HEWLETT PACKARD

Automated Signal Interpretation

Enrico Coiera
Information Management Laboratory
HP Laboratories Bristol
HPL-93-30
April, 1993

patient monitoring, signal processing, artifact rejection, smart alarms, anesthesia, artificial intelligence, expert systems, neural networks, model based reasoning

Most advances in monitoring have been associated with the development of new clinical measurements, or improvements in the processing of existing ones to enhance their information content or validity. It is also possible to enhance monitoring systems in quite a different way, by automating the interpretation of signals. Rather than simply displaying measurements for clinicians to interpret, monitoring devices can assist clinicians in the task of interpretation itself. Such advances are made possible through developments in the fields of signal processing, pattern recognition and artificial intelligence (AI). This paper presents an introduction for clinicians to the basic problems that need to be solved, and reviews the major technologies available at present.

1 Introduction

Most advances in monitoring have been associated with the development of new clinical measurements, or improvements in the processing of existing ones. It is also possible to enhance monitoring systems in quite a different way, by automating the *interpretation* of signals [11]. Rather than simply displaying measurements for clinicians to interpret, monitoring devices can assist clinicians in the task of interpretation itself. Such advances are made possible through developments in the fields of signal processing, pattern recognition and artificial intelligence (AI).

1.1 The need for automated interpretation

The motivations for automating signal interpretation are numerous. The most pressing arise from the difficulties clinicians face when they continuously monitor patient data, and are not unique to medicine. They are also an issue for example, in the design of systems used by airline pilots and nuclear power plant operators. These human factors include the problems of data overload, varying expertise, and human error [50].

It comes as no surprise that clinicians may have difficulty in interpreting information presented to them on current monitoring systems [49]. Not only may the amount of information available be greater than can be assimilated, but the clinical environment provides distractions with other tasks, reducing the effort that can be devoted to signal interpretation. Worse still, current monitors flood clinicians with false alarms, providing further unnecessary distraction [29].

It is also clear that the level of expertise that individuals bring to a task like the interpretation of signals varies enormously, and it is not always possible for such deficits to be remedied by consultation with more skilled colleagues. This frequently leads to errors in diagnosis and selection of treatment. Indeed the majority of complications associated with anaesthesia result from inadequate training or insufficient experience of the anaesthetist [12][46].

There are several ways in which computer based systems can assist in addressing such difficulties. One is to automate the process of *data validation*. At present it is up to the clinician 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 a waveform. Without quite specialised expertise, clinicians may misinterpret measured data as being clinically significant, when it in fact reflects an error in the measurement system. For example, changes in the height of a pressure transducer can significantly alter the measurements it produces.

The interpretations produced by a computer can be much more complex than an assessment of signal validity. It is also possible to design systems capable of diagnosing clinical conditions, to assist with identifying rare or complex cases [36]. Much of the research in medical artificial intelligence over the last two decades has been devoted to this area, and impressive diagnostic performances have been demonstrated in many specialised medical domains [5].

1.2 System Requirements

Before automated interpretation can begin to provide clinical benefit, there are several requirements that must be met. To ensure that monitored parameters are interpreted in clinical context, one may need access to clinical data other than the monitored signals themselves. This data may include the medical record, current medications, and values from other devices like the settings from an anaesthetic machine. Thus the first step in providing significant signal interpretation is to collate as many sources of clinical information as possible and present them in a uniformly accessible manner. Centralised anaesthetic record systems and anaesthetic workstations seek to do just this [19][31].

Along with the technical aspects of interpreting signals, there is an equally important human aspect to system design. It is essential that any system developed actually fulfils a relevant clinical role. The long lag in the introduction of computerised decision support into medicine is more probably due to failure on this point than because of technological limitations [43]. Systems must be developed to fit in with the work practices of clinicians [14], and to support decision making processes that are clinically relevant [11]. There is little advantage in developing a complex system that mimics interpretative skills already possessed by all clinicians. Rather, it should attempt to provide support for cognitive functions that clinicians perform poorly. Thus the development of intelligent systems is as dependent on developing an understanding of the cognitive patterns of the clinicians who will work with them as it is on advances in technology.

1.3 Stages of Development

Pragmatic considerations suggest that there will be a staged introduction of software capable of intelligent interpretation. Initial systems will be relatively simple, and require minimal interaction with the clinician, and minimal or no interaction with the patient record system. Most of these are likely to be *embedded*, residing within clinical devices. Such programmes exist within patient monitoring devices, filtering out artefacts and suppressing false alarms, or in laboratories where they interpret biochemical assays and produce reports that are sent to clinicians [26].

The next level of interpretive system will require *explicit interaction* with clinicians, offering some form of active decision support. Such systems will come into their own as integrated electronic patient record systems appear and the medical profession becomes accustomed to computer assistance. They will assist in the selection of tests and diagnosis, as well as the selection of optimal therapies. Interaction will be necessary because, although such systems are capable of drawing conclusions from patient data, they will not be privy to the complete clinical picture - the clinician must supply vital clinical context and therapeutic goals unavailable to the system.

The third stage of system introduction could consist of *autonomous intelligent systems*, capable of independent activity. These are at present experimental, but could eventually form the heart of closed-loop systems. For example, drug delivery systems could automatically measure a drug's level and administer doses based upon that measurement [1][35]. The development of closed loop systems is at present hampered as much by legal and ethical issues as it is by technological considerations.

2 Levels of interpretation

Signal interpretation can vary from a low level assessment of validity to a complex assessment of clinical significance. The levels of interpretation that a signal passes through are illustrated in Figure 1. A signal is first examined for evidence of artefact and the validated signal is then presented to the next layer in the interpretive hierachy. Where a single channel signal contains sufficient information for a diagnosis, this is made. In some circumstances however, several alternative explanations might be possible, and a single channel does not contain enough information to disambiguate them. In such circumstances, an interpretive system can look for cross signal correlations. In Figure 1., a flat portion of ECG trace is not diagnosed as an 'asystole' because examination of the corresponding arterial waveform reveals pulsatile behaviour consistent with normal cardiac function. A higher level of interpretation is also possible, taking into account relevant contextual patient information where this is available. This level is concerned with making decisions based upon signal interpretations, and may include recommendations for further investigations or therapeutic actions. The tasks of artefact detection, single and cross-channel interpretation and decision support will be examined in more detail below.

2.1 Artefact Detection

The first task in signal interpretation is to decide whether the values that are measured are physiologically valid. In other words, is the signal genuine or is it artefactual? An artefact is defined as a component of the measured signal that is unwanted. It may be caused by noise on the signal, or by distortions introduced through the measurement apparatus. Indeed, an artefact may be due to another physiological process that is not of interest in the current context - like a respiratory swing on an ECG trace. Thus *one man's artefact is another's signal* [39].

Artefact detection is important for several reasons. Firstly, an artefact may be misinterpreted as a genuine signal and lead to an erroneous therapeutic intervention. Next, invalid but abnormal values that are not filtered can cause alarm systems to register false alarms. Finally, artefact rejection improves the clarity of a signal when it is presented to a clinician for interpretation.

There are many sources of artefact in the clinical environment. False heart rate values can be generated by diathermy noise during surgery and by patient movement, and false high arterial blood pressure alarms are generated by flushing and sampling of arterial lines (Figure 2.). These forms of artefact have contributed significantly to the generation of false alarms on monitoring equipment. Koski et al. [29] found that only 10% of 1307 alarm events generated on cardiac postoperative patients were significant. Of these, 27% were due to artefacts e.g. sampling of arterial blood. The net effect of the distraction caused by these high false alarm rates has been that alarms have often been turned off intraoperatively, despite the concomitant increase in risk to the patient.

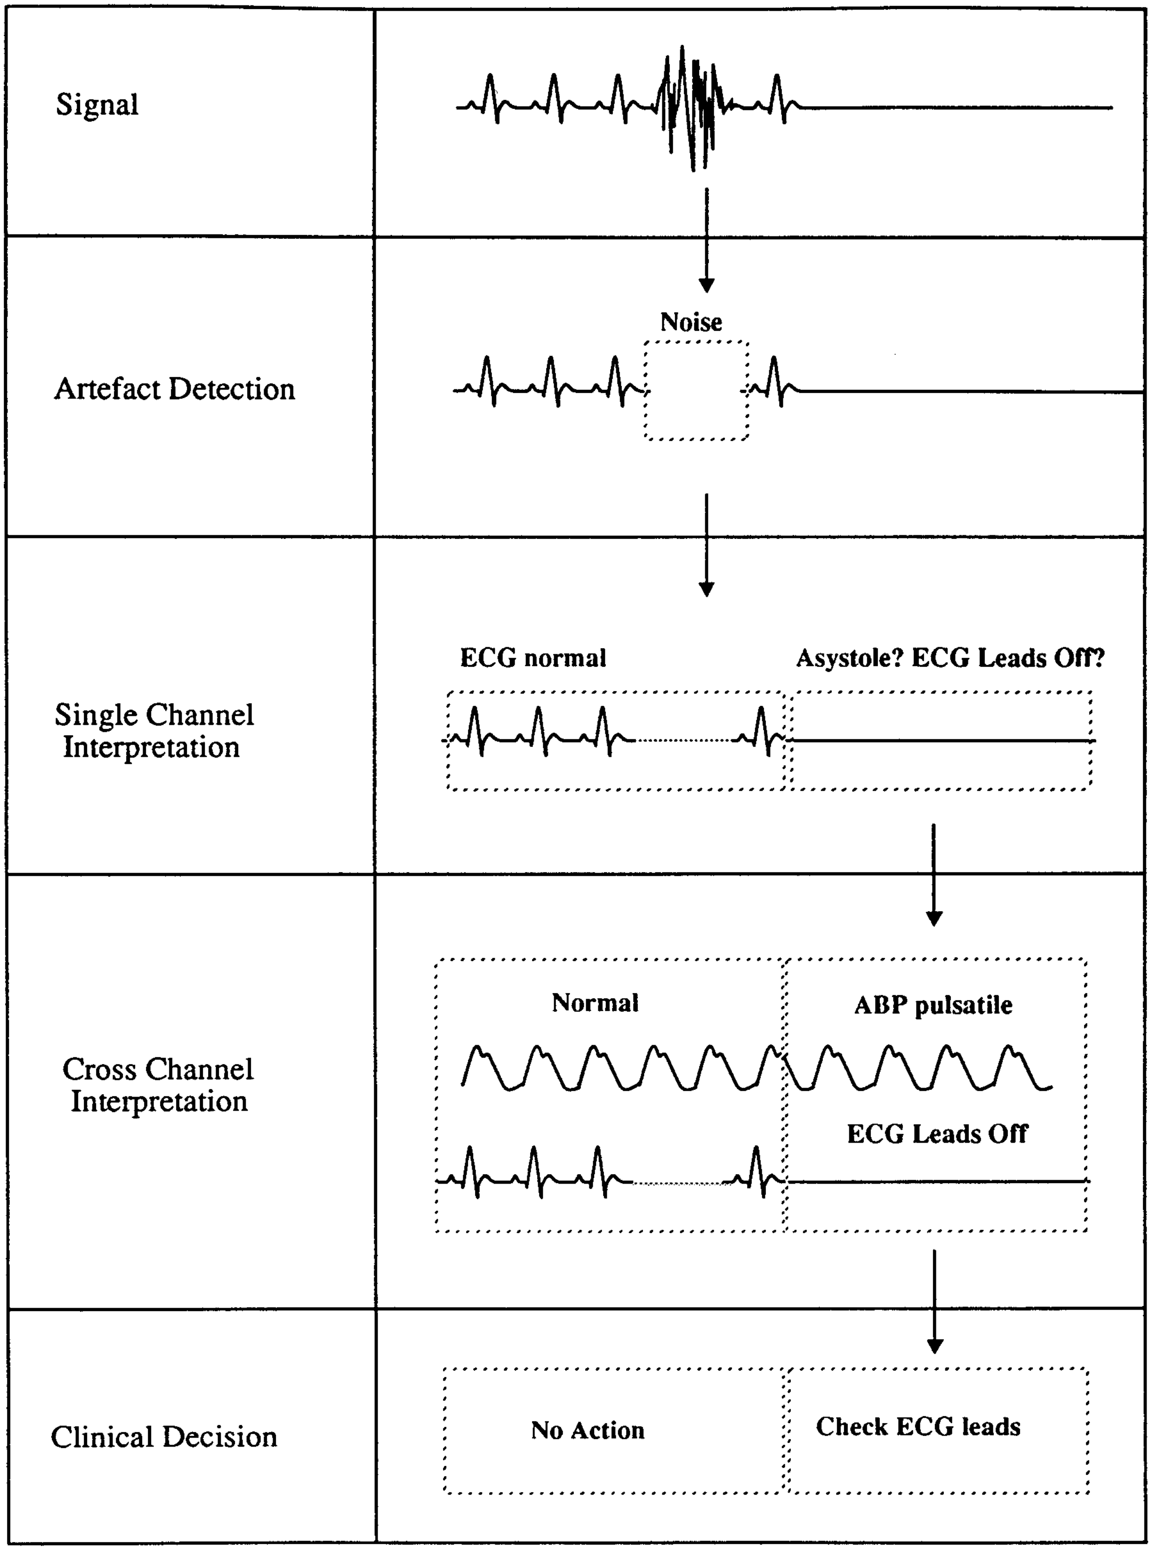

Figure 1. Increasing levels of interpretation from presentation of a physiological signal to the formulation of a clinical decision

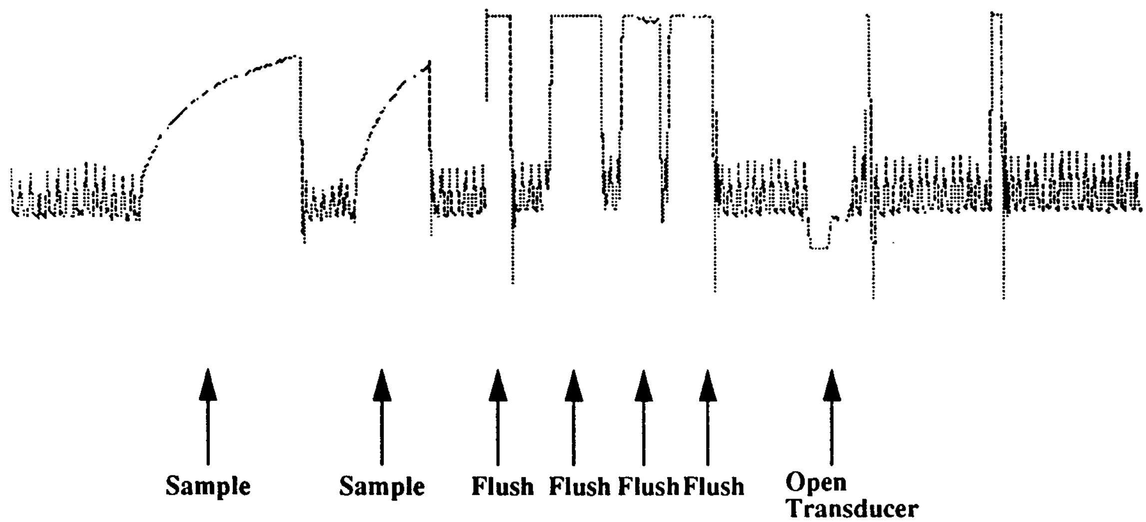

Figure 2. Examples of Artefact on the Arterial Blood Pressure Channel

While an artifact is best handled at its source through improvements in the design of the transducer system, or in low level signal processing, it is not always possible or practical to do so. The next best step is to filter out artefactual components of a signal or register their detection prior to using the signal for clinical interpretation. Many techniques have been developed to assist in this process [39] and include kalman filtering [44], rule-based expert systems [17], blackboard systems [2], and neural networks [41]. It is in the nature of artefact that it cannot always be eliminated on the basis of a single signal, and cross-channel correlations may be needed, making artefact detection a feature at all levels of signal interpretation.

2.2 Single Channel Interpretation

Having established that a signal is probably artefact free, the next stage in its interpretation is to decide whether it defines a clinically significant condition. This may be done simply be comparing the value to a predefined patient or population normal range, but in most cases such simple thresholding is of limited value. Firstly, clinically appropriate ranges cannot always be defined because the notion of the acceptable range for a patient is highly context specific [6]. One can in fact calculate statistically valid patient specific normal ranges [21][22] but these rely on a period of stability which may not be attainable in a dynamic clinical context. Further, the notion of an acceptable range is often tied up with expectations defined by the patient's expected outcome and current therapeutic interventions. Finally, even if one can decide upon an acceptable range for a specific parameter, the amount of information that a single out of range warning can convey is usually limited - even wildly abnormal values may have several possible interpretations. These limitations of simple threshold based alarm techniques have spurred on the development of more complex techniques capable of delivering 'smart alarms' [19].

Much more information can be obtained from the analysis of a single channel if is a time varying and continuous waveform, like arterial pressure. Specific pressure

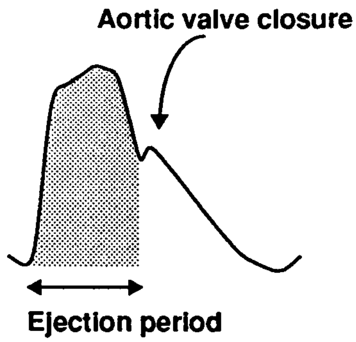

Figure 3. (a) Derivation of clinical information from waveform shape characteristics. The systolic ejection area beneath the arterial pressure waveform gives an indirect measure of stroke volume. It is demarcated by the beginning of systole and the dicrotic notch caused by the closure of the aortic valve.

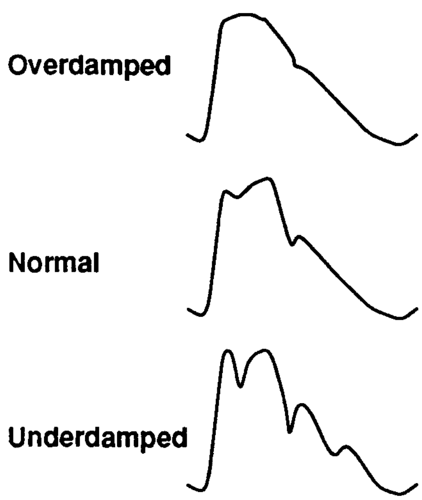

Figure 3. (b) The shape of arterial pressure waves varies with the dynamic response of the catheter transducer system. Analysis of the waveform frequency components following a fast flush can assist in detecting damping of the system, and hence assist in optimising pressure measurements (Gardner, 1981).

artefacts like sampling and flushing of the catheter line can be detected by their unique shape (Figure 2.). Estimates of clinically useful measures like stroke volume can be derived by analysing the area under the curve of the wave. It is even possible to analyse the frequency components of the pressure waveform to obtain information about the fidelity of the measurement system itself (Figure 3.).

Alterations in the behaviour of a repetitive signal can also carry information. Changes in the ECG are a good example. Features such as the height of the QRS peak help to label individual components within beat complexes. The presence or absence of features like P waves, and the duration and regularity of intervals between waves and complexes can carry diagnostic information about cardiac rhythm. However, it is not always possible to unambiguously label events in an ECG strip, and sometimes one needs to use additional contextual information to assist in the labelling process (e.g. [20]).

2.3 Cross-channel Interpretation

Often conditions can only be identified by examining the signals on several different channels. Such cross channel information is useful at several levels, starting with artefact detection and signal validation through to clinical diagnosis.

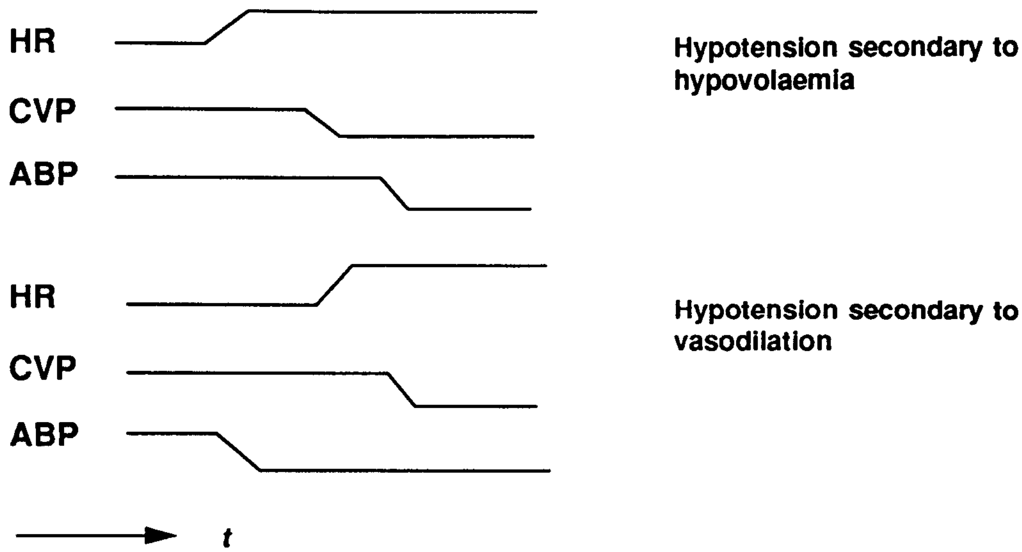

Figure 4. Time ordering of the onset of trends across different signals can help identify the onset of different clinical conditions. The cause of hypotension can be distinguished by the sequence of changes detected on heart rate (HR), central venous pressure (CVP) and arterial blood pressure (ABP).

Cross-correlation for signal validation can be made with a number of sources, depending on the signal being measured. The alternatives include correlating a signal with:

1. *Same Signal, different interpretation method.* If for example, an error is suspected when heart rate is derived by a simple peak detection algorithm, one could attempt to validate the value by comparing it to one derived using a different method on the same data.
2. *Different physical source, but same signal.* Comparing different ECG leads is a common technique for validating changes seen on one lead. Patterns across leads also have diagnostic importance.
3. *Different signal.* A flat ECG trace indicating asystole can be checked against the arterial pressure waveform or the plethysmograph, both of which should demonstrate pulsatile waves if the heart were contracting normally, and would loose their pulsatile characteristics if asystole was present (Figure 2.).

Cross channel information can also be used to identify conditions not detectable with single signal channels alone. Many clinical conditions can be distinguished by the time ordering of events in their natural history [7]. For example, the cause of a hypotensive episode may be deducible from the order in which changes occurred across heart rate, blood pressure and CVP (Figure 4.). In the presence of a vasodilator, the arterial blood pressure drop would precede the reflex tachycardia and CVP fall. In the presence of hypovolaemia, the first parameter to shift would be the heart rate followed by CVP and blood pressure.

2.4 Making Clinical Decisions

Once a computer system is able to diagnose complex disease patterns from measured signals and data stored in electronic patient records, it is in a position to assist clinicians in making therapeutic decisions. As noted earlier, intelligent interpretive systems will appear both as embedded systems hidden within instruments, and as explicit entities which can interact with a clinician.

The way in which an intelligent system is used affects its design [8]. Systems that need to interact with humans may need to justify their decisions in a way that an embedded system does not. Equally, a system that acts as an advisor to a human is placed in a less critical position than one that acts independently to manage a patient's therapy. While embedded interpretive systems are already starting to appear, those that require explicit interaction have yet to make a significant impact.

2.4.1 Decision Support Systems

The classic model of computer based decision support requires a clinician to input details of a patient's clinical state, with the machine then suggesting one or more possible diagnoses. The MYCIN system [3] is the archetype of this model, providing assistance with the selection of antibiotic therapy. Other examples include programs which assist in diagnosing abdominal pain [13] and chest pain [18]. In practice however, this model does not fit well with the realities of the clinical workplace. Clinicians are often unable to spend the time required to use such systems, and if they do, the types of problems that they would like assistance with are somewhat different [11]. As a consequence, systems which offer different models of decision support have been developed.

To support the decision making that is characteristic of anaesthetists in the operating room, work is currently underway to develop intelligent patient monitors [7] and anaesthetic workstations (e.g. [31]). Intelligent monitors will not only suppress spurious alarms generated by signal artefact, but will use cross-channel signal correlations to generate high level diagnostic alarms. They have a role in assisting with clinical vigilance of slowly evolving conditions, and of conditions which have been missed because of distraction. When integrated with the anaesthetic machine itself, the monitor system will also warn of faults within the gas delivery system, and possibly suggest corrective actions. Research is also underway exploring ways of integrating a predictive component into these systems. With such systems, a clinician could test changes to therapy before initiating them by simulating their effect on a mathematical model of the patient[40].

Other modes of decision support include therapy planning systems, which develop a treatment protocol based on a patient's clinical status [25], and therapy critiquing systems which examine treatment plans generated by clinicians, and attempt to suggest improvements [32].

2.4.2 Autonomous Therapeutic Devices

In contrast to decision support systems, autonomous systems can operate independently of human interaction on complex tasks. Such systems are at the present moment purely vehicles for research. Perhaps the most studied application area is ventilator management. Programs have been designed to automatically adjust ventilator settings in response to measurements of a patient's respiratory status. Early research into systems that could wean patients from ventilators [15] has lead to more ambitious projects that seek to take control of most of the tasks associated with ventilator management (e.g. [24]). While such projects may in the long term provide new classes of therapeutic devices, it is clear that their successful introduction will require continued advances in sensor design (e.g. implantable glucose sensors for insulin delivery systems), and in the technologies for signal interpretation.

3 Methods of Interpretation

Intelligent signal interpretation can be divided into two tasks. Firstly, distinct events within a signal are identified using *pattern recognition* methods e.g. detecting individual peaks in an ECG signal. Secondly a meaningful label is assigned to the detected events using *pattern interpretation* methods e.g. picking a QRS complex from a T wave, and interpreting its clinical significance. There have been significant advances with techniques for performing both these tasks, and new methodologies have emerged, several specifically from research in AI. Some of the more significant methods are reviewed below.

3.1 Pattern Recognition

Pattern recognition techniques are those that extract significant events from a signal. For example, they detect edges and curves in pictures, or letters, letter groups and words from speech. Pattern detection techniques vary in the way they model events within signals. For example they may be based on statistical models, in which the frequency of certain patterns is used in the recognition process. There are many classic recognition techniques that have clinical application - like blackboard systems [33] (initially developed for speech recognition) and markov models - but these will not be reviewed here. Neural Networks are another technique useful in pattern recognition, and they have become of increasing interest both to the medical and the AI community in recent years.

3.1.1 Neural Networks

Neural networks are based upon a simple computational model of the neuron [28]. Networks are composed of layers of neurons (or nodes) with interconnections between the nodes in each layer (Figure 5.(a)). The strength of the connections between nodes is modelled as a weight on that connection. A node in the network fires when the sum of its inputs exceeds a predetermined threshold (Figure 5(b)).

When a net is presented with a pattern on its input nodes, it will output a recognition pattern determined by the weights on the connections between layers. These weights are obtained by a period of training, in which a net is presented with examples of the signal patterns it is intended to recognise, and the weights in the net are slowly adjusted until it achieves the desired output. A neural network thus encodes within its weights a discriminating function that is optimised to

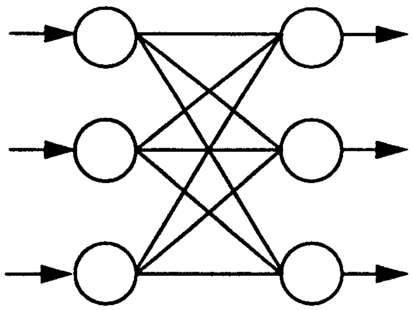

Figure 5. (a) A two layer neural network. Input nodes on the left receive a pattern to be classified, and output nodes on the right are triggered to produce a classification. The sum of the signals received at an output node determine whether it will fire or not.

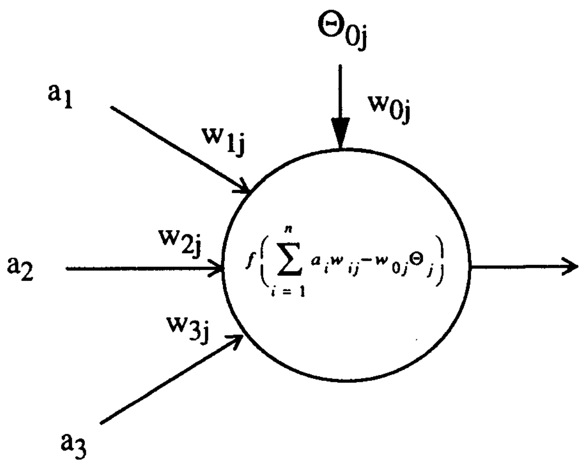

Figure 5(b) Structure of a typical node in a neural network. Inputs to the node (a_1 .. a_n) are assigned weights (w_1 .. w_n) and these weighted inputs are summed. Sometimes a threshold Θ is specified. When the sum exceeds the node threshold Θ, an output is produced. The function *f* determines the shape of the node output.

distinguish the different classes present within its training set. In theory, any such discriminant function can be approximated by a network [45].

Neural networks have been used to recognise ECG patterns [37], identify artefacts in arterial blood pressure signals [41], image recognition (e.g. ultrasound [34]) and in the development of clinical diagnostic systems [23].

Despite initial claims of uniqueness for the computational properties of neural nets, it is becoming clear that they have clear and important relationships with a number of more traditional discrimination methods including markov models [1], bayesian networks, and decision trees.

The properties of neural networks make them useful both for pattern recognition, and signal interpretation. The net not only recognises a pattern, but is able to associate it with a predetermined diagnostic class. While the interpretive facility of nets has found numerous application, it is limited by its inability to explain its conclusions. The reasoning by which a net selects a class is hidden within the distributed weights, and is unintelligible as an explanation. Nets are thus limited to interpreting patterns where no explanation or justification for selecting a conclusion is necessary. Since the need to justify a clinical diagnosis is recognised as an important part of the process of decision support, this limits the application of nets in such tasks.

3.2 Interpreting Patterns

Once patterns have been identified within a signal, 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.

3.3 Rule-based Expert Systems

An expert system is a program that captures elements of human expertise, usually in the form of situation recognition rules, and performs tasks that rely on specialist knowledge. Examples include programmes that can diagnose the cause of abdominal or chest pain based on clinical observations fed to the programme. 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 commonsense knowledge [9].

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.

Examples of rules which might be used to detect asystole and filter out false asystole alarms in the presence of a normal arterial waveform might be:

Rule ASY1:

> **If** heart rate = 0
> **then** conclude asystole

Rule ASY2:

> **If** asystole
> **and** (ABP is pulsatile and in the normal range)
> **then** retract asystole

In the presence of a zero heart rate, the expert system would first match rule *ASY1* and conclude that asystole was present. However, if it next succeeded in matching all the conditions in rule *ASY2* - that it had previously detected an asystole but could also detect a normal arterial waveform, then it would fire this second rule, which would effectively filter out the previous asystole alarm. If rule *ASY2* could not be fired because the arterial pressure was abnormal, then the initial conclusion that asystole was present would remain.

Rules tend to become much more complicated than the simple examples presented here, and the process of manual knowledge acquisition from human experts can become a drawn out affair. To counter this problem, much work has gone into developing techniques to automate the acquisition of knowledge in the form of rules or decision trees from databases of cases (e.g. [38]).

Rule based systems are more suited to cross-channel interpretation of signals than lower level signal processing. Clinically deployed expert systems perform a variety of tasks from the interpretation of ECGs [20] to analysis of laboratory results such as thyroid hormone assays [26]. 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 [15].

3.3.1 Model-based Systems

One of the important contributions of AI has been a growing understanding of the ways in which knowledge can be represented and manipulated. Rule-based representations of knowledge, as we have seen, are only appropriate for narrowly defined problems like diagnosing chest pain. Humans deal with a broader class of problems by invoking other types of knowledge than the rules-of-thumb that are typically stored in an expert system. Especially with difficult or rare problems, humans may attempt to reason from first principles, using models of pathophysiology or biochemistry to explain a set of clinical manifestations. For example, when several diseases are present at one time, it may only be possible to unravel the constellation of symptoms and signs by recourse to disease models. 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 makes such an enumeration impractical [7].

Model-based systems (sometimes called second-generation expert systems) are designed to utilise disease models in the hope that they will be able to cover a broader set of clinical problems than possible with rules [47]. These models may exist as mathematical descriptions of physiological relationships, as compartmental system models, or indeed as statistical models. As is often the case in medicine, formal models of disease phenomena are often not available or poorly formalised. In such cases, there is evidence that clinicians carry around looser models, expressible in non-numeric or qualitative terms [30]. Such qualitative representations of medical knowledge have been formalised, and can be used to capture useful portions of medical knowledge [9]. These representations find their uses in diagnosis (e.g. [27]), and patient monitoring (e.g. [7][48]).

Model-based systems are perceived as being better at explanation than "shallower" rule-based systems, and better at dealing with novel or complex problems. They are also however, more computationally expensive to run - it takes longer to reason a problem out from first principles than it does to simply recognise it from previous experience. Thus there is a move amongst researchers to build systems which combine the two sorts of system, having on the one hand the facility to invoke deep pathophysiological models should they be needed, but also being able to rely on efficient rules whenever they are applicable.

4 Limitations to Interpretation

While signal interpretation systems can perform at clinically acceptable levels, as with humans there are inherent difficulties in the reasoning process. These limitations need to be made explicit, and should be borne in mind by clinicians who use automated interpretation systems.

4.1 Data

Interpretive systems do not have eyes or ears, but are limited to accessing data provided to them electronically. While this constitutes a potentially enormous amount of data to work with, it does mean that critical pieces of contextual information may be unavailable. Thus interpretations need to be judged partly on the data available to the system when making its decision. This highlights the importance of designing an explanatory facility into expert systems, so that clinicians can understand the reasoning behind a particular recommendation by tracing the pieces of data that were used in its formulation.

As we have seen, the task of validating data is one of the first tasks that an interpretive system undertakes. While there are many clues available to suggest whether a datum represents a real measurement or is an error, this is not always decideable based solely on the electronic evidence. There may be no way for a machine to decide that a transducer is incorrectly positioned, or that blood specimens have been mixed up. Clinicians will always need to be wary of the quality of the data upon which interpretations have been made.

4.2 Knowledge

Knowledge is often incomplete, and this is an everyday reality in the practice of medicine. Clinicians deal with physiological systems they only incompletely understand, and have evolved techniques for dealing with this uncertainty. While a clinician is able to acknowledge that he is performing at the edge of his expertise, and adjust his methods of handling a problem accordingly, it is much harder to incorporate such a facility in a computer system. Computers at present treat all knowledge equally. While they are able to weight up probabilities that a set of findings represent a particular condition, they do not take into account they likelihood that some pieces of knowledge are less reliable than others.

Further, most present day systems are forced to utilise a static knowledge base. While there are many techniques which can be used to update knowledge bases, it will not be the case that a system necessarily incorporates the latest knowledge on a subject. Further, the technical problems associated with the process of knowledge acquisition mean that there are always potential mistakes in the system. Just as a normal computer program can contain "bugs", so a knowledge base can contain errors, since it is simply another form of program. Wherever possible, the explanation offered by a system should be examined, to ensure that the logical flow of argument reflects current clinical understanding.

5 Conclusion

It is already the case that much monitoring technology is poorly understood by the clinicians who use it, and that clinicians are often unaware of how to use or interpret their output correctly. As more intelligence is added to the devices that populate the clinical workplace, there is an even greater need to understand their advantages and limitations. There is no doubt that advances in AI will gradually change many of the ways clinicians handle day to day problems and that they will greatly improve many aspects of the clinical process. Of necessity, the high level at which these systems will perform, assisting both in diagnosis and therapy means that they have a direct impact on patient care. It will often only be the clinician who will be in a position to assess the conclusions of these systems. Rather than accepting these as a given however, the onus remains on those who use them to do so correctly. While it will not always be necessary to understand the details of the technologies used, and indeed these will continue to evolve, it is necessary to understand their nature and in particular the types of mistake that they are prone to make.

Further Reading

Artificial Intelligence

E. Charniak, D.V.McDermott, *Introduction to Artificial Intelligence*, Addison-Wesley, Reading MA, (1985).

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

R.S.Patil, Artificial Intelligence Techniques for Diagnostic Reasoning in Medicine, in H. Shrobe (ed), *Exploring Artificial Intelligence: Survey Talks from the National Conferences on Artificial Intelligence*, 347-380, Morgan Kaufmann, San Mateo, (1988).

G.Rennels, P.L.Miller, Artificial Intelligence Research in Anaesthesia and Intensive Care, *J. Clin. Monit.*, 4, (1988), 274-289.

Expert Systems

B.G.Buchanan, E.H.Shortliffe (eds.) *Rule-Based Expert Systems: the MYCIN experiments of the Stanford Heuristic Programming Project*, Addison-Wesley, Reading MA, (1984).

B.G.Buchanan, R.G.Smith, Fundamentals of Expert Systems, in A. Barr, P.R.Cohen, E.A. Feigenbaum (eds.), *The Handbook of Artificial Intelligence*, Vol. IV, 149-192,Addison-Wesley, Reading MA, (1989).

Neural Networks

T. Kohonen, An Introduction to Neural Computing, *Neural Networks*, 1,(1988), 3-16.

References

[1] J. A. Blom, Expert control of the arterial blood pressure during surgery, *International Journal of Clinical Monitoring and Computing*, 8, (1991), 25-34.

[2]M. Beech, S. Todd, V. Tombs, Knowledge-based techniques for alarm rationalisation in patient monitoring, *Proc. IEE Colloquium on AI in Medical Decision Making*,No 1990/062, (1990),

[3]B.G.Buchanan, E.H.Shortliffe (eds.) *Rule-Based Expert Systems: the MYCIN experiments of the Stanford Heuristic Programming Project*, Addison-Wesley, Reading MA, (1984)

[4]H. Bourland, C. Wellekens, Links between markov models and multilayer perceptrons, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12, (1990),1167-1178.

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

[6] E. Coiera, Modelling Clinical Expectation Over Time, *Proceedings of the Third Australian Conference on Applications of Expert Systems*, Sydney, (1987), 103-123.

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

[8] E. Coiera, Incorporating user and dialogue models into the interface design of an intelligent patient monitor, *Medical Informatics*, 16, (1991), 331-346.

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

[10]E. Coiera, Intermediate Depth Representations, *Artificial Intelligence in Medicine*, 4, (1992), 431-445.

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

[12]J.B. Cooper, R. S. Newbower, R. J. Kitz, An analysis of major errors and equipment failures in anaesthesia management: considerations for prevention and detection, *Anesthesiology*, 60,1, (1984), 34-42.

[13] F.T. de Dombal, D. J. Leaper, Computer-aided diagnosis of acute abdominal pain, *BMJ*,2, (1972), 9-13.

[14]D. Fafchamps, C. Young, P. Tang, Modelling Work Practices: Input to the design of a Physician's Workstation, *Proceedings of 15th Symposium on Computer Applications in Medical Care*, (1991).

[15] L. Fagan, E. H. Shortliffe, B. Buchanan, Computer-based Medical Decision Making: From MYCIN to VM, in J. Clancey, E. H. Shortliffe (eds), *Readings in Medical Artificial Intelligence* - The First Decade, 241-255,Addison-Wesley, Reading MA, (1984).

[16]R.M.Gardner, Direct Blood Pressure Measurements - Dynamic Response Requirements, *Anesthesiology*, 54, (1981), 227-236.

[17]D. Garfinkel, P. Matsiras, T.Mavrides, J. McAdams, S. Auckburg, Patient Monitoring in the Operating Room: Validation of Instrument Reading by Artificial Intelligence Methods, *Proceedings of 13th Symposium on Computer Applications in Medical Care*, (1989), 575-579.

[18]L. Goldman, E. F. Cook, et al., A computer protocol to predict myocardial infarction in emergency department patients with chest pain, *NEJM*, 318, (1988), 797-803.

[19]J. Gravesnstein, R. Newbower, A. Ream, N. Smith, (eds.), *The Automated Anaesthesia Record and Alarm Systems*, Butterworths, Boston, (1987).

[20]S. Greenwald, R. Patil, R. Mark, Improved Detection and Classification of Arrhythmias in Noise-Corrupted Electrocardiograms using Contextual Information, *Computers in Cardiology*, (1990).

[21]E. K. Harris, On the use of statistical models of within-person variation in long term studies of healthy individuals, *Clinical Chemistry*, 26, (1980), 383-391.

[22] E. K. Harris, Use of Statistical Models to detect subject specific changes, in *Progress in Health Monitoring - Proceedings of The International Conference on Automated Multiphasic Health Testing and Services*, T. Yasaka (ed), Excerpta Medica, Amsterdam (1981).

[23]A. Hart, J. Wyatt, Connectionist models in medicine: an investigation of their potential, in J. Hunter, J. Vookson, J. Wyatt (eds), *Lecture Notes in Medical Informatics*, volume 38, 115-124, Springer-Verlag, (1989).

[24]B. Hayes-Roth, R. Washington, D. Ash, et al., Guardian: A prototype intelligent agent for intensive care monitoring, *Artificial Intelligence in Medicine*, 4, (1992),165-185.

[25]D.H. Hickam, E. H. Shortliffe, et al., A study of the treatment advice of a computer-based cancer chemotherapy protocol advisor, *Ann. Intern. Med.*, 101, (1985), 928.

[26] K.Horn, P. Compton, et al., An expert system for the interpretation of thyroid assays in a clinical laboratory, *Australian Computer Journal*, 17, (1985), 7-11.

[27]L. Ironi, M. Stefannelli, G.Lanzola, Qualitative models in medical diagnosis, *Artificial Intelligence in Medicine*, 2, (1990), 85-101.

[28]T. Kohonen, An Introduction to Neural computing, *Neural Networks*, 1,(1988), 3-16.

[29]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, (1990), 129-133.

[30]B.J. Kuipers, J.P. Kassirer, Causal Reasoning in Medicine: Analysis of a Protocol, Cognitive Science, 8, (1984), 363-385.

[31] R. Loeb, J. Brunner, D. Westenskow et al., The Utah Anaesthesia Workstation, *Anesthesiology*, 70, (1989), 999- 1007.

[32] P.L. Miller, Critiquing anesthetic management: the ATTENDING computer system, *Anesthesiology*, 58, (1983), 362-369.

[33]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.,(1989), 1-82.

[34]M. Nikoonahad, D.C. Liu, Medical ultrasound imaging using neural networks, *Electronic Letters*, 26, (1990), 545-546.

[35]J. S. Packer, Patient care using closed-loop computer control, *Computing and Control Engineering Journal*, 1,1, (1990), 23-28.

[36]R.S.Patil, Artificial Intelligence Techniques for Diagnostic Reasoning in Medicine, in H. Shrobe (ed), *Exploring Artificial Intelligence: Survey Talks from the National Conferences on Artificial Intelligence*, 347-380, Morgan Kaufmann, San Mateo, (1988).

[37]E. Pietka, Neural nets for ECG classification, in *Images of the 21st century: IEEE Engineering in Medicine and Biology 11th annual Conference*, (1989),2021-2022.

[38] J. R. Quinlan, Induction of Decision Trees, *Machine Learning*, 1, (1986), 81-106.

[39]I. J. Rampil, Intelligent Detection of Artifact, in *The Automated Anaesthesia Record and Alarm Systems*, J. Gravenstein, R. Newbower, A. Ream, N. Ty Smith (eds.), 175-190, Butterworths, Boston, (1987).

[40]G. Rutledge, G. Thomsen, B. Farr et. al., The design and implementation of a ventilator-management advisor, *Artificial Intelligence in Medicine*, 5, (1993),67-82.

[41]A. Sebald, Use of neural networks for detection of artifacts in arterial pressure waveforms, in *Images of the 21st Century - IEEE Engineering in Medicine and Biology 11th Annual Conference*, 2034-35,(1989).

[42]J. Shavlik, R. Mooney, G. Towell, Symbolic and Neural Learning Algorithms: An Experimental Comparison, *Machine Learning*, 6, 111-143, (1991).

[43]Shortliffe E. H.,Computer Programs to Support Clinical Decision Making, *JAMA*, 258, (1987), 61-66.

[44]D. Sittig, M. Factor, Physiological trend detection and artifact rejection: A parallel implementation of a multi-state kalman filtering algorithm, *Proceedings of 13th Symposium on Computer Applications in Medical Care*, 569-574, (1989).

[45]M. Stinchombe, H. White, Multilayer Feedforward Networks are Universal Approximators, *Neural Networks*, 2, 359-366, (1989).

[46] M.K. Sykes, Essential Monitoring, *Br. J. Anaesth.*, (1987), 59, 901-912.

[47]S. Uckun, Model-Based Reasoning in Biomedicine, *Critical Reviews in Biomedical Engineering*, 19,4, (1992), 261-292.

[48]S. Uckun, B. Dawant, D. Lindstrom, Model-based diagnosis in intensive care monitoring: the YAQ approach, *Artificial Intelligence in Medicine*, (1993),31-48.

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

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

[51]L. E. Widman, A model-based approach to the diagnosis of cardiac arrhythmias, *Artificial Intelligence in Medicine*, 4, (1992), 1-19.