Two-Stage Interpretation of ICU Data Based on Fuzzy Sets

Friedrich Steimann and Klaus-Peter Adlassnig
Institut für Medizinische Computerwissenschaften der Universität Wien
Währinger Gürtel 18–20
A-1090 Wien
e-mail: steimann, kpa@imc.akh-wien.ac.at

Introduction

Intelligent on-line monitoring of ICU patients requires mechanisms to detect and interpret significant patterns in the flood of recorded data. Many monitoring methods have been proposed, ranging from simple surveillance of thresholds to complex reasoning about and forecasting of the patient's behaviour based on computational models of the human organism.

From the many components that constitute a comprehensive intelligent monitor we focus on two the combination of which results in a workable device capable of mapping collected data onto physiological and pathophysiological states. They are arranged as stages and form the kernel of a bigger, still evolving system called DiaMon-1, which is dedicated to the on-line interpretation of patients' status.

Stage one, which has also been more generally described as pre-processing in [7], deals with the transformation of incoming data streams into a sequence of events. In this paper we focus on a special issue of temporal abstraction: the detection of fuzzy trends. The event of detection is reported to stage two, a fuzzified deterministic automaton capable of assigning diagnoses to sequences of fuzzy events.

For the purpose of evaluation, ICU data of an eight month old female suffering from adult respiratory distress syndrome (ARDS) was distributed to the symposium participants. The data set contains continuous heart rates, blood pressures, O_2 saturation values, and ventilator settings sampled over a 12-hour period. In addition, discontinuous data from laboratory tests and the flowsheet was provided. In applying DiaMon-1 to the presented case we chose to concentrate on the events centred around hand bagging sessions where the mode of ventilation was changed from continuous mandatory ventilation (CMV) to hand bagging, accompanied by a temporary increase in the fraction of inspired oxygen (FIO_2) from 0.5 to 1 and the delivery of Ventolin.

Stage One: Trend Detection

For reasons of abstraction and subsequent symbolic processing, quantitative data is often transformed into qualitative terms or symbols such as *high*, *low*, or *normal*. This classification process can be implemented in many ways, e.g., through interval assignment, probabilistic or fuzzy classifiers. Trend detection resembles this process in principle, the main difference being that terms are extended by a temporal dimension covering the course of a parameter over time.

In DiaMon-1, the concept of a trend is designed to model imprecise notions of courses such as "O_2 saturation slowly rising within the last ten minutes" naturally, and its trend detection is based on fuzzy classification [3] rather than statistical methods. A *fuzzy trend* $\tilde{c}$ is defined as a fuzzy relation in the time-value space where $\mu_{\tilde{c}}(t_i, v_i)$ denotes the *degree of membership* or *membership grade* of a pair (t_i, v_i) in $\tilde{c}$. The *degree of match* of $\tilde{c}$ with a series $<(t_i, v_i)>$ of time-stamped values relative to a time of onset t_o is then defined as $\min_i \mu_{\tilde{c}}(t_i - t_o, v_i)$.

Table 1 defines a number of trends considered to be relevant to the presented monitoring case. Due to the lack of continuous arterial O_2 partial pressures, oxygenation is judged by arterial O_2 saturation (SaO_2). Numbers present thresholds for full compliance (membership grade equals one), while numbers in parentheses indicate the bounds beyond which membership grades are defined to be zero.

A trend may be better characterised by its parameter's first derivative than by its absolute course. Reasoning about derivatives of non-continuous functions, however, is nontrivial and not addressed here. For a more comprehensive specification of fuzzy trends and corresponding trend detection see [8].

For the detection of longer-term trends (more than approximately 10 data items) it may be advisable to pass the collected data through a moving average filter to compensate for short-term deviations from the course. In that case it is also meaningful to provide some kind of early warning mechanism, as it may not be affordable to hold back notification until the whole trend has been seen. Trend detection of DiaMon-1 is therefore enhanced by a preview mechanism indicating developing trends, which is presented in [8].

Table 1: Definition of trends

TREND	DEFINITION	DEPICTION
adequate oxygenation	SaO_2 above 97% (93%) for 5 minutes	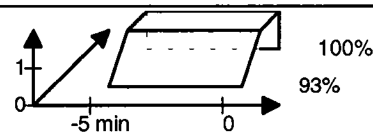
hypoxemia	SaO_2 between 90% and 93% (87% and 97%) for 2 minutes	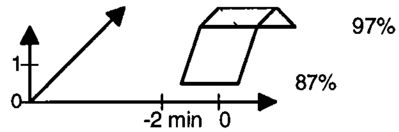
high FIO_2	FIO_2 above 60% for 30 seconds	trivial
low FIO_2	FIO_2 below 60% for 30 seconds	trivial
rapidly improving oxygenation	SaO_2 increasing from 87–95% (85–99%) to 97–100% (93–100%) within 30–90 seconds (5–120 seconds)	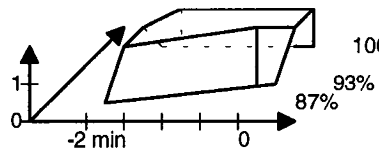
slowly decreasing oxygenation	SaO_2 above 96% (91%) steady or decreasing to 94% (89%) within 25 minutes	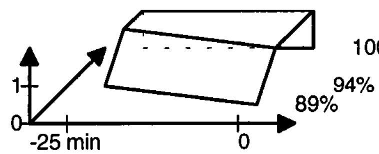

Figure 1 depicts the results of trend detection applied to the provided data sample. In each frame, the upper half displays SaO_2 in a range of 80–100%, while the lower half depicts the concurrent degree of match (ranging from 0–1) with the respective trend. SaO_2 of *adequate oxygenation*, *hypoxemia*, and *slowly decreasing oxygenation* was smoothened by a moving average filter of two, one, and five minutes, respectively. All output was generated by DIAMON-1; the trends on FIO_2 were regarded trivial and hence omitted.

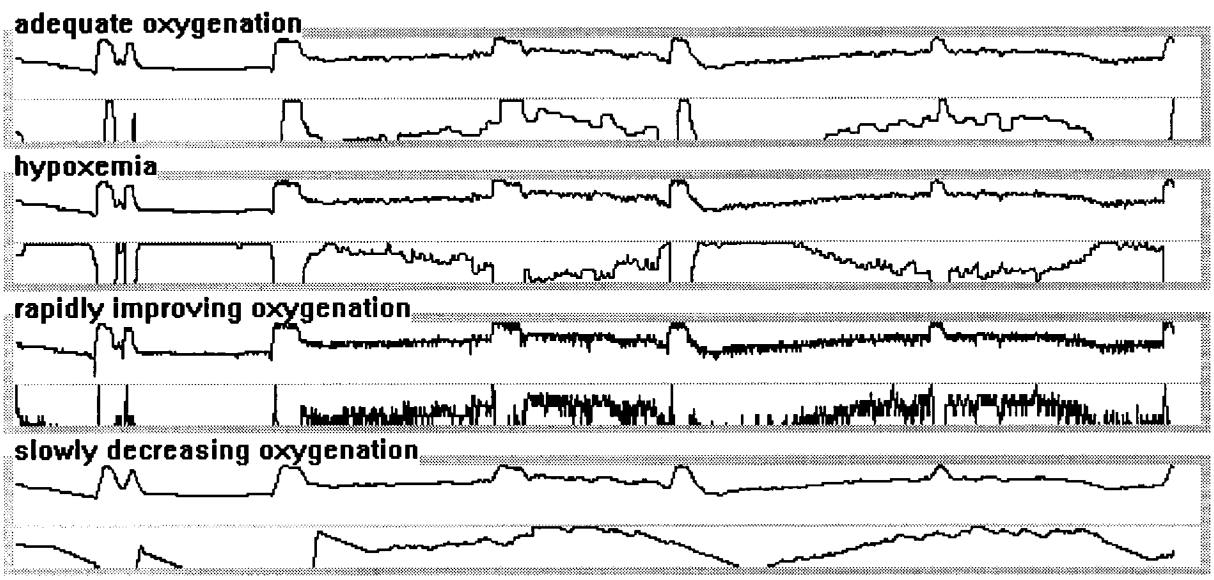

Figure 1: Trend detection applied to SaO_2 (in part filtered through moving average), 12-hour period

The example demonstrates the sensitivity of the proposed method: all sudden increases in SaO_2 were correctly identified. Also note the smoothening effect of the trend detection itself: the longer the trend to be detected is, the less volatile is the curve indicating the degree of match. However, trend detection merely increases the total amount of information presented—a clearer picture of the case is not yet provided.

Stage Two: Event Interpretation

While the detection of a trend itself may not be of very high diagnostic value, it can trigger the transition from one (physiological or pathophysiological) state to another. Such a state may then be associated with a diagnostic interpretation taking not only the current condition of the patient, but also her/his history into account. In some cases, these state-space trajectories are even indispensable to arrive at the correct diagnosis (cf. [1, 9]).

An eligible framework of states and transitions is provided by one of the classical formalisms of theoretic computer sciences: automata theory. Deterministic finite automata are capable of analysing sequences of events or symbols forming so-called regular expressions. If an automaton is designed as a model of the system under observation, then its current state reflects the actual status of that system. Applied to patient monitoring, an automaton can model possible courses and complications of a disease, allowing to follow the patient on an abstract level. Following such an automaton is defined in an exemplary fashion.

In the presented ICU case we focus on the states and transitions related to the change of mode of ventilation from CMV to hand bagging and vice versa. The expected effect of hand bagging along with delivery of pure oxygen is an immediate increase in oxygenation reflected in a rise of SaO_2 to values close to 100%. In addition, if manually exerted higher pressures and the effects of Ventolin helped in recruiting occluded alveolar spaces, a persistent improvement in oxygenation may be expected.

Tables 2 and 3 present the states and transitions used to obtain the results presented in Fig. 3. Figure 2 graphically depicts the provided transitions in a state transition diagram.

Table 2: Definition of states under consideration

State	Interpretation (Diagnosis)
start	initial state, system warm up, undecided
normal	oxygenation is satisfactory without additional effort such as increased FIO_2
hypoxic	oxygenation is too low and should be improved
responding to high FIO_2	high FIO_2 has affected oxygenation positively
not responding to high FIO_2	high FIO_2 does not have the desired effect
improved after hand bagging	hand bagging has persistently improved oxygenation
not improved after hand bagging	hand bagging shows no satisfactory effect

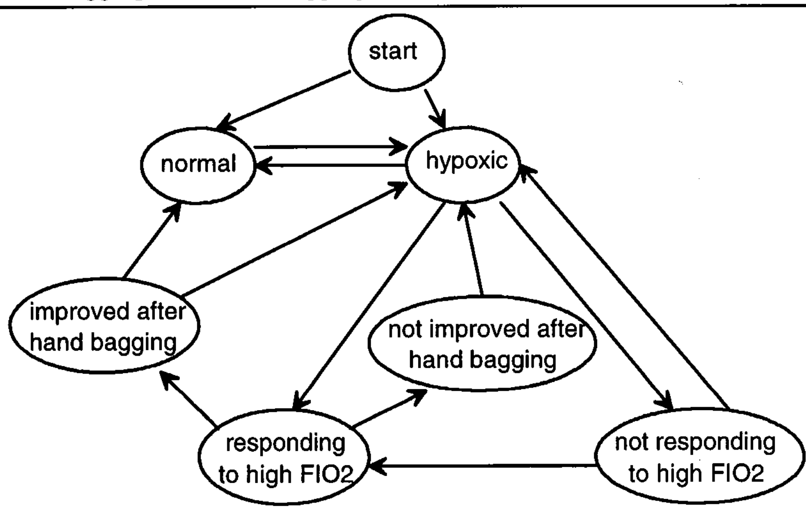

Figure 2: State transition diagram of the automaton

The automaton employed is a fuzzified version described in [7]. It is best characterised as the superposition of a fuzzy and a crisp automaton maintaining a current state (called *active*) as well as a set of gradually satisfied states. It transitions on fuzzy events, fuzzy sets associating a degree with every event of the automaton´s input alphabet. To obtain fuzzy events, the trends and the degrees to which they have been detected in stage one are comprised in one fuzzy set.

Table 3: Definition of transitions

FROM STATE	ON EVENTS	TO STATE
start	*adequate oxygenation*	*normal*
	hypoxemia	*hypoxic*
normal	*hypoxemia*	*hypoxic*
hypoxic	*low* $FIO_2 \wedge$ *adequate oxygenation*	*normal*
	high $FIO_2 \wedge$ *rapidly improving oxygenation*	*responding to high* FIO_2
	high $FIO_2 \wedge$ *hypoxemia*	*not responding to high* FIO_2
responding to high FIO_2	*low* $FIO_2 \wedge$ *slowly decreasing oxygenation*	*improved after hand bagging*
	low $FIO_2 \wedge$ *hypoxemia*	*not improved after hand bagging*
not responding to high FIO_2	*low* $FIO_2 \wedge$ *hypoxemia*	*hypoxic*
	high $FIO_2 \wedge$ *adequate oxygenation*	*responding to high* FIO_2
improved after hand bagging	*adequate oxygenation*	*normal*
	hypoxemia	*hypoxic*
not improved after hand bagging	*hypoxemia*	*hypoxic*

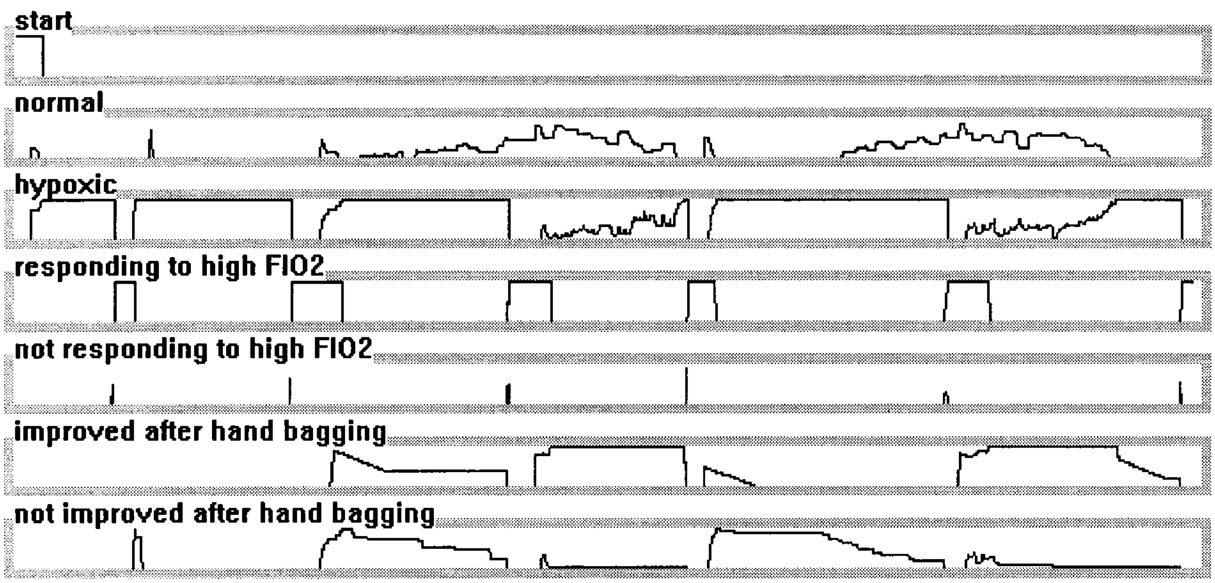

Figure 3: Distribution of states over time

Figure 3 comprises definite and vague statements in one diagram: note that there is always one (active) state with a membership grade of 1, which is a consequence of the automaton "waiting" for a full transition to a successor state being triggered by sufficient evidence; on the other hand, candidate successors may show rising membership grades reflecting a change in the patient´s condition. This models human decision making naturally: once a decision has been made, it is usually pursued rather uncritically until there is sufficient evidence for another to be favoured.

Discussion

Employing automata for sequence interpretation implements implicit memory avoiding expensive history look-ups as described in [9]. However, the expressiveness of automata is practically limited, particularly when complex interrelations are to be expressed. This drawback will be alleviated in the future by allowing to cascade automata, thereby introducing an automata hierarchy implementing arbitrary levels of abstraction and parallelism.

Unlike Shahar´s approach to temporal abstraction [4], DIAMON-1 can only abstract data to predefined trends. It is not equipped to actively create its own abstractions comprising data from several disjoint intervals. Its automaton can, however, assign interpretations to sequences of trends provided that respective states and transitions have been specified. Corrective consideration of delayed data, although possible in principle, presents severe problems to time-constrained on-line monitoring processes and is not accounted for.

Haimowitz' $TrenD_x$ [2] uses so-called trend templates to define and detect trends in series of time-stamped data. A trend template is a collection of temporal intervals each of which constrains a number of parameters. The temporal intervals can be of indeterminate length, the bounds are then related to other intervals or landmark points through temporal constraints.

Trends are detected by assigning time-stamped data to suitable intervals. For this purpose, $TrenD_x$ recursively generates and prunes hypotheses, instantiations of trend templates whose intervals are adapted to fit the data's times and values.

Although a direct comparison of $TrenD_x$ with DIAMON-1 is difficult, fuzzy trends of DIAMON-1 find their counterparts in the intervals of $TrenD_x$, while the automaton best corresponds to the trend template. The fundamental difference between the two arises from the initial approach taken: $TrenD_x$ is based on the verbal description of trends as uttered by experts, while DIAMON-1 mimics a visual pattern recognition process and therefore uses graphical representations of trends. However, unlike many other AI systems modelled after verbal problem descriptions, $TrenD_x$ is not rule-based; instead, it classifies data streams through a set of constraints. In that it resembles DIAMON-1 in principle.

While $TrenD_x$ may be superior in its capability to detect complex trends involving different phases of variable length, DIAMON-1's strengths are intuitive trend definition as well as detection of trends en passent, i.e., starting at any time in the input data stream. $TrenD_x$ could certainly benefit from softening its hard constraints, for example by introducing probability distributions or fuzzy sets. DIAMON-1 in turn would profit from explicitly and separately maintaining alternative hypotheses instead of tracking them all with one automaton.

Sittig et al. have worked on the detection of trends using Kalman filtering and linear regression [5, 6]. However, based on statistical models of the problem domain these approaches lead to results with an entirely different meaning [3]. DYNASCENE [1], an earlier project in which Kalman filtering was planned to be integrated, also implemented states called clinical "scenes". It is DIAMON-1's most direct competitor, however, it was not pursued further.

Conclusion

DIAMON-1 was presented as a two-stage approach to diagnostic monitoring. However, this structuring is not imperative, as in- and output of both stages is general enough to allow for easy integration of additional levels. Eventually, DIAMON-1 will evolve into an ontology the constituents of which may be arranged individually to meet problem-specific requirements rather than one general design consideration.

References

[1] Cohn, A.I., Rosenbaum, S., Factor, M., Miller, P.L. (1990) "DYNASCENE: An Approach to Computer-Based Intelligent Cardiovascular Monitoring Using Sequential Clinical 'Scenes'", Meth Inform Med 29 122–131

[2] Haimowitz, I.J. and Kohane, I.S. (1993) "Automated Trend Detection with Alternate Temporal Hypotheses", Proc. of the 13th IJCAI 146–151

[3] Pedrycz, W. (1990) "Fuzzy Sets in Pattern Recognition: Methodology and Methods", Pattern Recognition 23 121–146

[4] Shahar, Y., Tu, S.W., and Musen, M.A. (1992) "Knowledge Acquisition for Temporal Abstraction Mechanisms", Knowledge Acquisition 1:4 217–236

[5] Sittig, D.F. and Cheung, K.-H. (1992) "A Parallel Implementation of a Multi-State Kalman Filtering Algorithm to Detect ECG Arrhytmias", Int J Clin Mon Comp 9 13–22

[6] Sittig, D.F., Cheung, K.-H., and Berman, L. (1992) "Fuzzy Classification of Hemodynamic Trends and Artifacts: Experiments with the Heart Rate", Int J Clin Mon Comp 9 251–257

[7] Steimann, F. and Adlassnig, K.-P. (1994) "Clinical Monitoring with Fuzzy Automata", Fuzzy Sets & Systems 61:1 37–43 (to appear)

[8] Steimann, F. (1994) "Diagnostic Monitoring Based on Fuzzy State Transitions", Technical Report 1-94, Department of Medical Computer Sciences, University of Vienna

[9] Uckun, S., Dawant, B.M., and Lindstrom, D.P. (1993) "Model-Based Reasoning in Intensive Care Monitoring: The YAQ Approach", Artificial Intelligence in Medicine 5:1 31–48