoring diseases with empirical and model-generated histories*, Artificial Intelligence in Medicine, 2, 135-147, 1990

[10] S. Uckun, *Model-Based Reasoning in Biomedicine*, Critical Reviews in Biomedical Engineering, 19, 4, 1992

[11] K.J. Cios, L.S. Goodenday, K.K. Shah and G. Serpen, *A novel algorithm for classification of SPECT Images of a human heart*. In:Proc. Of the CBMS'96, Ann Arbor, 1996

[12] Anzai Y., (Ed.), *Pattern Recognition and Machine Learning*, Academic Press, 1992