Intelligent alarms for patient supervision

A. Otero, P. Félix[1], F. Palacios[2], C. Pérez-Gandía, and COS. Sorzano
Dept. of Software and Knowledge Engineering. University San Pablo CEU. 28668 Madrid, SPAIN (abraham@dec.usc.es).
[1]Dpto. de Electrónica e Computación, Universidade de Santiago de Compostela.
[2]Universitary Hospital of Getafe. 28905 Madrid, SPAIN

Abstract **– This work presents a proposal for creating intelligent alarms that offer more efficient support to medical staff than the threshold alarms currently available in commercial monitors. Our alarms make it possible to handle the uncertainty and imprecision that are characteristic of the medical domain, reason over the temporal evolution of physiological variables, and incorporate information from a number of variables into one single alarm.**

The proposal is based on a structural pattern recognition model (the MFTP model) which allows certain monitoring criteria in a computational representation to be captured and identified over the evolution of patients' physiological variables. The description of the criteria is carried out using a graphical tool (TRACE), which is sufficiently intuitive to be used by physicians without the need for assistance.

Keywords **– Cognitive Overload, Intelligent Alarms, Fuzzy Constraint Satisfaction Problems, Patient Monitoring.**

I. INTRODUCTION

Advances made in electronic monitoring devices, and in the information and communications technologies, place an increasing amount of data on patients at the disposal of health-care personnel. In the bibliography it is a recognized fact that if the volume of data available exceeds the cognitive capabilities of medical staff, this situation may indeed be counter-productive [1], [2], [3], [4], as they may be forced to ignore some of the data in the decision-making process. Of all the information available on patients (x-rays, ultrasound scans, laboratory analyses, data from examinations, etc.) that which places the greatest burden on health-care staff is information from the analysis of physiological variables: electrocardiogram, heart rate (HR), oxyhemoglobin saturation (SatO2), respiratory rate (RR), blood pressure (BP), ST deviation, etc. These variables evolve over time, and often physiopathological variables appear over them, requiring rapid intervention to reduce or avoid life-threatening situations for the patient. Thus they require continuous attention.

The only support health-care staff have for monitoring the physiological variables are threshold alarms in commercial monitors; these are triggered each time the value of a parameter leaves a preestablished range. The signals usually have high levels of artifacts, often due to the movement of patients, resulting in a high number of false positives. Consequently, health-care staff may lose faith in threshold alarms, and not respond as quickly as they could in situations where intervention is really required and, in extreme situations, they may ignore and even disconnect the alarms [5], [6]. On the other hand, establishing the ranges entails searching for a balance between sensitivity and specificity, which would keep the number of false positives within reasonable limits. With these ranges it is often not possible to monitor all events that may be indicative of possible life-threatening situations for the patient; hence, the limitations of these alarms must be offset by the continuous supervision of the health-care staff.

This situation results in an increasing imbalance between the volume of data available on patients and improvements in health care quality that these data produce. There is an increasing need for a new generation of intelligent alarms, supplying more efficient support in the task of monitoring pathological signs over the physiological variables of patients and, of course, allowing the increase in the volume of data to come to fruition in the form of improved health care.

In the following section we analyze the characteristics required by this new generation of alarms, and, in Section 3 we describe a proposal (comprising the MFTP model [7] and the tool TRACE [8]) for creating alarms with these characteristics. Section 4 presents the results obtained when various alarms created with our proposal were applied over a total of 175 hours of recordings of parameters from 71 patients. Finally, we present the salient conclusions from this work.

II. A NEW GENERATION OF INTELLIGENT ALARMS

The enhancements presented below are based on the experience accumulated over 20 years of solving patient-monitoring problems. They supply two main benefits: reducing the number of false positives in threshold alarms, and supplying a greater amount of diagnostic evidence on the illness being monitored.

1-4244-0830-X/07/$20.00 ©2007 IEEE.

A. Supplying a degree of certainty

The threshold alarms currently in use have an all-or-nothing behaviour. This contrasts with the nature of illnesses in the clinical domain: often their presence or absence cannot be considered as a binary problem, rather as a question of degree [10]. The alarms, which aim to automatically identify manifestations of pathologies, must show the degrees between the maximum certainty of the absence of a given sign, and the maximum certainty of its presence. Using certain artificially precise criteria in alarm definition can lead us to commit important errors when evaluating a set of findings that are on the borderline between values that are clearly normal and those which are not.

The alarms must make it possible to handle the uncertainty and imprecision that are characteristic of the medical domain, reflecting the gradual transition between those states considered to be normal, and those considered to be abnormal. To this end, they should be based on imprecise criteria and, using visual metaphors that are intuitive for health-care personnel, they should supply a measurement of the certainty of the presence of the signs that they monitor. Given that definitions of normality and abnormality often rely on experience-based heuristic knowledge, fuzzy set theory - a tool which has proved its value for suitably representing and handling this type of knowledge - would seem to be one of the most suitable solutions for the construction of these alarms [5], [6], [7], [10], [11], [12].

B. Reasoning over the temporal evolution of parameters

One of the principal limitations of the alarms that are currently available, and which gives rise to a great number of false positives, is that of restricting their activation to the instantaneous value of a determined physiological parameter, and to the membership or not of a range of normality. Thus, any artifacts producing a value outside the range will trigger the alarm. By employing knowledge on the dynamics of the parameters and reasoning over their temporal evolution, it would be possible to identify any inconceivable rates of change in the corresponding physiological variables and identify them as artifacts. By way of example, for certain parameters, such as SatO2, a very sharp fall from a normal value to a null or very low one is not possible. The values of others may not exceed a given rate of change; e.g. certain increases in the heart rate are impossible due to them being sharper than the heart's response capacity.

On the other hand, the literature on Medical Informatics highlights the interest of medical personnel in alarms capable of identifying findings over the temporal evolution of physiological parameters, such as, for example, "sustained increase in the HR of at least 15 beats over approximately half a minute" [13]. Among the solutions appearing in the bibliography which tackle this problem, those that indicate the compatibility of the evolution of a physiological variable with a fuzzy trajectory are considered to be of special interest [7], [11], [12].

C. Integrating information from different parameters

The capacity for integrating information originating from various parameters allows alarms to be generated on the basis of findings which, taken in isolation, are irrelevant, but which may endorse the hypothesis of the occurrence of pathological processes of clinical interest if they can be related with other findings over other variables, which on their own may also be irrelevant. This also makes it possible to reduce the margins of the abnormality values monitored over each physiological variable, keeping the number of false positives within reasonable levels thanks to the merging of information originating from more than one parameter. Furthermore, these alarms supply strong diagnostic evidence on the pathology that they monitor, given the large quantity of information that they contain.

Our experience in the medical domain has shown us that physicians, especially the most experienced, use monitoring criteria based on multiparametric patterns, allowing pathologies to be identified in the early stages, while still not life-threatening, and their effect on the physiological variables is too subtle to be identified by means of threshold of alarms. The lack of devices capable of monitoring these patterns means that the burden of this task falls squarely on the shoulders of health-care staff. Nevertheless, the overwhelming nature of this task means that monitoring is only carried out in those scenarios were they are more likely to arise, and not in all those where it would be desirable.

One example of this type of scenario is hemodialysis, a therapy requiring the introduction of a catheter, via the femoral route, into the patient's body. When the patient's blood starts flowing through the tubes and filters of the hemodialysis machine, the reduction of blood in the body may result in hypovolemia. In its early stages this pathology is evident in the parameters that are commonly monitored in the Medical Unit in the form of a slight but sustained increase in heart rate that is simultaneous with a light, sustained increase in the systolic blood pressure. This and other similar patterns are constantly monitored by the physician in charge of the hemodialysis, especially during the initial minutes.

On the other hand, reasoning simultaneously over the evolution of more than one parameter makes it possible to identify artifacts when the behaviour of one or more of them is not consistent with the rest. For example, a null or extremely low value for systolic blood pressure is an artifact if the mean and diastolic blood pressure values are normal, or if the heart continues to beat at a normal rate.

D. Editing of the monitoring criteria

The physiological variability among human beings prevents the definition of generic monitoring criteria that can be applied to any patient. To complicate the situation even further, the physical state of patients in the medical unit differs radically from that of healthy patients, and often they are under the effects of drugs. Thus, any proposal attempting to solve the problem of the cognitive overload must take into consideration the monitoring context associated with each patient.

Proposals in the literature habitually tackle this problem by incorporating contextual information, which modifies the monitoring criteria [14]. This was the path that we followed for many years in our monitoring techniques. With the passing of time, we came to realize that there are no monitoring criteria that are unique to a given pathology: two physicians may use different threshold criteria to monitor the same patient without it being possible to assert that the criteria of one is more suitable than that of the other. It is simply the case that each physician wishes to be alerted on different deviations from normality.

Our research has led us to resolve the problem of context in our proposal by supplying the physician with certain "template"monitoring patterns which the physician understands and may edit in a simple way in order to represent the criteria that he/she may consider opportune. The solution we propose does not differ greatly from the current situation: when patients are admitted to a medical unit, the first thing medical personnel do is adjust the ranges of the threshold alarms on the basis of their state and the monitoring objectives.

Enabling the physician to edit the monitoring criteria of an alarm imposes certain constraints on the monitoring techniques that can be used. On one hand, there must be a computable representation of the monitoring criteria that is understandable for medical personnel, and which they must be able to edit in a simple manner. On the other hand, it must be possible to generate detection procedures for certain criteria without the need to carry out any implementation task, and without a learning process. Thus, the technique employed must separate the representation of monitoring criteria from the matching procedures that make it possible to identify them.

Thus, structural pattern recognition techniques [15] would seem to be the most suitable for this new generation of alarms. Unlike statistical or connectionist ones, these techniques work directly on the input data, and not on a properties space that is created on the basis of them. Hence, in this technique the representation of the monitoring criteria is close to a physician's mental representation of the criteria. On the other hand, in structural techniques, input data are compared with a set of primitives, which combine with each other through a series of relations to give rise to more complex patterns. This comparison is carried out by segmenting the input data, which allows explanations to be given on the matching results and means that there is no need for training: the matching algorithms only require a definition of of the primitives, and this is obtained directly from the medical staff.

III. Representation of Monitoring Criteria

We shall start by introducing certain prior concepts on which the MFTP model is based.

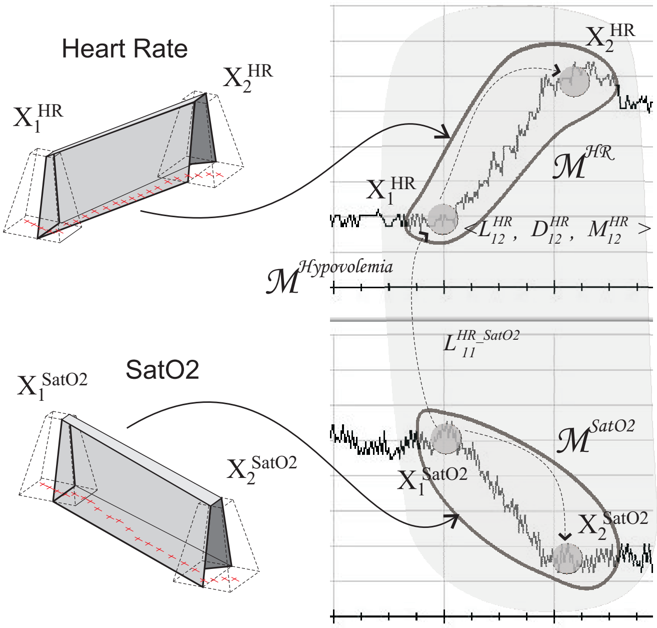

Fig. 1. Graph of the hypovolemia MFTP drawn over a real occurrence of the pattern.

A. *Fuzzy fundamentals*

We shall consider time as being projected onto a one-dimensional discrete axis $\tau = \{t_0,\ t_1,\ ...,\ t_i,\ ...\}$, where t_i represents a *precise* instant. We consider that for every $i \in \mathbb{N}$, $t_{i+1} - t_i = \Delta t$, where the constant Δt is the minimum step of the temporal axis.

Given as discourse universe the set of real numbers $\mathbb{R}$, a *fuzzy number* A is a normal ($\exists\ v \in \mathbb{R},\ \mu^A(v) = 1$) and convex ($\forall v, v', v'' \in \mathbb{R},\ v' \in [v, v''],\ \mu^A(v') \geq \min\ \{\mu^A(v), \mu^A(v'')\}$) fuzzy subset of $\mathbb{R}$. We obtain a fuzzy number A from a flexible constraint given by a possibility distribution π^A, which defines a mapping from $\mathbb{R}$ to the real interval $[0, 1]$. Given a precise number $v \in \mathbb{R}$, $\pi^A(v) \in [0, 1]$ represents the possibility of A being precisely v. By means of π^A we define a fuzzy subset A of $\mathbb{R}$, which contains the possible values of A.

B. *The MFTP model*

The MFTP model [7] makes it possible to represent monitoring criteria, comprising a set of morphologies defined over the evolution of a set of physiological variables $\mathcal{P} = \{P^1, ...,\ P^s\}$, where each variable P^p is obtained by means of an acquisition and sampling process: $P^p = \{(v_1^p, t_1^p), ..., (v_i^p, t_i^p), ...\}$, and temporal and magnitude relations between these morphologies. The MFTP model is based on the CSP formalism and on fuzzy set theory. An MFTP allows a pattern to be represented by means of a network of fuzzy constraints between a set of significant points.

We define *significant point* on a physiological parameter P^p, X_i^p, as the pair formed by a variable from the domain V_i^p and a temporal variable T_i^p. A significant point $X_i^p = < V_i^p, T_i^p >$

represents an unknown value for P^p at an unknown temporal instant. In the absence of constraints, V_i^p and T_i^p may take any precise value v_i^p and t_i^p, respectively, where $(v_i^p, t_i^p) \in P^p$. By A_i^p we denote the assignment of precise values from the evolution P^p to the variables of X_i^p; i.e., $A_i^p = (v_i^p, t_i^p)$.

In principle, nothing restricts the form of the constraints that make up a MFTP. However, experience has shown that a set of constraints limiting the fuzzy increment, fuzzy temporal extension and fuzzy slope between a pair of significant points is able to capture the majority of monitoring criteria habitually used by medical staff. Thus we define a *constraint* L_{ij}^{pq} between two significant points X_i^p and X_j^q by means of a normal, convex possibility distribution $\mu^{L_{ij}^{pq}}(X_i^p, X_j^q) = \pi^{L_{ij}^{pq}}(h)$, $h \in \tau$, which represents the possibility of the *fuzzy temporal extension* between X_i^p and X_j^q being h. The assignations $T_i^p = t_i^p$ and $T_j^q = t_j^q$ are possible if $\pi^{L_{ij}^{pq}}(t_j^q - t_i^p) > 0$. In Fig. 1 L_{12}^{HR} models the linguistic description "*approximately more than half a minute*".

A *constraint* D_{ij}^{pq} between a pair of significant points X_i^p and X_j^q is defined by means of a normal and convex possibility distribution $\mu^{D_{ij}^{pq}}(X_i^p, X_j^q) = \pi^{D_{ij}^{pq}}(d)$, $d \in \mathbb{R}$, which represents the possibility of the *fuzzy increase* between X_i^p and X_j^q being d. The assignations $V_i^p = v_i^p$ and $V_j^q = v_j^q$ are possible if $\pi^{D_{ij}^{pq}}(v_j^q - v_i^p) > 0$. In Fig. 1 D_{12}^{HR} models the description "*increase of more than approximately 15*".

A *constraint* M_{ij}^p between a pair of significant points X_i^p and X_j^p, defined over the same parameter P^p, is defined by means of a normal and convex possibility distribution $\mu^{M_{ij}^p}(X_i^p, X_j^p) = \pi^{M_{ij}^p}(m)$, $m \in \mathbb{R}$, which represents the possibility of the *fuzzy slope* between X_i^p and X_j^p being m. The assignations $V_i^p = v_i^p$, $V_j^p = v_j^p$, $T_i^p = t_i^p$ and $T_j^p = t_j^p$ are possible if $\pi^{M_{ij}^p}((v_j^p - v_i^p)/(t_j^p - t_i^p)) > 0$. In Fig. 1 M_{12}^{HR} models the description "*...sustained...*", where "*sustained*" is modelled by means of an approximately constant slope value.

We define a *Multivariable Fuzzy Temporal Profile* (MFTP) $\mathcal{M} =< \mathcal{W}^{\mathcal{M}}, \mathcal{X}^{\mathcal{M}}, \mathcal{R}^{\mathcal{M}} >$ as a finite set of MFTPs $\mathcal{W}^{\mathcal{M}} = \{\mathcal{M}_1^{\mathcal{M}}, ..., \mathcal{M}_s^{\mathcal{M}}\}$, a finite set of significant points $\mathcal{X}^{\mathcal{M}} = \{X_{i_1}^{p_1}, X_{i_2}^{p_2}, ..., X_{i_g}^{p_g}\}$ and a finite set of constraints $\mathcal{R}^{\mathcal{M}} = \{R_1, ..., R_f\}$ amongst the points of $\mathcal{W}^{\mathcal{M}}$ and $\mathcal{X}^{\mathcal{M}}$.

The recursive structure of the MFTP model is based in the way that humans define patterns; a complex pattern is often made up of a set of findings and a set of relations between them. Each of the findings of the pattern may also be a pattern, and may comprise a set of findings and relations between them, and so on, successively. Thus, for example, the hypovolemia pattern (see Section II C) is made up of two findings for which certain temporal relations between them must be satisfied: $\mathcal{M}^{Hypovolemia} =< \{\mathcal{M}_{HR}^{Hypovolemia}, \mathcal{M}_{BP}^{Hypovolemia}\}, \varnothing, \{L_{11}^{HR\ BP}\} >$. Each of the findings is in turn an MFTP that is defined over its corresponding parameter; thus $\mathcal{M}_{HR}^{Hypovolemia} =< \varnothing, \{X_1^{HR}, X_2^{HR}\}, \{L_{12}^{HR}, D_{12}^{HR}, M_{12}^{HR}\} >$, and $\mathcal{M}_{BP}^{Hypovolemia} =< \varnothing, \{X_1^{BP}, X_2^{BP}\}, \{L_{12}^{BP}, D_{12}^{BP}, M_{12}^{BP}\} >$.

An MFTP can be represented by a graph in which nodes correspond to significant points, and arcs correspond to constraints (see Fig. 1). The MFTP model also enables us to restrict the evolution of a parameter P^p between each pair of significant points X_i^p and X_j^p (see Fig. 1) by means of a membership function $\mu^{\mathcal{S}_{ij}^p}(\mathcal{A}_i^p, \mathcal{A}_j^p)$ which defines a fuzzy course within which the temporal evolution of the parameter must remain in order to satisfy the constraint [16].

C. Pattern recognition procedures

The MFTP definition allows the matching task to be structured hierarchically, where a pattern $\mathcal{M}$ constitutes a processing level that incorporates a set of findings detected in the previous processing level. Identifying a pattern $\mathcal{M}$ over the evolution $P = \{P^1, ..., P^s\}$ of the patient's physiological variables is equivalent to finding a solution to the fuzzy constraint network defined by $\mathcal{M}$ [17]. A network solution is built by means of the assignment A_i^p of a sample of the evolution of P^p to each significant point X_i^p. A solution of $\mathcal{M}$ is defined as a set of assignments $A = \{A_0, A_1, ..., A_n\}$ that satisfy the set of constraints that make up $\mathcal{M}$, with a degree higher than zero. The degree of satisfaction of A is given by the minimum degree of satisfaction of all the constraints of $\mathcal{M}$, that is:

$$\pi^{\mathcal{M}}(\mathcal{A}) = min\{ \min_{\mathcal{M}_h^{\mathcal{M}} \in W^{\mathcal{M}}} \{\pi^{\mathcal{M}_h^{\mathcal{M}}}(A^{\mathcal{M}_h^{\mathcal{M}}})\}, \min_{\mathcal{R}_k \in \mathcal{R}^{\mathcal{M}}} \{\pi^{\mathcal{R}_k}(A^{\mathcal{R}_k})\}\}$$

where $A^{\mathcal{M}_h^{\mathcal{M}}}$ is the projection of A over the set of significant points involved in $\mathcal{M}_h^{\mathcal{M}}$, and $A^{\mathcal{R}_k}$ is the projection of A over the set of significant points that are involved in R_k. $\pi^{\mathcal{R}_k}$ is the degree of satisfaction of $R_k \in R^{\mathcal{M}}$, and $\pi^{\mathcal{M}_h^{\mathcal{M}}}$ is the degree of satisfaction of $\mathcal{M}_h^{\mathcal{M}} \in W^{\mathcal{M}}$. The solutions to each $\mathcal{M}_h^{\mathcal{M}}$, and hence $\pi^{\mathcal{M}_h^{\mathcal{M}}}$, are calculated in a previous recognition stage and then assembled to find a solution for $\mathcal{M}$. $\pi^{\mathcal{M}}(A)$ represents the compatibility between a fragment of the evolution of patient's physiological variables with the description made in $\mathcal{M}$ of a given association between findings of physiological interest.

For example, in order to match the hypovolemia pattern, we start by searching for occurrences of the two findings that make it up. In order to calculate the degree of compatibility of $\mathcal{M}^{HR}$ with $\mathcal{A}^{\mathcal{M}^{HR}} = \{A_1^{HR}, A_2^{HR}\}$ the following expression is used:

$$\pi^{\mathcal{M}^{HR}}(\mathcal{A}^{\mathcal{M}^{HR}}) = min\{\pi^{L_{12}^{HR}}(t_2^{HR} - t_1^{HR}),$$
$$\pi^{D_{12}^{HR}}(v_2^{HR} - v_1^{HR}), \pi^{M_{12}^{HR}}((v_2^{HR} - v_1^{HR})/(t_2^{HR} - t_1^{HR}))\}$$

where the assignations A_1^{HR} and A_2^{HR} are taken from the values registered for the heart rate. A similar expression applies for $\mathcal{M}^{BP}$. Solutions are then sought searched for $\mathcal{M}^{Hypovolemia}$ over the previously found occurrences for $\mathcal{M}^{HR}$ and $\mathcal{M}^{BP}$. In order to calculate the degree of membership of $\mathcal{A}^{Hypovolemia}$ the following expression is used:

$$\pi^{\mathcal{M}^{Hypovolemia}}(\mathcal{A}^{Hypovolemia}) = min\{\pi^{\mathcal{M}^{HR}}(\mathcal{A}^{HR}),$$
$$\pi^{\mathcal{M}^{BP}}(\mathcal{A}^{BP}), \pi^{L_{11}^{HR\ BP}}(t_1^{BP} - t_1^{HR})\}$$

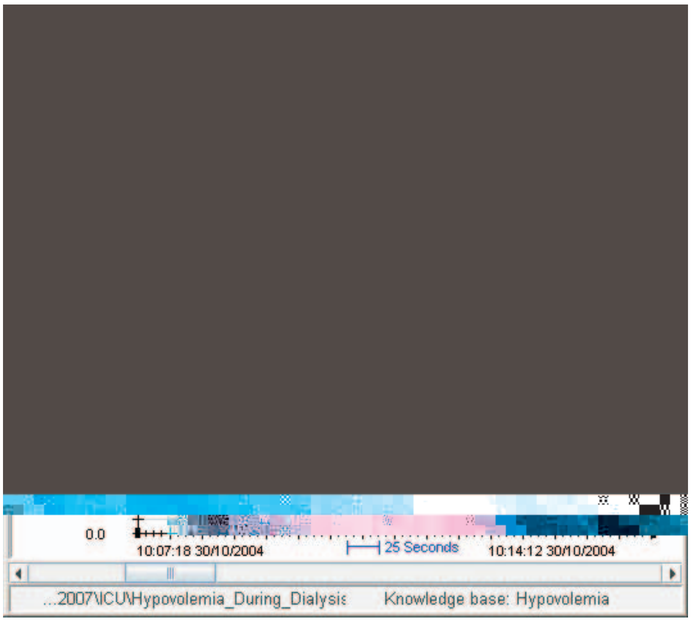

Fig. 2. TRACE SHOWING THE PATTERN DETECTION FOR HYPOVOLEMIA OVER A RECORDING TAKEN DURING DIALYSIS.

where $\mathcal{A}^{HR}$ and $\mathcal{A}^{BP} \subset \mathcal{A}^{Hypovolemia}$.

Despite the theoretically high computational complexity of matching an MFTP, the modular breakdown of the problem along with a small number of heuristics that exploit the continuity properties of real signals allow the real-time requirements of the medical domain to be fully satisfied [7].

D. Knowledge acquisition: TRACE

Based on the MFTP model we have constructed the Tool foR anAlyzing and disCovering pattErns, TRACE [8], [9], a tool for creating, editing and validating alarms based on the MFTP model. The tool makes use of the very graph that represents the MFTP (the form of which is reminiscent of that of the monitoring pattern it represents) as a visual metaphor to assist in editing knowledge relating to alarms. This editing can be carried out in an entirely visual manner, using only the mouse.

The tool also allows the matching procedures for the MFTP model to be executed, and its results viewed. Each detection colours the fragments of each physiological variable that has demonstrated compatibility with the morphology defined over it, and adds a signal to the environment called detection, which represents the compatibility of the global pattern. TRACE has proved to be sufficiently intuitive for use by medical teams without the need for assistance [8].

IV. EXPERIMENTAL RESULTS

Through the use of TRACE, and with the support of a medical team, we have defined a set of alarms that check monitoring criteria that cannot be suitably represented by means of thresholds. Alarms do not necessarily correspond with a manifestation of a pathology, although they all identify occurrences of events that are of interest to the physician.

TABLE I.
RESULTS OF THE DETECTION.

Alarm	C	FP	FN	%C	%FP	%FN
LV_SatO2	35	2	3	95	5.7	7.9
LV_HR	35	2	0	95	5.7	0
LV_BP	32	1	3	97	3.1	8.6
HV_HR	62	1	0	98	1.6	0
I_BP	74	1	0	99	1.4	0
I_RR	48	4	1	92	8.3	2.0
D_RR	19	2	0	90	10	0
I_HR-D_BP	31	1	0	97	3.2	0
I_HR-D_SatO2	20	2	0	91	10	0
I_HR-I_SatO2	32	3	2	91	9.4	5.9
I_HR-I_BP	64	3	3	96	4.7	4.5
I_RR-D_SatO2	13	0	0	100	0	0
Total	465	22	12	97	4.7	2.5

C: Correct; FP: False Positive; FN: False Negative.

A number of these alarms supervise the occurrence of episodes, of at least four minutes, with moderately abnormal values in a physiological variable: an abnormally low (1) blood oxygen level (LV_SatO2), (2) heart rate (LV_HR) and (3) blood pressure (LV_BP) values; and an abnormally high heart rate value (HV_HR). Others supervise trends in a parameter; i.e. sustained moderate increases/decreases for at least 45 seconds: rise in (1) blood pressure (I_BP) and (2) respiratory rate (I_RR); and a fall in the respiratory rate (D_RR).

Some alarms incorporate information originating from two parameters: a rise in the heart rate which is approximately simultaneous with (1) a drop in blood pressure (I_HR-D_BP), (2) a drop in the blood oxygen level (I_HR-D_SatO2), (3) an increase in the blood oxygen level (I_HR-I_SatO2) and (4) and increase in blood pressure (I_HR-I_BP); and an increase in the respiratory rate which is approximately simultaneous with a decrease in the blood oxygen level (I_RR-D_SatO2). In this case there is a reduction in the abnormality values that are supervised over each parameter with respect to alarms that only consider information from one parameter.

Approximately 175 hours of recordings of physiological parameters from 71 different patients admitted to the Intensive Care Unit were used to validate the alarms. Recordings varied in length (from barely 20 minutes to over 12 hours) and not all of them contained the same parameters. TRACE was used to run the matching procedures, and the results of the validation are given in Table I. The generation of an alarm in a context which physicians "would not have wished to be alerted" due to their considering a more detailed examination of the patient unnecessary is taken as a false positive. A false negative is taken to be a situation in which the physician would have carried out a more detailed examination and it was not reported as an alarm. A correct detection is taken to be the generation of an alarm in a context in which physicians would have wished to be alerted due to their considering a more detailed examination of the patient necessary.

The manner in which the tests were carried out, defining certain unique monitoring criteria and employing them on 71 different patients, differs from the way in which MFTP-based alarms will most likely be used in the medical domain routine. Ideally, before commencing with monitoring, physicians should use TRACE to define the criteria they wish to supervise for each patient on the basis of a template MFTP. Tests were carried out in this manner as, with the exception of the signal recording, no information with which to contextualize the monitoring criteria was available on the vast majority of the 71 patients.

The number of false positives generated was excellent compared with those appearing with threshold alarms. The false negatives were due to exceptionally sharp rates of change in the values of patients' physiological variables. Our monitoring criteria avoid reporting certain rates of change as alarms by considering that they are too high, and must must thus be artifacts. Although this was crucial in obtaining a low number of false positives, in the medical domain a false negative results in disregarding an alert of a possible life-threatening situation for the patient, with potentially disastrous results. By relaxing the rates of change admissible for the parameters, the false negatives could be eliminated, even though the number of false positives would probably increase; nevertheless, these would foreseeably remain within an acceptable range, and would be much lower than for threshold alarms.

V. Conclusions

In this work we have presented a proposal for constructing alarms that mitigate the cognitive overload on medical staff arising due to the large quantities of physiological variables monitored over each patient and the rudimentary nature of the alarms available in commercial monitoring devices. These alarms (1) show the different degrees between the maximum certainty of the presence of manifestations of pathologies that they monitor and the maximum certainty of their absence; (2) they allow reasoning over the temporal evolution of the physiological variables; (3) they can integrate information originating from more than one variable into a single parameter; and, in spite of their high semantic content, (4) they can be edited by physicians in a simple manner without the need for assistance thanks to the use of visual metaphors in their definition.

Currently, further research still needs to be carried out to ensure that these alarms permit the supervision of all situations of abnormality that can be identified with threshold alarms. Nevertheless, the results obtained to date are highly promising: the alarms proposed herein show a significant reduction in the production of false positives; the diagnostic evidence they supply on the pathologies they supervise is significantly superior than that supplied by threshold alarms; and they are capable of identifying situations of abnormality that cannot be supervised using thresholds. Thus we believe that the coexistence of the alarms proposed in the present study with those currently in use in medical units would result in a significant increase in the quality of health care.

One problem that we have encountered in this work has been the lack of a catalogue of monitoring criteria based on patterns of the temporal evolution of physiological variables or in the combination of multiparametric information - in spite of the fact that health care staff habitually employ these types of criteria. The most likely reason for this is the lack of tools for automatically identifying these criteria, making it difficult to carry out rigorous studies on them.

Our future work will be aimed at continuing to evaluate the alarms dealt with herein, along with any others that may be of interest for health care staff, and with them to construct a monitoring criteria catalogue. In a pilot experiment, this catalogue will be applied to the routine of an ICU, employing an intelligent patient supervision system in which alarms will be configured using TRACE's pattern editor.

Acknowledgments

This work was supported by the Spanish MEC and the European FEDER, and by the Xunta de Galicia under the grants TIN2006-15460-C04-02 and PGIDIT04SIN206003PR, respectively.

References

[1] R. Bellazzi, C. Larizza, and A. Riva, "Temporal abstractions for interpreting diabetic patients monitoring data," *Intelligent Data Analisis*, vol. 2, pp. 97–122, 1998.

[2] E.E. Milios and S.H. Nawab, "Signal abstraction in signal processing software," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 37, pp. 157–180, 1989.

[3] V. Moret-Bonillo and A. Alonso-Betanzos, "Uncertainty based approach for symbolic classification of numeric variables in intensive care units," *Journal of Clinical Engineering*, vol. 15, no. 5, pp. 361–370, 1990.

[4] Y. Shahar and M.A. Musen, "Knowledge-based temporal abstraction in clinical domains," *Artificial Intelligence in Medicine*, vol. 8, pp. 267–298, 1996.

[5] A. Jungk, B. Thull, and G. Rau, *Fuzzy Logic in Medicine*, chapter Intelligen Alarms for Anaesthesia Monitoring Based on Fuzzy Logic Approach, Physica-Verlag, 2002.

[6] F.A. Mora, G. Passariello, G. Carraylt, and J.L. Pichon, "Intelligent patien monitoring and management systems: A review," *IEEE Engineering in Medicine and Biology*, 1993.

[7] A. Otero, P. Félix, S. Fraga, S. Barro, and F. Palacios, "Lecture notes in artificial intelligence," 2006, vol. 4177, pp. 31–41.

[8] A. Otero, P. Félix, S. Barro, and F. Palacios, "A tool for the analysis and synthesis of alarms in patient monitoring," in *The 7th International Conference on Information Fusion*, 2004, pp. 951–958.

[9] A. Otero. General overview of trace. https://www.dec.usc.es/trace/, 2006. Flash presentation.

[10] F. Steimann, "Fuzzy set theory in medicine," *Artificial Intelligence in Medicine*, vol. 11, pp. 1–7, 1997.

[11] A. Lowe, M.J. Harrison, and R.W. Jones, "Diagnostic monitoring in anaesthesia using fuzzy trend templates for matching temporal patterns," *Artificial Intelligence in Medicine*, vol. 16, pp. 183–199, 1999.

[12] F. Steimann, "The interpretation of time-varying data with DIAMON-1," *Artificial Intelligence in Medicine*, vol. 8, pp. 333–357, 1996.

[13] E.T. Keravnou, "Temporal reasoning in medicine," *Artificial Intelligence in Medicine*, vol. 8, pp. 187–191, 1996.

[14] Y. Shahar, "A framework for knowledge-based temporal abstraction.," *Artificial Intelligence*, vol. 90, pp. 79–133, 1997.

[15] S.N. Srihari and V. Govindaraju, *Encyclopedia of Computer Science*, chapter Pattern Recognition, Van Nostrand-Reinhold, IEEE Press, 1993.

[16] P. Félix, S. Barro, and R. Marín, "Fuzzy constraint networks for signal pattern recognition," *Artificial Intelligence*, vol. 148, pp. 103–140, 2003.

[17] R. Dechter, *Constraint Processing*, Morgan Kaufmann Publishers, 2003.