The Interpretation of Time-Varying Data with DIAMON-1

Friedrich Steimann, Sallstraße 65, D-30171 Hannover
e-mail: 100607.704@compuserve.com

Abstract: Applying the methods of Artificial Intelligence to clinical monitoring requires some kind of signal-to-symbol conversion as a prior step. Subsequent processing of the derived symbolic information must also be sensitive to history and development, as the failure to address temporal relationships between findings invariably leads to inferior results. DIAMON-1, a framework for the design of diagnostic monitors, provides two methods for the interpretation of time-varying data: one for the detection of trends based on classes of courses, and one for the tracking of disease histories modelled through deterministic automata. Both methods make use of fuzzy set theory, taking account of the elasticity of medical categories and allowing discrete disease models to mirror the patient's continuous progression through the stages of illness.

Keywords: diagnostic monitoring, trend detection, disease tracking, fuzzy sets, automata

1. Introduction

It is a widely appreciated fact that much of the clinically relevant information is conveyed in the change and development of a patient's physiological variables—in the presence of a trend, absolute values may indeed be of subordinate significance [7]. Utilization of this information would seem imperative for diagnostic systems, as the inability to relate consecutive findings is a serious handicap that cannot be compensated for [21]. However, the modelling of time in most early diagnostic support systems—if at all existent—has been implicit [22], and so has been the account of change and development. Later adding explicit temporal awareness to an atemporal system design is nontrivial: it does not only boost complexity in the problem space, but also affects data representation and processing, usually making a complete system redesign necessary. Diagnostic monitoring systems with a full-fledged temporal account are thus still few and far between.

DIAMON-1 is a monitoring framework that has incorporated time into its data model. It allows the construction of problem-oriented diagnostic monitors out of standardized components. The components analyse sequences of explicitly time-stamped samples rather than instantaneous parameter values, thereby making monitoring sensitive to history and development [34].

As pointed out in [36], developments in a patient can be observed on different levels. At the low end of the spectrum the change in one or a number of sampled physiological variables can hint at an alteration in the patient's present condition. Here, methods of trend detection lend themselves to identifying clinically relevant developments. At the high end of the spectrum the patient's progression through the natural or therapy-induced stages of a disease is reflected in a sequence of characteristic conditions; formalized models of disease histories allow their automated tracking and prediction.

Covering both ends of the spectrum this article presents two of DIAMON-1's time-sensitive methods, one for the detection of trends and one for the tracking of disease histories, and demonstrates their effectiveness by interpreting recorded data from an on-line monitored case in critical care.

appeared in: *Artificial Intelligence in Medicine* 8:4 (1996) 343–357

2. Notation and conventions

Theories of clinical monitoring, borrowing from several disciplines, employ a wealth of concepts and notations; the ones used throughout this article are briefly presented in the following.

2.1. Variables and samples

I call all physiological parameters observed of a patient as well as personal attributes and apparatus settings *variables*. Each variable has a label and a value. All variables are time-dependent; the value of a variable x is a total function of time, $x : T \to V_x$, to the variable's value space, its instantaneous (point) value at time t_0 being denoted by $x(t_0)$.

Practically, a variable is seldom available in its continuous-time form. Instead, variables are sampled at distinct points in time, their availability thus being discrete-time. Because sampling in the clinical environment can generally not be assumed to be regular [16], the actual sampling time must be provided with the sample. I do so by writing $x[t_0]$ to denote the sample taken at t_0, indicating (by use of the square brackets) that sampling is discrete. Accordingly, a sequence of n samples is denoted by $\langle x[t_1], ..., x[t_n] \rangle$.

2.2. Fuzzy sets

A *fuzzy subset* $\tilde{a}$ of a set M is specified by its *membership function*, $\mu_{\tilde{a}} : M \to [0, 1]$, assigning a *degree of membership* to the elements of $\tilde{a}$. The membership function is a generalization of the characteristic function of ordinary sets, fuzzy subsets (called *fuzzy sets* hereinafter) thus being generalizations of ordinary sets.

Fuzzy sets are related to ordinary sets through α-cuts. The α-cut of a fuzzy set $\tilde{a}$ is defined as $\{x | \mu_{\tilde{a}}(x) \geq \alpha\}$. A fuzzy set is convex if and only if all its α-cuts are convex. [13]

Fuzzy sets are usually employed to express diffusion in the extension of concepts[1]. The degree of membership of an element in a fuzzy set then expresses the grade of compliance or compatibility of that element with the concept the set denotes. Examples of this will be given in context below.

An element of a fuzzy set is given by a pair (x, μ), where μ is the degree of membership of x in that set. Finite fuzzy sets can thus be specified by the listing of such pairs in curly brackets, as in $\{(a, 0.2), (b, 0.4), (c, 0.9)\}$.[2]

Just like ordinary sets, fuzzy sets can be combined by the set-theoretic operations union and intersection, the resulting fuzzy set's membership function being defined as the maximum of the operands' membership functions or as the minimum, respectively.

2.3. Automata

Automata specify sequences of symbols or discrete events[3]. Their regard of time is implicit in the notion of sequence: one event comes after the other, the distance between events being ignored.

Formally, a *deterministic finite automaton* is an abstract machine consisting of a finite set of distinct states, a finite set of input symbols (the events), and a transition function mapping the automaton's state and an event onto its next state [11]. The transition function connects states through transitions; a transition is triggered whenever the corresponding event occurs. Every automaton has an initial state from which it departs upon occurrence of the first event; the set

[1] Zadeh usually terms these concepts *linguistic variables* [37].

[2] Elements with zero degree of membership are usually omitted.

[3] As a rule of thumb, one speaks of symbols if sequence is spatial, and of events if sequence is temporal; the latter is the case in the context of this article.

of final states found in most definitions of automata is of no significance in the context of clinical monitoring [32].

3. Detecting trends in a sequence of samples

Although broadly used, the term *trend* does not come with one generally accepted definition. In the biomedical field, a trend has been defined as

- "presence of a slow, consistent, unidirectional change in the value of a variable" [6];
- "general direction of the mean level of the data" [2];
- "any change in the underlying signal dynamics slower than the system's time constants" [3];
- "steadily rising or steadily falling pattern" [8]; and
- "clinically significant pattern in a sequence of time-ordered data" [18].

These informal definitions suggest that a trend is either regarded as a feature that can be extracted from or as a pattern that can be recognized in a signal. I will give a different, more stringent definition below.

Like the term trend itself the problem of trend detection is a very general one. It can be formulated as "identify all the trends occurring in a signal together with their times of onset". If the trend specifications are parameterized, trend detection may also derive the parameters as part of the answer.

The general problem of trend detection has several more specific instantiations, two of them being "which trend is the signal currently following?" and "did a certain trend occur in the signal and when was its time of onset?". While the latter is the subject of another article [33], the former is characteristic of the critical care environment and will be dealt with in the following.

3.1. Classification of sequences of samples based on sets of absolute courses

Classification of instantaneous quantitative observations based on intervals (convex sets), fuzzy or non-fuzzy, is easy to specify and implement; because it has served its purpose well in several prominent projects (including the non-fuzzy MYCIN [28] and the fuzzy CADIAG-2 [1]), adding a temporal dimension to cover the course of a variable seems a worthwhile extension. Following this idea I define a *trend* as a *set of courses or continuous-time functions having identical meaning with respect to the given monitoring problem*.

The detection of a trend in discrete-time sequences of samples usually relies on one implicit assumption, namely that the variable does not deviate significantly from the hypothesized trend at times at which it is not sampled. In other words, it is assumed that the observed variable follows a continuous-time course that both explains the sampled findings and is compliant with the trend to be detected. To make this assumption explicit and for reasons of derivability (see below) I base the trend detection in sequences of samples on the classification of continuous-time courses.

With the above definition of trend at hand, a given *sequence of samples is compatible with a trend* if there exists a continuous-time course going through all samples (called an *explanatory course* hereafter) which is also a member of the trend. More formally, if the trend is defined as a fuzzy set $\widetilde{C}$, the degree of compatibility, γ, of a sequence of samples, $\langle x[t_1], ..., x[t_n]\rangle$, with $\widetilde{C}$ is defined by

$$\gamma = \sup_{x_e(t)} \mu_{\widetilde{C}}(x_e(t)), \quad x_e(t_i) = x[t_i] \text{ for all } 1 \le i \le n \quad .^4 \tag{1}$$

[4] Throughout this discourse I ignore the temporal extent (duration) of a trend and its time of onset in the potentially infinite sequence of samples for the sake of simplicity; refer to [34] for a complete account.

The explanatory course required in the definition of compatibility presents a guess about the true course of the variable. Trend detection as defined in (1) is optimistic in that it chooses the explanatory course best suiting the trend to be detected. Quite clearly, if sampling is sparse, trend detection is subject to uncertainty and hence speculation: if gaps between samples allow a variable to have taken significantly differing courses, the method cannot differentiate trends competing in the explanation of the findings. Consequently, if these trends are mutually exclusive, they must be understood as alternative hypotheses, differentiation between which being left to other methods or the user.

A trend $\widetilde{C}$ is specified by its membership function. This membership function can be constructed using a parametrized description of the trend's elements, the courses, as in

$$\widetilde{C} = \{(x(t), \mu) | x(t) = mt + b, m, b \in \Re, \mu = \min(\mu_{\widetilde{m}}(m), \mu_{\widetilde{b}}(b))\} ,$$

specifying the fuzzy set of linear courses with some fuzzy slope $\widetilde{m}$ and intercept $\widetilde{b}$. However, naturally occurring courses rarely lend themselves to parametrization. Instead, the following definition of fuzzy courses allows intuitive specification of trends and makes classification of sequences of samples straightforward.

Let $\widetilde{x}(t)$ be a function mapping time to convex fuzzy subsets of a variable's range V_x. I call $\widetilde{x}(t)$ a *fuzzy course*; it may be viewed as a fuzzy band denoting the allowable spread in the course of a variable still considered following a trend. The trend $\widetilde{C}$ is then defined as the set of courses whose points all lie within the range of $\widetilde{x}(t)$; specifically,

$$\widetilde{C} = \{(x(t), \mu) | \mu = \inf_{t_0} \mu_{\widetilde{x}(t_0)}(x(t_0))\} . \qquad (2)$$

As shown in [34], the degree of compatibility γ of a sequence of samples $\langle x[t_1], ..., x[t_n] \rangle$ with $\widetilde{C}$ is then determined by the least degree of membership of all $x[t_i]$ in $\widetilde{x}(t_i)$, i.e.,

$$\gamma = \min_{1 \le i \le n} \mu_{\widetilde{x}(t_i)}(x[t_i]) ,$$

because the explanatory course $x_e(t)$ required in (1) can always be constructed so that

$$\mu_{\widetilde{x}(t_0)}(x_e(t_0)) \ge \mu_{\widetilde{x}(t_i)}(x_e(t_i)) = \mu_{\widetilde{x}(t_i)}(x[t_i])$$

for t_0 in proximity of each t_i.

This framework has a nice graphical metaphor depicted in Fig. 1: if the fuzzy course $\widetilde{x}(t)$ is viewed as a tunnel extending in time and whose height at any point (t_0, x_0) is given by its membership function's value $\mu_{\widetilde{x}(t_0)}(x_0)$, then γ is the clearance of the best path going through all $x[t_i]$, which is equal to the minimum of the heights at all $(t_i, x[t_i])$ with $1 \le i \le n$.

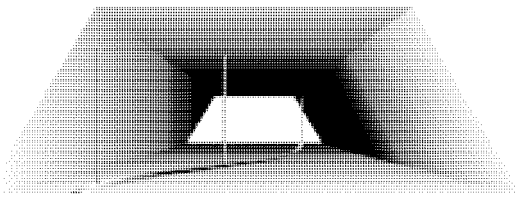

Figure 1: Trend detection visualized: the tunnel represents the trend, the vertical bars denote the samples and their respective degrees of membership, and the rail corresponds to the explanatory course; the clearance (minimum height) above the explanatory course is the degree of compatibility of the findings with the trend

The feasibility of the approach is demonstrated on a recorded data set taken from a case of adult respiratory distress syndrome (ARDS) in an eight-month-old female. For a period of 12 hours, several physiological variables including arterial oxygen saturation (SaO2), mean arterial blood pressure and heart rate were sampled and recorded in intervals of approximately 20 seconds.

ARDS, a form of acute respiratory insufficiency, requires substantial respiratory support. Intervening mechanical ventilation, this particular patient was hand-ventilated (with a hand bag) approximately every two hours for roughly 15 minutes; during that time, the fraction of inspired oxygen (FIO2) was increased from 50 to 100 per cent.

One expected effect of such a drastic increase in FIO2 is a corresponding increase in oxygenation. Oxygenation as measured by SaO2, however, is naturally bounded by 100 per cent. Because ventilator therapy aims at (and, in the given case, succeeded in) keeping SaO2 above 90 per cent, the variable's potential for change is within fixed bounds. The expected increase in SaO2 is modelled by the fuzzy course shown in Fig. 2 a).

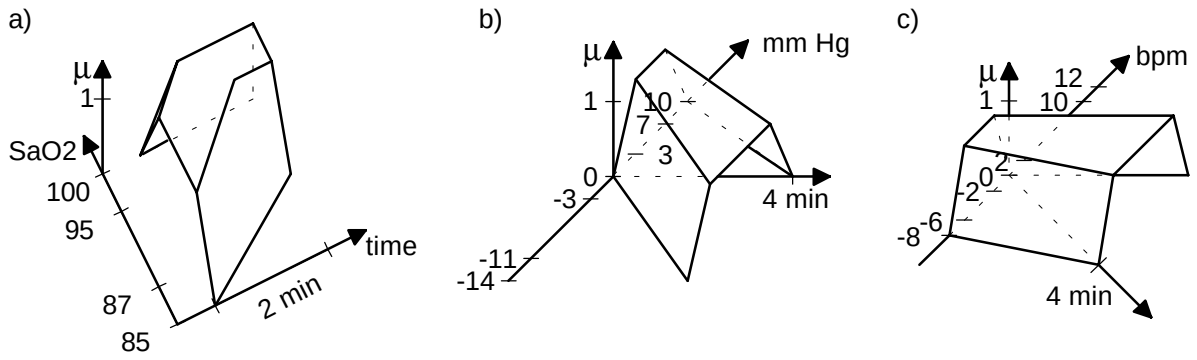

Figure 2: Fuzzy courses specifying trends
a) *sharply rising SaO2*; b) *falling blood pressure*; c) *increasing heart rate*

Based on this expectation the trend detection method produces the results shown in Fig. 3. It fully recognizes sharply rising SaO2 seven times, six times at the beginning of a hand-bagging session and once during mechanical ventilation.

Figure 3: detection of a sharp increase in oxygenation; upper frame shows SaO2 in the range of 80–100 per cent, lower frame shows degree of compatibility with the trend; vertical dotted bars indicate periods of hand-bagging (high FIO2)

The fact that trend detection based on absolute courses works well for the expected increase in SaO2 is partly due to the fact that it is naturally bounded by 100 per cent as a landmark value. Other trends lack this absolute orientation; a method appropriate for their detection will be presented next.

3.2. Floating level trend detection

Trends specified in terms of absolute courses constrain absolute variable values rather than relative changes, a property that makes them unsuited for many clinically relevant patterns. For example, the trend "heart rate stable for ten minutes" is characterized by a fairly constant course of the heart rate, no matter at which level.

Mathematically, the first derivative is a measure for the rate of change in a signal. A trend may thus more adequately be specified in terms of its courses' first derivatives. However, a sequence of samples is not continuous, and, unless the sampling frequency is high enough to allow (near) perfect reconstruction of the original signal, the sequence's derivative cannot be derived. Matching a sequence of samples with a trend specified through its courses' first derivatives thus builds on explanatory courses: if there is a continuous-time course which explains all findings and whose first derivative satisfies the trend's specification, then the findings

are compatible with the trend. Obviously, such an explanatory course is not as easily constructed as in the absolute case; yet, there is a solution with computational effort linear in the number of samples as long as the defining fuzzy course satisfies certain conditions [34].

However, there is a simpler solution to the problem: if the change in a signal making up a trend can be specified in absolute terms, the trend detection method based on absolute courses can be adapted to match signals at variable levels, leaving the representation of trends unaltered. All that needs to be done is to derive an offset that when added to the trend yields the best match.

The method works as follows. Assume that $x[t_1]$ is the first sample to be matched with a trend specified as in (2) by a fuzzy course $\tilde{x}(t)$. $\mu_{\tilde{x}(t_1)}(x[t_1])$ then denotes the degree of compatibility of that sample with the trend. Obviously, there is always an offset b making the match perfect so that $\mu_{\tilde{x}(t_1)}(x[t_1]-b)=1$. In fact, the degree of compatibility of the one-sample sequence with the trend can be specified as a function of b, namely

$$\mu_{\tilde{b}_1}(b) := \gamma = \mu_{\tilde{x}(t_1)}(x[t_1]-b),$$

with t_1 and $x[t_1]$ being constants. $\mu_{\tilde{b}_1}(b)$ specifies the fuzzy set of offsets, $\tilde{b}_1$, making a match.

Clearly, the same holds for the second sample, $x[t_2]$, so that for the pair of samples, $\langle x[t_1], x[t_2]\rangle$, the set of offsets resulting in a positive match is given by the intersection $\tilde{b}_1 \cap \tilde{b}_2$, the offsets suiting both the first and the second sample. The degree of compatibility for any chosen offset b is then given by

$$\gamma = \mu_{\tilde{b}_1 \cap \tilde{b}_2}(b).$$

It follows that for any sequence of samples $\langle x[t_1], ..., x[t_n]\rangle$ and $\tilde{b} = \tilde{b}_1 \cap ... \cap \tilde{b}_n$ the degree of compatibility of the samples with the trend and offset b is $\mu_{\tilde{b}}(b)$. Consequently, for the matching to produce the best result, b only needs to be chosen so that $\mu_{\tilde{b}}(b)$ is maximum.

While scanning over the findings, this trend detection method continuously adapts the offset of the fuzzy course (thereby keeping its level floating, hence its name) so it best matches the findings. Surprisingly, doing so adds only little computational effort; in fact, given that fuzzy set intersection can be done in constant time, computational effort is still linear in the number of samples, which must be attributed to the fuzzy set representation of the offset b.

Again, the practicability of the approach is demonstrated on the ARDS case. As pointed out in [18], blood pressure and heart rate of the ARDS patient show noticeable fluctuations. In particular, several coordinated drops in blood pressure and increases in heart rate can be observed. Because these changes set off from widely differing levels, they are typical candidates for floating level trend detection. Fig. 4 shows the outcome of this procedure based on the fuzzy courses of Fig. 2 b) and c). Note how most of the events are detected during hand-bagging.

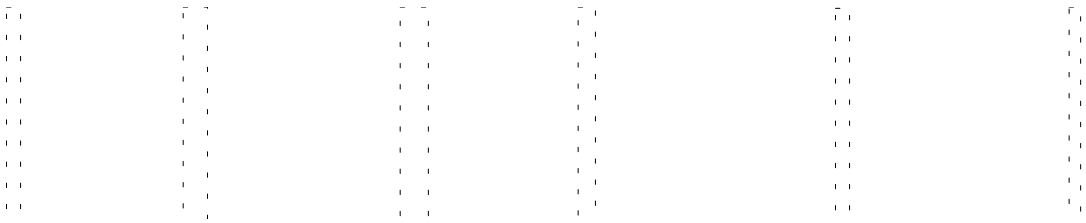

Figure 4: Detecting changes in 12-hour period with floating level trend detection (input smoothed by a median filter); range of blood pressure is 50–90 mm Hg, range of heart rate is 150–190 bpm

4. Tracing developments on the symbolic level

The development and course of a disease is sometimes described as a sequence of characteristic stages, possibly with alternative branches and turning points. Such a disease naturally lends itself to being modelled by a state transition diagram. A patient's progression through the stages can then be traced by means of a corresponding finite state machine or automaton; a state of the automaton corresponds to the stage of the disease or status of the patient, a transition links two potentially consecutive stages, and the event causing the transition corresponds to a clinically significant condition or development found in the patient.

In cases where the pathogenesis is so simple that it does not require a model the employment of automata may still be worthwhile: whereas the input of a monitor is usually rather unsteady, its output, the diagnosis, should show temporal stability. In particular, a minor fluctuation in the input should not cause a diagnosis to be revised. The required inertia in behaviour—also known as *hysteresis*—can easily be implemented by means of automata; it is achieved if the reversal of transition from one state to another, if at all allowed, is triggered by differing findings.

The automaton as a disease model is flawed by its digital nature: upon fulfilment of a condition the current state of the automaton leaps from one state to the next, discretely and without any indication of the forthcoming event. This abrupt change is not natural; rather, most transitions take their time, taking place gradually and continuously. Indeed, an appropriate framework of disease tracking would have to

- report the tendency of the patient to change state and the change's continuous progress to the observer so countermeasures can be taken as soon as deemed necessary (smooth transitioning)[5], and
- report and record mild occurrences of pathophysiological states and regard them in the derivation of future states of the patient (gradation in severity).

The latter is particularly important as not all patients exhibit the full symptomatology of a disease, and mild occurrences may be just as noteworthy as fully developed.

Smooth transitioning requires intermediate states, snapshots of the patient's being in between states. An intermediate state can be specified by means of a fuzzy subset of the automaton's set of states. Such a fuzzy state is reached upon partial fulfilment of the condition triggering the transition, specified in the form of a fuzzy event. For example, a state between *normal* and *reduced blood pressure* can be specified by {(*normal blood pressure*, 0.4), (*reduced blood pressure*, 0.6)}. It is reached as the result of the automaton formerly being in the state {(*normal blood pressure*, 1)} and then transitioning on detection of the trend *falling blood pressure* with a degree of compatibility of 0.6. Even if there is no physiological measure to determine the patient's true status relative to the states of the model, continuous transition from one state to the next as implied by the increasing fulfilment of the triggering condition could well be visualized as a spatial transition in the diagram, the interpretation being left to the observer [34].

Formally, a fuzzy automaton is obtained by application of the extension principle [13] to the automaton's transition function. The extended function maps a fuzzy state and input onto the next fuzzy state [32]. The advantage of this approach over other fuzzy automata (e.g. [11]) is that the specification of the automaton itself is not fuzzy—the fuzziness is brought into the automaton solely through its fuzzy input. Disease models can thus be designed in their ideal form without worrying about fuzzy aspects.

[5] This is particularly true for (closed-loop) therapy control, where early low-dosed counteraction can help to keep the patient stable by avoiding drastic measures and, consequently, reactions.

The employment of fuzzy automata in the dense data environment of critical care requires a few adjustments described in [32, 34]. The adjustments include measures to maintain a full degree of membership of at least one state in the current (fuzzy) state in order to avoid membership depletion and oscillation. In case of the following example, the automaton has been modified to leave its current state only if its successor state has gained full set membership.

Revisiting the previous example (Fig. 4), the falling blood pressure during hand bagging can be attributed to reduced venous return as a result of increased intrathoracic pressure. This interpretation suggests that hand bagging is a little too vigorous. The rise in heart rate observed at the same time can be explained as a compensatory mechanism maintaining cardiac output; whereas a rise in heart rate alone may be indicative of some other perturbation, in this case it is a direct consequence of the drop in blood pressure and thus indirectly linked to hand bagging.

The state transition diagram of an automaton classifying a rise in heart rate following falling blood pressure during hand bagging as compensating is shown in Fig. 5; apart from two cardiovascular states, it also maintains states to indicate the period of mechanical ventilation (*idle*) and that of unperturbed hand bagging (*alert*).

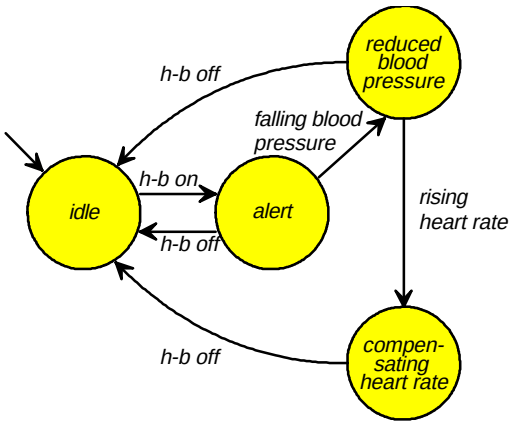

Figure 5: Finite automaton monitoring the cardiovascular response to hand bagging (initial state is *idle*)

When fed with the output of Fig. 4 complemented by two variables indicating the onset (*h-b on*) and ending of hand bagging (*h-b off*), the automaton's state changes as shown in Fig. 6. Each line depicts over time on a scale from 0 to 1 the degrees of membership of the corresponding state in the automaton's sequence of fuzzy states.

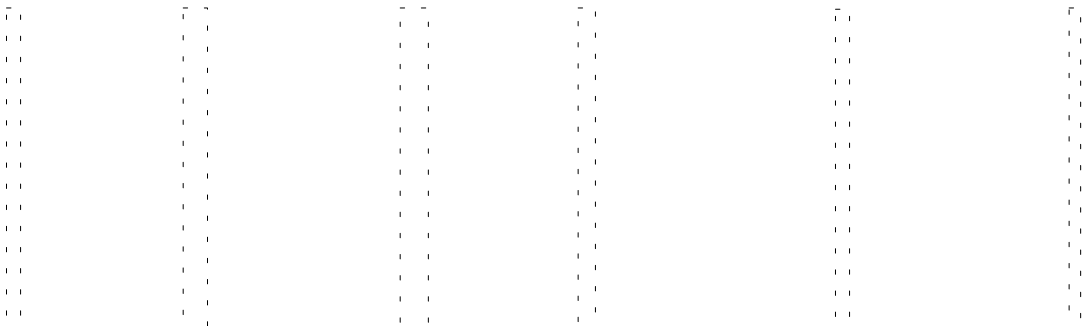

Figure 6: Interpretation of the trends of Fig. 4 through the automaton of Fig. 5

During the first four episodes of hand bagging the patient exhibits reduced blood pressure as well as a compensating heart rate, albeit the latter only to a certain degree during the second and third episode (gradation in severity). Note how one can observe the state *reduced blood pressure* building up slowly during session number two, reflecting a smooth transition to that state. The falling blood pressures in context of hand bagging session number five and six do

not trigger a transition because the trend is detected only after hand bagging has ended, the automaton therefore returning to *idle* before a transition to *reduced blood pressure* can take place. Note the temporal stability of the output as a sign of effective information reduction.

The chosen example is debatable to some respect, as the automaton does not model the behaviour of *one* variable; rather, it combines the internal states of a device specially designed to avoid multiple alarms pertaining to the same cause with the values of an abstract physiological variable *cardiovascular status*. If needed, the automaton can be split into two, the state of the first (tracking the mode of ventilation) serving as input for the second.

5. Related work and discussion

Performing signal-to-symbol conversion and state-based sequence analysis, the presented methods are at the intersection of signal processing, time series analysis, and AI. Although the output of both is symbolic, it may diverge from what one might expect: rather than producing propositions such as "temperature rising" assigned to temporal intervals, the output too has signal character, i.e., consists of a sequence of samples, only that the samples are on a higher level of abstraction. This is a required property within the framework of DIAMON-1, where all systems' output must be in a form consumable by other systems as input. If desired, interval-based abstractions can be generated from the output, for example by use of temporal abstraction methods such as the ones presented in [20, 27].

The methods presented find analogies in several other monitoring projects discussed in the following. It must be understood, however, that formal comparison is difficult, as each project focuses on its own set of problems. In fact, the problems encountered in clinical monitoring are so numerous that no single approach can be expected to solve all of them.

VM [15] was not only one of the first attempts to tackle the clinical monitoring problem, it also made use of the basic concepts picked up in this paper: the classification of time-varying data by extending parameter intervals by a temporal dimension; and the employment of finite state machines to model admissible state transitions. In fact, the TIMEEXP premise function of VM is a special case of trend detection as presented in this article. VM's finite state machine, however, is used to specify possible changes in therapy, namely the change of ventilation modes. This is in contrast to the tracking of disease histories put forward in this article.

Direct high-level diagnostic use of automata-like structures has been made in DYNASCENE [9] and ICM [14]. Both systems identify clinically meaningful constellations and interpret the sequence of their occurrence. However, they both lack concepts of smooth progression or gradual illness as made possible by fuzzy automata. As a result, their all-or-nothing nature renders them likely to miss out or overvalue developments in individual patients (sensitivity/specificity trade-off).

Other approaches use Chomsky-type grammars to classify temporal patterns that can be coded as a sequence of discrete events, for example [5, 19]. In [5] a context-free grammar is employed to detect heart arrhythmia by analysing sequences of fuzzy beat labels. The fuzzy beat labels are derived in a fuzzy classification process described in [4]; the result of parsing is the certainty with which a derived sequence of labels belongs to an arrhythmia. Another two-level approach to real-time diagnosis and control, in this case of fermentation processes, is taken in [23]; the upper level, basically a finite state machine, traces the stage of the fermentation process and so determines the choice of applied control strategies. The transition between stages is triggered by the detection of trends via temporal shape analysis [24].

Generally, the employment of automata in monitoring has several benefits. It allows conclusions only to be drawn in consequence of others, thus making interpretation of findings sensitive to history. This may be necessary in situations where the same condition found in a

patient can have different meanings depending on what happened before; the increase in heart rate as a consequence of a drop in blood pressure is an example of this. Also, as shown in [34], automata can reduce the volatility in the output of a monitor, producing temporally stable interpretations of unsteady input.

Last but not least, the history an automaton has encountered is comprised in its current state. This is particularly beneficial if the monitor operates in a transient display mode: while basic variables such as blood pressure only report momentary readings, the current state of an automaton conveys the significant events and turning points of the past, all at one glance.

However, all approaches based on automata suffer from their limited expressive power. Sequence is their only temporal relationship, ignoring the fact that the temporal interval between findings can be of great importance. YAQ [35] overcomes this deficiency by linking state transitions to explicit history lookups which can reference past values and query additional temporal information. The price being paid is that the regular formalism is left, requiring all history lookups to be hard-coded. Attributed formalisms such as augmented transition networks may be a way out of this dilemma.

Some work has also been done at the intersection of high-level disease tracking and low-level trend detection. For example, some approaches regard the course of a perturbation as a sequence of primitive trend segments. The problem of trend detection is then decomposed into two subproblems:

1. the detection of trend primitives; and
2. the identification of the transition from one primitive to the next.

GUARDIAN employs a real-time, on-line segmented trend detection method based on a fuzzy temporal pattern recognition (TFPR) [12]. TFPR performs a strictly sequential, non-optimal segmentation of trends based on local maxima in the matching process. Specifically, TFPR considers a segment switch for each new sample based on a comparison of its degree of membership in the current and the next segment. An expected higher degree of match gained by a segment switch, however, may turn out to be a dead end, in which case continuation of the previous segment might have produced a better result. Despite this imperfection, the optimal segmentation algorithm presented in the same work is considered too expensive to be used in practice.

Another sophisticated trend detection method addressing the temporal variability and segmentation of trends is TRENDX [18]. It uses so-called trend templates to define and detect trends in sequences of samples. A trend template is a collection of temporal intervals each of which constrains a number of parameters (either through value constraints or by regression, see below). 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 in TRENDX by assigning samples to suitable intervals. For this purpose, TRENDX maintains competing hypotheses, instantiations of trend templates whose intervals' parameters are set to fit the samples' times and values. Each hypothesis is then a possible interpretation of the findings.

It may be objected that the segmentation of trends, although an appealing theoretical issue, renders trend detection a highly time-consuming task, the practical necessity of which is not entirely plausible. Indeed, in [33] it is shown that a single fuzzy course can cover a wide spectrum of individual developments including considerable variability in the duration of the trend.

The trend detection method based on fuzzy courses has an obvious competitor in statistics: regression. Regression is made the basis of trend detection in several monitoring projects, for example [24, 25, 31, 18]. Interestingly, in earlier versions of TRENDX function-based value constraints similar to fuzzy courses, only non-fuzzy, were employed [17]. A more recent

version favours low order regression, and the author notes that regression-based TRENDX, yielding a gradual measure of fit, is more robust and allows ranking of hypotheses, as opposed to its constraint-based predecessor, where a single sample out of bounds sufficed to reject a hypothesis [18]. It may be added that using fuzzy constraints such as fuzzy courses would have had a similar effect.

There is a certain relationship between trend detection via regression and trend detection via fuzzy courses as defined above. Suppose that $x_a(t)$ is a polynomial with fixed coefficients and that

$$E = \max_n |x_a(t_n) - x[t_n]|$$

is the measure of deviation. Then

$$x_a(t) + \widetilde{E},$$

where $\widetilde{E}$ is a convex fuzzy set of reals symmetrically centred around zero, defines a fuzzy course, and

$$\mu_{\widetilde{E}}(E)$$

is the degree of compatibility of a sequence with the trend. However, the reader will agree that this has little in common with common regression. Besides, as suggested by the tunnel metaphor, trend detection based on fuzzy courses puts forward a graphical approach, where specification of a trend is intuitive and highly flexible at the same time.

Kalman filters and their multi-state extensions have also successfully been employed to detect trends in biomedical signals, for example in [7, 16, 26, 29, 30]. It appears that Kalman filters are particularly good at very early detection of a trend (or, rather, deviation from a trend), a property they have in common with more primitive forecasting methods such as cumulative sums [2, 3]. This makes them particularly suited for the critical care environment, where other trend detection methods including the ones based on regression are certainly slower to react. However, because of the filter process model's recursive definition involving a limited past, the effectiveness of trend detection through Kalman filtering is restricted to domains where a trend manifests itself in very few samples, and not in overall developments involving long sequences of samples. By contrast, trend detection based on fuzzy courses has been shown to work well in all cases [34].

Finally, it must be noted that none of the discussed approaches properly address the problem of superposition and interacting diseases: an observed variable may be influenced by different underlying developments the effects of which superimpose on one another. Sophisticated diagnostic monitoring systems should be able to take account of possible superposition, for example by utilizing correlation with other, non-superimposed variables or by employing explicit models of superposition. The latter is explored in [10]; however, because it is tightly geared to qualitative simulation, its potential influence on quanitative or symbolic methods is rather limited.

6. Acknowledgements

Much of the presented work has been conducted at the Department of Medical Computer Sciences, University of Vienna. I thank Klaus-Peter Adlassnig for introducing me to the fuzzy set theory and for his demonstration of its aptness.

7. References

[1] K.P. Adlassnig, G. Kolarz, W. Scheithauer, H. Grabner, Approach to a hospital-based application of a medical expert system, *Medical Informatics* 11 (1986) 205–223.

[2] R. Allen, Time series methods in the monitoring of intracranial pressure part I: problems, suggestions for a monitoring scheme and review of appropriate techniques, *J Biomed Eng* 5 (1983) 5–17.
[3] R.K. Avent, J.D. Charlton, A critical review of trend-detection methodologies for biomedical monitoring systems, *Critical Reviews in Biomedical Engineering* 17 (1990) 621–659.
[4] S. Barro, R. Ruiz, J. Mira, Fuzzy beat labeling for intelligent arrythmia monitoring, *Comp Biomed Res* 23 (1990) 240–258.
[5] S. Barro, R. Ruiz, J. Presedo, J. Mira, Grammatic representation of beat sequences for fuzzy arrythmia diagnosis, *Int J Biomed Comput* 27 (1991) 245–259.
[6] J.E.W. Beneken, J.A. Blom, F.F. Jorritsma, A. Nandorff, J. Spierdijk, Prognosis, trend and prediction in patient management, *J Biomed Eng* 1 (1979) 185.
[7] J.A. Blom, J.A.F. de Ruyter, N. Saranummi, J.E.W. Beneken, Detection of trends in monitored variables, in: E.R. Carson, D.G. Cramp, eds., *Computers and Control in Clinical Medicine* (Plenum Press, New York 1985) 153–174.
[8] R.E. Challis, R.I. Kitney, Biomedical signal processing (in four parts). Part 1: time domain methods, *Medical & Biological Engineering and Computing* 28 (1990) 509–524.
[9] A.I. Cohn, S. Rosenbaum, M. Factor, P.L. Miller, DYNASCENE: an approach to computer-based intelligent cardiovascular monitoring using sequential clinical 'Scenes', *Methods of Information in Medicine* 29 (1990) 122–131.
[10] E.W. Coiera, Monitoring diseases with empirical and model-generated histories, *Artificial Intelligence in Medicine* 2 (1990) 135–147.
[11] E.R. Dougherty, C.R. Giardina, *Mathematical Methods for AI and Autonomous Systems* (Prentice Hall, Englewood Cliffs, 1988).
[12] J. Drakopoulos, B. Hayes-Roth, *A Sigmoidal Based Fuzzy Pattern Recognition System for Multi-Variate Time-Dependent Patterns*, Stanford University, Stanford 1993.
[13] D. Dubois, H. Prade, *Fuzzy Sets and Systems* (Academic Press, New York, 1980).
[14] M. Factor, D.H. Gelernter, D.F. Sittig, The multi-trellis software architecture and the Intelligent Cardiovascular Monitor, *Methods of Information in Medicine* 31 (1992) 44–55.
[15] L.M. Fagan, J.C. Kunz, E.A. Feigenbaum, J.J. Osborn, Extensions to the rule-based formalism for a monitoring task, in: B.G. Buchanan, E.H. Shortliffe, eds., *Rule-Based Expert Systems* (Adisson-Wesley, Reading, MA, 1984) 397–423.
[16] K. Gordon, The multi-state Kalman Filter in medical monitoring, *Computer Methods and Programs in Biomedicine* 23 (1986) 147–154.
[17] I.J. Haimowitz, I.S. Kohane, Automated trend detection with alternate temporal hypotheses, in: *Proceedings of the 13th International Joint Conference on Artificial Intelligence IJCAI-93* (Chamberey, France, 1993) 146–151.
[18] I.J. Haimowitz, *Knowledge-Based Trend Detection and Diagnosis*, PhD Thesis, MIT, Cambridge, MA, 1994.
[19] M. Juhola, T. Grönfors, A scheme of inference of regular grammars for the syntactic pattern recognition of saccadic eye movements, *Artificial Intelligence in Medicine* 3 (1991) 87–94.
[20] M.G. Kahn, Combining physiological models and symbolic methods to interpret time-varying patient data, *Methods of Information in Medicine* 30 (1991) 167–178.
[21] M.G. Kahn, Modelling time in medical decision support programs, *Medical Decision Making* 11 (1991) 249–264.
[22] E. Keravnou, Medical temporal reasoning, guest editorial, *Artificial Intelligence in Medicine* 3 (1991) 289–290.

[23] K.B. Konstantinov, T. Yoshida, A knowledge-based pattern recognition approach for real-time diagnosis and control of fermentation processes as variable structure plants, *IEEE Transactions on Systems, Man, and Cybernetics* 21 (1991) 908–914.
[24] K.B. Konstantinov, T Yoshida, Real-time qualitative analysis of the temporal shapes of (bio)process variables, *AIChE Journal* 38 (1992) 1703–1715.
[25] E.M.J. Koski, A. Mäkivirta, T. Sukuvaara, A. Kari, Development of an expert system for haemodynamic monitoring: computerized symbolization of on-line monitoring data, *International Journal of Clinical Monitoring and Computing* 8 (1992) 289–293.
[26] P. Penczek, W. Grochulski, J. Grzyb, M. Kowalczyk, The use of a multi-chanel Kalman filter algorithm in structural analysis of the epileptic EEG, *Int J Bio-Medical Computing* 20 (1987) 135–151.
[27] Y. Shahar, M.A. Musen, Knowledge-based temporal abstraction in clinical domains, *Artificial Intelligence in Medicine* (in press)
[28] E.H. Shortliffe, *Computer-Based Medical Consultation: MYCIN* (Elsevier, New York, 1976)
[29] D.F. Sittig, M. Factor, Physiologic trend detection and artifact rejection: a parallel implementation of a multi-state Kalman filtering algorithm, *Computer Methods and Programs in Biomedicine* 31 (1990) 1–10.
[30] D.F. Sittig, K.H. Cheung, A parallel implementation of a multi-state Kalman filtering algorithm to detect ECG arrhytmias, *International Journal of Clinical Monitoring and Computing* 9 (1992) 13–22.
[31] D.F. Sittig, K.H. Cheung, L. Berman, Fuzzy classification of hemodynamic trends and artifacts: experiments with the heart rate, *International Journal of Clinical Monitoring and Computing* 9 (1992) 251–257.
[32] F. Steimann, K.P. Adlassnig, Clinical Monitoring with fuzzy automata, *Fuzzy Sets and Systems* 61 (1994) 37–42.
[33] F. Steimann, M. Hayde, B. Panzenböck, K.P. Adlassnig, A. Pollak, Fuzzy support for serodiagnosis: the ONSET program, *IEEE Engineering in Medicine and Biology Magazine* 13 (1994) 705–709.
[34] F. Steimann, *Diagnostic Monitoring of Clinical Time Series*, Doctoral Dissertation, Technical University of Vienna, Austria, 1995.
[35] S. Uckun, *An Ontology for Model-Based Reasoning in Physiological Domains*, PhD Thesis, Vanderbilt University, Nashville, Tennessee, 1992.
[36] S. Uckun, Intelligent systems in patient monitoring and therapy management, *International Journal of Clinical Monitoring and Computing* 11 (1994) 241–253.
[37] L.A. Zadeh, The concept of a linguitic variable and its application to approximate reasoning, *Information Sciences* 8 (1975) 199–249.