Self-organising discovery, recognition and prediction of haemodynamic patterns in the intensive care unit

R. G. Spencer[1] C. S. Lessard[1] F. Davila[2] B. Etter[1]

[1]Department of Bioengineering, Texas A&M University, Texas, USA
[2]Texas A&M Medical School and Scott & White Memorial Hospital, Temple, Texas, USA

Abstract—*To care properly for critically ill patients in the intensive care unit (ICU), clinicians must be aware of haemodynamic patterns. In a typical ICU, a variety of physiological measurements are made continuously and intermittently in an attempt to provide clinicians with the most accurate and precise data needed for recognising such patterns. However, the data are disjointed, yielding little information beyond that provided by instantaneous high/low limit checking. Although instantaneous limit checking is useful for determining immediate dangers, it does not provide much information about temporal patterns. As a result, the clinician is left to sift manually through an excess of data in the interest of generating information. In the study, an arrangement of self-organising artificial neural networks is used to automate the discovery, recognition and prediction of haemodynamic patterns in real time. It is shown that the network is capable of recognising the same haemodynamic patterns recognised by an expert system, DYNASCENE, without being explicitly programmed to do so. Consequently, the system is also capable of discovering new haemodynamic patterns. Results from real clinical data are presented.*

Keywords—*Adaptive resonance theory, Adaptive spectral timing, Artificial neural networks, Haemodynamic pattern recognition, Hebbian learning, Intensive care unit, Leaky integrators, Medical information processing, Procedural knowledge, Short-term memory*

Med. & Biol. Eng. & Comput., 1997, **35**, 117–123

List of symbols

A = decay rate of leaky integrator nodes
$C_{ijk}(n)$ = variable resting potential of long-term memory traces of (LTM) adaptive spectral timing (AST) connections
K = number of unidirectional delays between AST nodes (predictive reach of the model)
$I_p(n)$ = occurrence (digital) variable of pth physiological event at time $t = nT$
M = maximum number of haemodynamic patterns
N = total number of synapses in the network
P = number of physiological events
$s_p(n)$ = recency of pth physiological event at time $t = nT$
T = sample period
$y_j(n,n)'$ = probability of occurrance of jth hamodynamic pattern at time $t = (n + n')T$
v_{pi} = strength of pth physiological event component of ith haemodynamic pattern vector
$w_{ijk}(n)$ = strength of kth delay synapse from ith to jth AST node at time $t = nT$
$x_i(n)$ = sensory activation level of ith haemodynamic pattern node (ART2) at time $t = nT$
$X_i(n)$ = laterally inhibited activation of ith haemodynamic pattern node (ART2) at time $t = nT$
ε = excitatory learning constant of LTM traces between AST nodes
γ = inhibitory learning constant of LTM traces between AST nodes
δ_{ll} = binary Hebb condition variable

First received 10 October 1994 and in final form 19 August 1996

1 Introduction

TO CARE properly for patients in the intensive care unit (ICU), clinicians must be aware of haemodynamic patterns. Although many of the processes involved in recognising these patterns are perfunctory, they have not yet been automated on a large scale. Instead, the information must be manually extracted by clinicians. This burdensome and time-consuming work makes the clinician's job more difficult, distracts from critical tasks, and reduces the chance of it being done correctly. Consequently, there is a great need for automatic recognition of haemodynamic patterns.

Several experimental expert systems have been created for the purpose of integrating the data and automating the recognition of several diagnoses (FAGAN, 1980; DAVIS *et al.*, 1984), but they have not become a permanent part of the ICU. Many of these systems exist within the symbolic framework of traditional rule-based paradigms and emulate the reasoning

processes of clinical experts. However, problems exist with such traditional expert systems owing to the heavy dependence on rules. As traditional rule-based systems usually do not have mechanisms for extracting rules from data, they require much declarative information *a priori*. As a result, rules must be manually extracted and explicitly programmed into the system. There are several disadvantages with such a system. First, the extraction of rules from experts is not always consistent. Secondly, the task of programming a rule-based system by itself is not very feasible, considering the patchy nature of rules themselves. Finally, losses in translation can occur between the expert and the computer, even if the experts are able to articulate consistent rules in the first place.

Some rule-based systems, such as CN2 (CLARK and NIBLETT, 1989), can learn the rules that drive them, thus overcoming the tedious task of declaring them; however, they often produce so many rules that they cannot be easily interpreted. As the data are complex and rules are such simple representations of knowledge, one rule tends not to mean much by itself; therefore, the mass action of many rules is required to represent the complexities of the data. Consequently, the knowledge is distributed, and such a system has lost much of the fundamental characteristic of rule-based systems, i.e. the ability to express its function in succinct common terms, yet it is still restricted to rule-based representations of knowledge.

An alternative type of expert system, the connectionist expert system, is the subject of much interest (GALLANT, 1988; EBERHART and DOBBINS, 1991; FU, 1991; POLL *et al.*, 1991; RIALLE *et al.*, 1991; PAPADOURAKIS *et al.*, 1992). Connectionist expert systems learn procedural-like knowledge *a posteriori*. By nature, these systems are highly adaptive and massively parallel. They are usually very 'trainable' and have a tendency to discover their own internal representations of knowledge, much like learning rule-based systems; however, they are not restricted to a rule-based framework and are free to discover more complex spatio-temporal patterns. These features make connectionist expert systems very attractive in data-rich environments (MILLER *et al.*, 1992).

The system proposed in this work is a haemodynamic pattern recognition system capable of discovering, recognising and predicting the same haemodynamic patterns as DYNASCENE (COHN *et al.*, 1990) without being explicitly programmed to do so. No assumptions were made beforehand about what the system should and should not see. The system is able to learn from its environment, discovering its own haemodynamic patterns. Such a system could advance medical knowledge by discovering unknown haemodynamic patterns, in addition to recognising well known ones.

2 Background

Recognition of haemodynamic patterns requires the recognition of sequences of physiological events. Events such as vasoconstriction, vasodilation, hypotension, hypertension, intravascular volume overload/depletion and increased pericardial pressure can occur in different orders, or permutations, indicating quite different haemodynamic trends. The physiological events themselves are recognised by observing raw (low level) physiological variables such as heart rate (HR), cardiac output (CO) and blood pressure (BP). Many of these variables are continuously measured in the ICU by automatic devices to ensure that they stay within normal limits. Although instantaneous limit checking is useful for determining immediate dangers, it does not provide much information about temporal patterns. These data could easily be processed by more advanced algorithms.

2.1 *Temporal pattern recognition*

When describing or tracking a dynamic system, a simple snapshot is not always adequate. As dynamic systems evolve along complex trajectories that have location and direction, a single snapshot does not usually contain enough information to define the state of the system. As snapshots convey location information only, two very different states of the system can have the same snapshot. To distinguish between two such states, directional information must be extracted by observing two or more successive snapshots in the order of occurrence.

The task of recognising haemodynamic patterns is a temporal pattern recognition problem. For example, to discriminate between congestive heart failure due to fluid overload and congestive heart failure due to pump failure, different permutations of the same combination of three physiological events must be distinguished. In each case, the same three events occur, but in different orders. Consequently, attention must be paid to temporal order.

2.2 *DYNASCENE*

Cohn *et al.* reported on a practical expert system called DYNASCENE, which exhibited a connectionistic-like macrostructure (COHN *et al.*, 1990). It was shown that DYNASCENE could recognise temporal patterns of physiological events and associate them with the corresponding haemodynamic disorders:

(*a*) congestive heart failure with fluid overload: intravascular volume overload → vasoconstriction → hypotension
(*b*) congestive heart failure due to pump failure: myocardial ischaemia → vasoconstriction → hypotension → intravascular volume overload
(*c*) hypovolaemia: intravascular volume depletion → vasoconstriction → hypotension
(*d*) sepsis: vasodilation → hypotension → vasoconstriction
(*e*) cardiac tamponade: increased pericardial pressure → vasoconstriction → hypotension

Although DYNASCENE represents a leap in the direction of connectionism, it still requires explicit programming by experts. As a result, the system is not able to recognise any haemodynamic patterns other than those programmed into it.

3 Methods

A self-organising haemodynamic pattern recognition and prediction system was created and simulated. The complete model is shown in Fig. 1. It consists of three layers: (*i*) a short-term memory (STM) layer, (*ii*) a temporal pattern recognition (categorisation) layer; (*iii*) a temporal pattern prediction layer. Each layer is described below.

3.1 *Short-term memory layer*

A bank of seven leaky integrators was used as a front-end STM. This type of STM was used for two reasons: its ability to perform current time processing, and its bias towards recency of events. Each leaky integrator corresponded to one of the physiological conditions denoted by I_p, as shown in Fig. 1. Using the same STM as Gjerdingen, the activation of a leaky

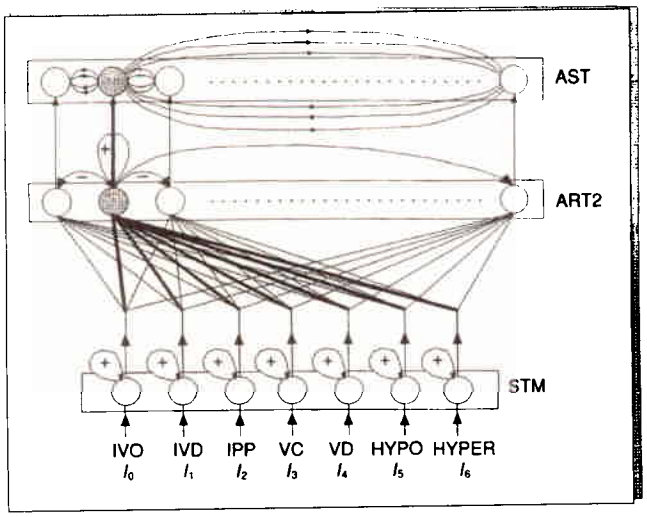

Fig. 1 *Self-organising temporal pattern recognition and prediction network; three distinct layers are present: a STM layer composed of leaky integrator neurons; a temporal pattern recognition layer composed of laterally inhibited winner-take-all neurons; a temporal pattern prediction layer composed of AST neurons (CARPENTER and GROSSBERG, 1991); IVO = intravascular volume overload; IVD = intravascular volume depletion; IPP = increased pericardial pressure; VC = vasoconstriction; VD = vasodilation; HYPO = hypotension; HYPER = hypertension*

integrator at time $t = (n+1)T$ (where T is the sample period) is (GJERDINGEN, 1992)

$$s_p(n+1) - (1-A)s_p(n) + [1-(1-A)s_p(n)]I_p(n) \qquad (1)$$

where $0 < A < 1$. When physiological event p occurred, $s_p(t)$ was driven to saturation in one iteration. The membrane potential $s_p(t)$ remained saturated until physiological event p ceased to exist, at which time it began to decay at a rate determined by the forgetting term A. Using the activations of all STM nodes, a normalised STM vector was created. A typical state of the STM is shown in Fig. 2.

Currently, there is no formal method for setting the forgetting term A. However, it can be chosen according to some common sense guidelines. A should be such that the STM usually reflects the last few events. If A is too large, traces of the events decay too quickly and there is not enough emphasis on primacy. On the other hand, if A is too small, too much of the past is stored and there is not enough emphasis on recency. If A is either too large or too small, the STM vector tends to look the same, regardless of what is happening. A good starting range for A is between 0·1 and 0·5.

3.2 *Haemodynamic pattern recognition layer*

The temporal patterns in the STM were categorised according to the degree of similarity with previously recognised temporal patterns using a continuous-valued adaptive resonance theory (ART2) network (CARPENTER and GROSSBERG, 1987). The state of the STM at any instant was represented by one of M haemodynamic pattern categories (nodes) in the ART2 layer. The equations used were effectively the same as the ART2 equations; i.e. the STM vector $\boldsymbol{s}$ was categorised by calculating the degree to which it excited each node in the ART2 network and choosing the one with the maximum activation. The ith activation was calculated by the dot product between the ith haemodynamic vector $\boldsymbol{v}_i$ and the STM vector $\boldsymbol{s}$

$$x_i(n) = \sum_{\forall p} s_p(n)v_{pi}(n) = s(\boldsymbol{n}) \cdot v_i(\boldsymbol{n}) \qquad (2)$$

The winning node then laterally inhibited all other nodes such that its own state was binary one and those of the others were binary zero, as shown in eqn. 3, and it then adapted its own haemodynamic pattern vector towards $\boldsymbol{s}$ using a winner-take-all learning rule (HECHT-NIELSEN, 1987). As a result, the ART2 nodes self-organised to recognise the STM patterns conveyed by $\boldsymbol{s}$ over time.

$$X_i(n) = \left\{ \begin{array}{ll} 1 & \text{if } x_i(n) = \max(x_i(n)), \forall i \\ 0 & \text{otherwise} \end{array} \right\} \qquad (3)$$

As a general rule, the capacity to remember haemodynamic patterns, M, should be as high as possible to allow the system to declare new categories to characterise new patterns adequately; however, the vigilance parameter of the ART2 network should be low enough not to start too many; otherwise, performance can be compromised.

3.3 *Temporal pattern prediction layer*

For prediction purposes, a layer of neurons with adaptive spectral timing (AST) connections (GROSSBERG and SCHMAJUK, 1989) was created on a one-to-one basis with each ART2 category node. In addition to being connected to a single ART2 node, each AST node was connected to every other node in the AST layer by $2K$ synapses (K synapses from i to j, and K synapses from j to i), each denoted by w_{ijk}, for $i = 1,2,\ldots I$, $j = 1,2,\ldots, I$, and $k = 1,2,\ldots, K$, where I was the number of ART2 categories and K was the maximum number of discrete time delays between two AST nodes. A set of unidirectional pathways is shown in Fig. 3.

Discrete versions of the AST equations were used. They functioned as follows: when the ith haemodynamic pattern occurred at time $t = nT$, synapse w_{ijk} was activated at time $t = (n+k)$T. The synaptic gain w_{ijk} was modified by an adaptrode-based learning rule (MOBUS, 1990):

$$w_{ijk}(n+1) = w_{ijk}(n) + \varepsilon\delta_{Hebb-ijk}(n)[1 - w_{ijk}(n)] - \gamma[w_{ijk}(n) - C(n)] \qquad (4)$$

where $\delta_{Hebb-ijk}$ represents the Hebb condition (HEBB, 1949) for synapse w_{ijk}, the variables ε and γ are excitatory and decay rates, respectively, and $C(n)$ is a variable resting potential

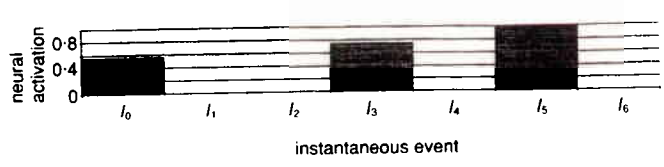

Fig. 2 *Typical snapshot of STM; as shown by the activation levels, I_1 occurred first because it has had the most time to decay, followed by I_3 and I_5*

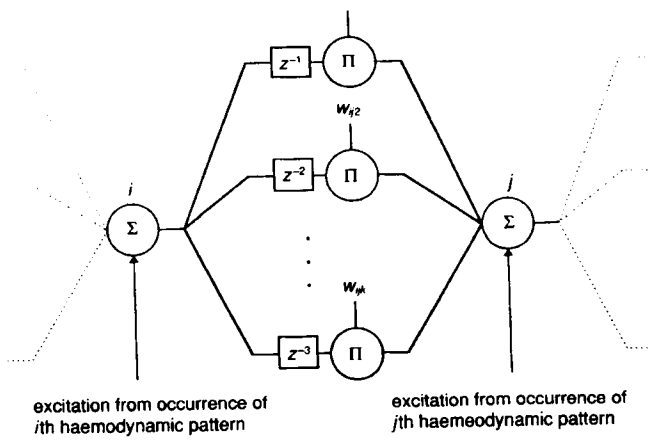

Fig. 3 *AST pathways from node i to node j*

$(0 \leqslant C(n) \leqslant 1)$ described by Mobus (MOBUS, 1990). The aptrode-based synapse was used to protect the synaptic gain against transient associations. The Hebb condition was satisfied when any one of the delayed recognitions of the ith haemodynamic pattern occurred at the same time as the jth haemodynamic pattern, as given by eqn. 5. When this condition was satisfied, the synaptic gain between node i

$$\delta_{Hebb-ijk}(n) = X_i(n-k)X_i(n) \tag{5}$$

and j was increased for the kth pathway, in a classical conditioning sense (PAVLOV, 1927).

To predict future haemodynamic patterns, the incoming signals to each neuron in the AST network were summed n' time units in advance:

$$y_i(n, n') = \sum_{i=1}^{I} \sum_{k=n'}^{n+n'} X_i(n + n' - k) w_{ijk}(n) \tag{6}$$

The resulting quantity $y_j(n,n')$ represents an aggregate conditional probability that the jth haemodynamic pattern will occur at time $t = (n + n')T$, given that certain haemodynamic patterns were recognised k time units in the past.

The parameter K defines the predictive reach of the network, i.e. the length of the longest temporal pattern that can be recognised and predicted. This parameter should be set according to the period between samples and the longest temporal pattern. The product KT should be no smaller than the time that it takes for the longest observed haemodynamic pattern to unfold.

4 Results

The network was implemented in UNIX C on a Sparc 20 workstation and simulated using two sets of data: a fabricated set for calibration and actual clinical data. For the purpose of recognising instantaneous physiological conditions, four physiological variables were measured (clinical data set) or fabricated (calibration set): heart rate (HR); systolic arterial pressure (SAP); diastolic arterial pressure (DAP); and diastolic pulmonary artery pressure (PAD). The instantaneous values and rates of change of these variables were used to recognise seven physiological events: intravascular volume overload, intravascular volume depletion; increased pericardial pressure; vasoconstriction; vasodilation; hypotension; and hypertension. Of these events, six were the same as those recognised by DYNASCENE.

4.1 *Calibration*

For calibration purposes, the model was trained on the same five haemodynamic sequences recognised by DYNASCENE. Raw data were provided to produce the required sequence of events. The main parameters of the simulation were $A = 0{\cdot}25$, $K = 32$, and $M = 78$. By setting $A = 0{\cdot}25$, haemodynamic events could not decay too quickly from the STM, so that sequences of adequate length could be recognised while not placing too much emphasis on primacy information. By setting K to 32, the system was able to recognise and predict events that occur within 32 time units of each other (also referred to as predictive reach). Finally, by setting M equal to 78, the system had a capacity to recognise, remember and refine 78 different haemodynamic pattern exemplars.

Some of the results of the calibration are illustrated in Figs. 4 and 5, which show how the system reacted to two haemodynamic sequences: congestive heart failure with fluid overload, and congestive heart failure due to pump failure. Of the five sequences used for calibration, these two were chosen because they represent two different permutations of the same combination of physiological events.

Figs. 4*a–c* show the progressive states of the STM during the unfolding of congestive heart failure with fluid overload. Fig. 4*a* shows the occurrence of intravascular volume overload after a long period of seeing no physiological events. Fig. 4*b* shows the occurrence of vasoconstriction, and Fig. 4*c* shows the occurrence of hypotension. Note that the memory traces of intravascular volume overload and vasoconstriction have decayed by amounts that correlate to the elapsed time since their occurrence. As a result, the state of the STM conveyed the order in which the events occurred. Fig. 4*d* shows the prediction of haemodynamic pattern C after seeing only the first two events. Haemodynamic pattern C was one of the temporal patterns that represented congestive heart failure with fluid overload. Fig. 4*e* shows the eventual recognition of haemodynamic pattern C after seeing the third and final event.

Fig. 5, which shows how the system reacted to congestive heart failure due to pump failure, is structured in the same manner as Fig. 4. Note that in Fig. 5*d* and *e* the predicted and eventually recognised haemodynamic pattern was not pattern C, although the same combination of physiological events occurred.

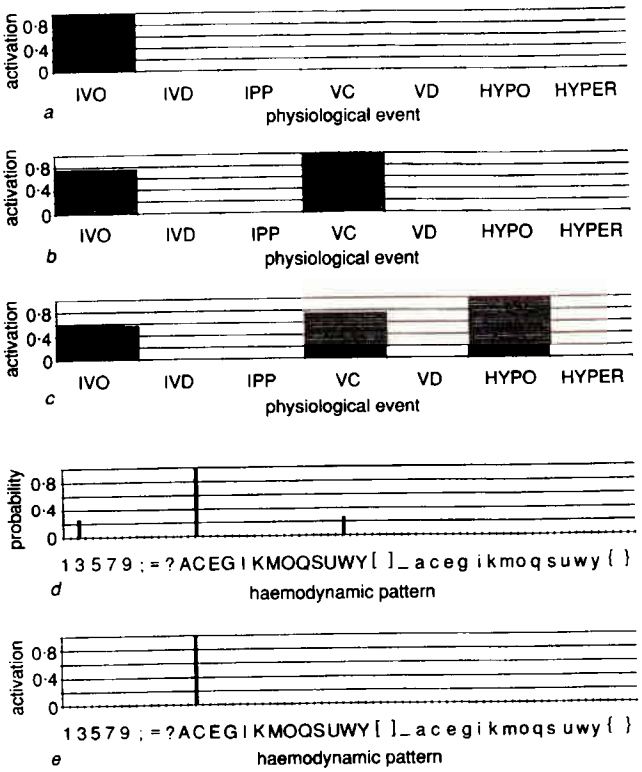

Fig. 4 *Prediction and recognition of congestive heart failure with fluid overload: (a–c) state of STM is shown as the condition unfolds: (a) onset of intravascular volume overload; (b) onset of vasoconstriction; (c) onset of hypotension; (d) prediction of congestive heart failure with fluid overload before hypotension after the system has seen only the first two events; (e) recognition of congestive heart failure with fluid overload after hypotension after the system has seen the third and final event: (d,e) x axis represents all haemodynamic patterns; IVO = intravascular volume overload; IVD = intravascular volume depletion; IPP = increased pericardial pressure; VC = vasoconstriction; VD = vasodilation; HYPO = hypotension; HYPER = hypertension*

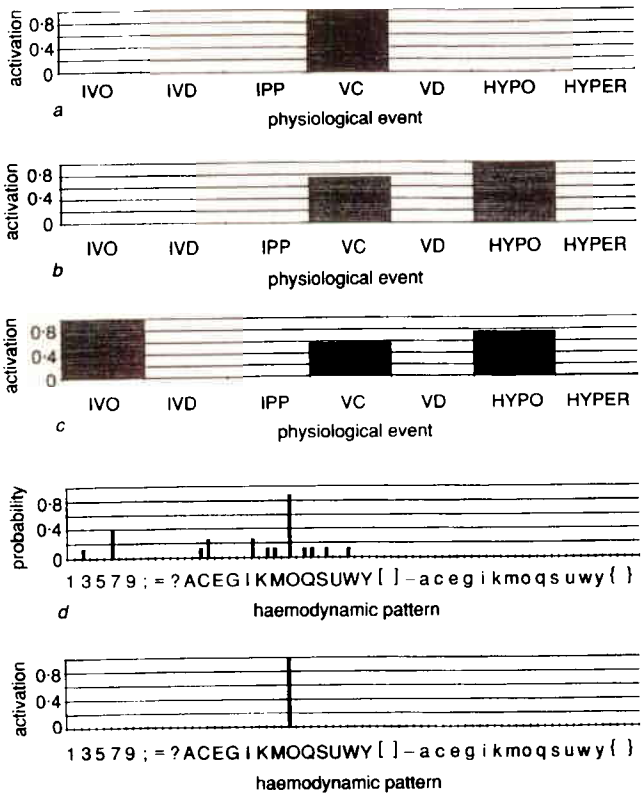

Fig. 5 *Prediction and recognition of congestive heart failure due to pump failure: (a–c) state of STM is shown as the condition unfolds: (a) onset of vasoconstriction; (b) onset of hypotension; (c) onset of intravascular volume overload; (d) prediction of congestive heart failure due to pump failure before intravascular volume overload after the system has seen only the first two events; (e) recognition of congestive heart failure due to pump failure after the system has seen the third and final event: (d,e) x axis represents all haemodynamic patterns; IVO = intravascular volume overload; IVD = intravascular volume depletion; IPP = increased pericardial pressure; VC = vasoconstriction; VD = vasodilation; HYPO = hypotension; HYPER = hypertension*

4.2 *Application*

The system was also tested on actual clinical data measured at 5 min intervals using the same parameters: $A = 0{\cdot}25$, $K = 32$, and $M = 78$. A full 24 h of data (288 5 min samples) from a cardiac patient were analysed. As expected with 5 min interval samples, a large amount of variability was observed from one sample to the next. Three events occurred repeatedly: intravascular volume overload, vasodilation and hypertension. From this sequence, several interesting haemodynamic patterns were discovered:

1 vasodilation and hypotension → hypertension → intravascular volume overload
2 hypotension → hypertension → intravascular volume overload and vasodilation
3 hypotension → hypertension → vasodilation → intravascular volume overload
4 hypotension → vasodilation → intravascular volume overload and hypertension
5 hypotension → vasodilation → hypertension → intravascular volume overload
6 hypotension → hypertension → intravascular volume overload → vasodilation

Although the physiological events occurred in the orders indicated, many did not occur in a contiguous fashion. Instead, they were mixed in with other events. As a result, more than one ART2 node could represent the same haemodynamic pattern. Variations of the same permutation were represented by sets of ART2 nodes. Although not shown in the list, the time-between-event information was retained by each haemodynamic pattern category, such that two haemodynamic patterns with identical permutations but different interstimulus intervals (ISIs) would still be distinguished. For example, haemodynamic pattern 6 shows four distinct physiological events occurring in serial order. As it turned out, the occurrence of hypotension and hypertension were separated from intravascular volume overload and vasodilation by several time units. This can be seen by the actual synaptic vector of haemodynamic pattern 6, as shown in Fig. 6. Note the difference in height between intravascular volume overload and hypertension. The difference between activations indicates that these two sets of events were separated by more than a single time unit. If the contiguous version of the same sequence had occurred, the STM vector would have been different enough from this vector to justify a new category, thus distinguishing between the two variations of the same sequence.

5 Limitations

In addition to the need for much more clinical testing and calibration, there are three limitations to this design. First, the number of connections needed to connect the nodes that represent haemodynamic patterns grows faster than linearly with the number of remembered patterns:

$$N(M, P, k) = MP + 2k \sum_{x=1}^{M-1} x \tag{7}$$

For example, to increase the number of nodes from $M = 78$ to $M = 79$, with seven events ($P = 7$) and a predictive reach of $k = 32$, would increase the number of synapses from 195 170 to 200 169 (an increase of 4999 synapses). Alternatively, to increase k to 33 from the same initial capacity would increase the number of synapses to 201 252 (an increase of 6082 synapses).

The second limitation is the learning time for a given patient. To achieve a truly customised network for each patient, the network needs to observe the temporal patterns of that patient for some time. A reasonable amount of time would be that needed to fully observe each temporal pattern five or six times. Of course, a network that has been previously trained on other patient(s) could be used as a starting point, and then the network would slowly adapt itself to the new patient. Having an experienced system monitor the patient during the initial adaptation period of a new system would certainly be more useful, however, such a system would have to have a much higher capacity than an unexperienced system.

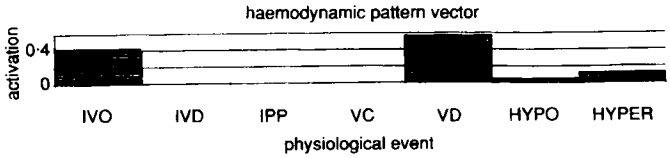

Fig. 6 *Synaptic weight vector of haemodynamic pattern 6; this pattern category shows the sequence hypotension → hypertension → intravascular volume overload → vasodilation with a lag of more than one time between the first and last two events; IVO = intravascular volume overload; IVD = intravascular volume depletion; IPP = increased pericardial pressure; VC = vasoconstriction; VD = vasodilation; HYPO = hypotension; HYPER = hypertension*

The third limitation of this approach is that it has no mechanism for categorising two temporal patterns played out at different speeds (time-scales) as the same pattern. Currently, these two patterns are stored as two different patterns, which causes two potential problems: the capacity of the network must be higher to store multiple scales of the same pattern; and the network cannot recognise a time-scaled version of a previously stored pattern without seeing it first. There are several ways to accomplish time-scale invariance that were not employed in this study: employing multiple delay sets that are scaled versions of each other, with a lateral inhibition to choose the best representation, employing adaptive delays; and using expectation and fulfilment models to set the tempo of recognition and recall.

6 Discussion

In addition to the ability to recognise haemodynamic patterns, the model is able to predict several haemodynamic patterns before they completely unfold in time. A haemodynamic pattern that is defined by a permutation of several physiological events, for example, can be predicted after the first couple of events. Other haemodynamic patterns that occur regularly, or even semi-regularly, can be predicted if they fall within the predictive reach of the network.

A very interesting property of the interconnected adaptive spectral timing network has emerged: the ability to predict temporal patterns based on rhythm. At first, this property was alarming, but on further reflection it seemed natural for such a network, which is so deeply rooted in discovering interstimulus intervals, to become caught up in rhythm. For example, if haemodynamic pattern B was frequently made to occur k time units after haemodynamic pattern A, the AST network learned to predict pattern B based on seeing pattern A. If haemodynamic pattern B was suddenly removed from the input stream altogether, pattern B was still predicted by the recognition of pattern A. Even though the first couple of events of pattern B did not occur, the network still predicted the last physiological events of pattern B. This result is interesting because the system predicted pattern B without seeing any of the events in pattern B. If pattern B represented congestive heart failure, for example, then the system would have predicted congestive heart failure without seeing the events associated with congestive heart failure, even if the state of the STM had completely decayed since the onset of pattern A. Essentially, the AST network detached itself from external cues, confident in its own internal momentum. Of course, the system never recognised pattern B because it never actually occurred; it just predicted it. This feature could prove very useful in discovering haemodynamic patterns that are longer than those simulated in this study.

Another interesting property of the predictive network was its ability to predict haemodynamic patterns that occurred at varying frequencies. For varying ISIs between two patterns, the first derivative of the predictive potentials; $y_j(n, n')$, with respect to time, sometimes correlated more closely to the probability of occurrence than absolute magnitude. For example, if pattern B usually followed sequence A 20–30 time units later, then the predictive potentials of the nodes that recognised variations of pattern B tended to grow more (high first derivative) than other potentials during the 20–30 time-unit window. It is expected that this feature would be extremely useful for discovering real-world haemodynamic patterns in the ICU.

7 Conclusions

The collection of artificial neural networks presented in this study is able to recognise haemodynamic patterns without being explicitly programmed to do so. As the system learns self-organised representations of input data, it requires no declarative knowledge *a priori*, other than the preprocessing required to recognise fundamental physiological events. As a result, it is able to discover unknown haemodynamic patterns, in addition to those that are well known. It is expected that such a system, combined with robust methods for preprocessing and cross-verifying raw data, could be used in the clinical environment in real time, to assist clinicians in caring for critically ill patients.

References

CARPENTER, G.A., and GROSSBERG, S. (1991): 'Pattern recognition by self-organizing neural Networks,' (MIT Press, Cambridge, Massachusetts)

CARPENTER, G.A., and GROSSBERG, (1987): 'ART 2: Self-organization of stable category recognition codes for analog input patterns,' *Appl. Opt.*, **26**, pp. 4919–4930

CLARK, P., and NIBLETT, T. (1989): 'The CN2 induction algorithm,' *Machine learning*, **3**,(4), pp. 261–283

COHN, A.I., ROSENBAUM, S., FACTOR, M., and MILLER, P.L. (1990): 'DYNASCENE: An approach to computer-based intelligent cardiovascular monitoring using sequential clinical 'scenes',# *Meth. Inf. Med.*, **29**, pp. 122–131

DAVIS, R., BUCHANAN, B., and SHORTLIFFE, E. (1984): 'Production rules as a representation for knowledge-based consultation program,' *in*: 'Readings in medical artificial intelligence: the first decade,' (CLANCEY, W., and SHORTLIFFE, E. (Eds.) (Addison-Wesley, Reading, Massachusetts)

EBERHART, R.C., and DOBBINS, R.W. (1991): 'Using neural networks in hybrid medical diagnostic systems', Proc. 13th Ann. Int. Conf. of IEEE Engineering in Medicine & Biology Society, Section 4, pp. 1470–1471

FAGAN, L (1980): 'VM: Representing time dependent relations in a medical setting.' PhD. Thesis, Stanford, California

FU, L. (1991): 'A hybrid medical expert system.' Proc. 13th Ann. Int. Conf. of IEEE Engineering in Medicine & Biology Society, Section 4, pp. 1290–1291

GALLANT, S.I. (1988): 'Connectionist expert systems,' *Commun. ACM*, **2**, (31), pp. 152–169

GJERDINGEN, R.O. (1992): 'Learning syntactically significant temporal patterns of chords: a masking field embedded in an ART3 architecture,' *Neural Netw.*, **5**, pp. 551–564

HEBB, D.O. (1949): 'The organization of behavior,' (Wiley, New York)

HECHT-NIELSEN, R. (1987): 'Counterpropagation networks,' *Appl. Opt.* **26**, (23), pp. 4979–4984

MILLER, A.S., BLOTT, B.H., and HAMES, T.K. (1992): 'Review of neural network applications in medical imaging and signal processing.' *Med. Biol. Eng. Comput.*, **30**, pp. 449–464

MOBUS, G.E. (1990): 'The adaptrode learning model: applications in neural network computing.' Technical Report CRPDC-90-5, Center for Parallel and Distributed Computing, University of North Texas, Denton

PAPADOURAKIS, G.M., GAGA, E., VARELTZIS, G., and BEBIS, G. (1992): 'Use of artificial neural networks for clinical decision-making (maldescensus testis).' Proc. Int. Joint Conference on Neural Networks (IJCNN'92), Baltimore, Maryland, USA, June 1992

PAVLOV, I.P. (1927): 'Conditioned reflexes,' (Oxford University Press)

POLI, R., CAGNONI, S., LIVI, R., COPPINI, G., and VALLI, G. (1991): 'A neural network expert system for diagnosing and treating hypertension,' *IEEE Comput.*, March, pp. 64–71

RIALLE, V., OHAYON, M., AMY, B., and BESSIERE, P. (1991): Medical knowledge modeling in a symbolic-connectionist perspective,' *Proc. 13th Ann. Int. Conf. of IEEE Engineering in Medicine & Biology Society*, Section 4, pp. 1109–1110

Author's biography

Ronald G. Spencer gained his BSc in Electrical Engineering from GMI Engineering and Management Institute in Flint, Michigan, in 1991. After graduation, he worked for General Motors, as a Programmer/Associate Engineer. After obtaining his MSc in Bio-engineering from Texas A&M University in 1994, he joined Dr. Edgar Sanchez-Sinencio at Texas A&M and is now pursuing a PhD in Electrical Engineering, with emphasis on VLSI design for neural networks and image processing. His current research interests include monolithic implementation of multiresolution analysis wavelet filter banks for face recognition.