

Combining diagnosis and treatment using ASBRU

Andreas Seyfang^{a,*}, Silvia Miksch^a, Mar Marcos^{b,c}

^a *Institute of Software Technology and Interactive Systems, University of Technology, Favoritenstraße 9-11/188, A-1040 Vienna, Austria*

^b *Department of AI, Faculty of Sciences, Vrije Universiteit, Amsterdam, The Netherlands*

^c *Department of Computer Engineering and Science, Universitat Jaume I, Castelló, Spain*

Abstract

Traditionally, diagnosis and treatment have been seen as two distinct tasks. Consequently, most approaches to computer supported health care focus on one of the two—mostly on diagnosis or rather on the interpretation of measurements which is much better understood and formalised. However, in practice diagnosis and treatment overlap and influence each other in many ways. Combinations range from repeatedly going through the diagnosis-treatment loop over a period of time to permanent monitoring of the patients' health condition as it is done in intensive care units. In this article we describe how to model these combinations using the clinical protocol-representation language ASBRU. It implements treatment steps in a hierarchy of skeletal, time-oriented plans. Diagnosis can either be described in a declarative way in the conditions, under which treatment steps are taken or it can be modelled explicitly as plans of their own right. We demonstrate our approach using examples taken from the American Association of Paediatricians' guideline for the treatment of hyperbilirubinemia in the new-born.

© 2002 Elsevier Science Ireland Ltd. All rights reserved.

Keywords: Guideline-; Protocol-based care; Diagnosis; Treatment management; ASGAARD; Hyperbilirubinemia; ASBRU

1. Introduction

Clinical guidelines and protocols are gaining increasing acceptance in medical practice as a means to support both diagnosis and treatment. While it is clear that no computer system can replace the physician and her complex reasoning about the patient's state, the positive effect of clinical guidelines on medical practice has been demonstrated by several studies [1].

There are various representations for modelling guidelines which are compared with our approach in the [Section 7](#). In this paper we describe the implementation of scenarios of diagnosis and treatment in our skeletal plan-representation language ASBRU based on a guideline for the treatment of jaundice in neonates.

While diagnosis has a much broader meaning in medicine and in the Artificial Intelligence technique Model-Based Diagnoses, in this paper we refer to diagnosis as the application of a set of rules on information about the patient in order to derive abstrac-

* Corresponding author

E-mail address: seyfang@asgaard.tuwien.ac.at (A. Seyfang).

tions of the data which are closer to the complex medical concepts. This abstractions may but need not always match the diagnosis of, e.g. a certain disease or syndrome.

2. Scenarios of diagnosis and therapy

Going from the simplest case to the most complex, we can distinguish the following settings.

2.1. *Diagnosis considered alone*

In the simplest case, only the diagnostic task is considered, without explicitly mentioning treatment options. This class of guidelines is the oldest and most widely spread. The exact parameters for a certain diagnosis have been explored in many fields and well-known standards are established. The output of the diagnosis process is only the name of a disease, a syndrome, or a symptom—there is no suggestion for treatment.

2.2. *Initial diagnosis with treatment recommendations*

This is the case, if the above case is augmented with recommendations for treatment. Most often only the initial diagnosis is covered, later re-examination of the patient is assumed but not mentioned.

2.3. *Repeated loops of diagnosis and treatment*

In this case the re-examination of the patient is explicitly denoted by giving the interval of time between two examinations, the conditions for reassessing the patient's health state, or the task after whose completion another diagnostic pass has to be performed.

2.4. *Continual diagnosis during treatment*

Diseases which either last for an extended period of time or which threaten the patient's life or health mandate an even more extensive observation than the scenarios above. For technical reasons one can distinguish high-frequency domains such as intensive care units and low-frequency domains such as diabetes. While in the first new values are obtained e.g. every second, in domains of the second category values are obtained several times per day or week. Concerning the topic of this paper we can consider both domains together ignoring technical differences.

3. Modelling diagnosis and treatment

3.1. *Modelling diagnosis*

There are two opposite ways of modelling a set of questions leading to a set of actions or diagnostic findings: decision trees and decision tables. Decision trees are most suitable for situations, in which the question to be asked secondly depends on the answer to the first question. For example, if the question for the sex of the patient is answered with 'female', the second question might deal with breast cancer while for men the second question might ask for prostate gland cancer. There is no point in inquiring breast cancer in men or prostate gland cancer in women.

In other cases, questions will be independent, e.g. queries for a set of risk factors. These are best modelled using a decision table, which assigns a diagnostic finding or a similar statement to each of all the possible combinations of inputs.

It is clear to see that modelling a certain combination of questions using the wrong method leads to tiresome redundancy. However, practical problems always bring both

types of question sets—such to model with decision trees and such to model with decision tables. The solution lies in combining the questions in groups each of which is either modelled by a tree or a table and by using the findings of one group as an input to other groups.

3.2. *Modelling treatment*

Although many guidelines available today focus on describing which treatment should be considered and leave the details of its pursuit open, computer support during the treatment is desired in many fields of medicine. A system satisfying this demand must model the characteristic features of treatment: modular decomposition into interchangeable subtasks, flexible temporal properties of both the actions themselves and their intentions and effects, and complex conditions under which a task is started, suspended, resumed, successfully completed, or aborted.

3.3. *Interleaving diagnosis and treatment*

In some cases, only the failure or success of a certain treatment is a solid basement for diagnosis. For example a favourable response to phototherapy is a necessary indication for its appropriateness. This means that a treatment step can form an integral part of the diagnosis process. Often, the tasks of diagnosis and treatment are seen as distinct and only loosely coupled in guidelines and textbooks, while in clinical practice they cannot be separated.

4. The need for data abstraction

Some of the information needed in the diagnosis process is entered in the same form it is needed, e.g. by answering questions with

yes or no. Other parts, such as complex series of measurements must be filtered to eliminate measuring errors, compared with each other to detect trends, and abstracted to turn a host of data into information useful to the physician. Not only high-frequency domains such as monitoring in an ICU, but also low frequency domains in which parameters are measured a few times per day often demand complex abstractions from the numbers delivered by the laboratory, namely their association with one of several qualitative values such as ‘normal’, ‘too high’, or ‘too low’. This mapping is often context dependent, i.e. the same number can represent a normal value in one situation and a pathologically increased value in another, which calls for context-sensitive data abstraction.

5. Materials and methods

In this section, we describe the plan-representation language ASBRU used for modelling, as well as the medical background of the guideline presented in the next section.

5.1. *Introduction to ASBRU and ASGAARD*

ASBRU is a skeletal plan-representation language for the modelling of time-oriented, hierarchical treatment schemata. Diagnosis is modelled as conditions under which treatment steps are performed. ASBRU is part of the ASGAARD project in which a set of task-specific problem-solving methods for the design, execution, and critiquing of treatment plans are developed [3].

An ASBRU plan library is built on skeletal plans which are plan schemata at various levels of detail, capturing the essence of the procedure, but leaving enough room for execution-time flexibility in the achievement of particular goals [4]. Thus, they are usually

reusable in various contexts. In ASBRU, we have enriched the idea of skeletal plans by adding knowledge roles, a rich set of ordering of actions and plans, and temporal dimension of states, actions, and plans. ASBRU enables the designer to represent a clinical protocol in computer-processable form. It also enables the physicians to execute the protocol in a flexible way, based on the patient's state. This leads to a better acceptance of both protocols and computer support by the medical staff.

An ASBRU plan is identified by its name and consists of five components: *preferences*, *intentions*, *conditions*, *effects*, and a *plan body* (*layout*) which describes the actions to be executed. The major features of ASBRU are that:

- prescribed actions and states can be continuous;
- intentions, conditions, and world states are temporal patterns;
- uncertainty in both temporal scopes and parameters can be flexibly expressed by bounding intervals;
- a rich set of ordering of plans (all or some plans might be executed in sequence, in parallel, unordered, or periodically);
- particular conditions and operators are defined to control the plans' execution; and
- explicit intentions and preferences can be stated for each plan separately.

The basic syntactic construct is the temporal pattern. All conditions for the transition from one plan state to another are expressed in terms of temporal patterns. They consist of one or more parameter propositions or plan-state descriptions. Each parameter proposition contains a parameter name, a value description, a context description and a time annotation. The time annotations used allow the representation of uncertainty in starting time, ending time, and duration. The time

annotation supports multiple time lines (e.g. different zero-time points and time units) by providing arbitrary *reference annotations*. Temporal shifts from the reference annotation are used to define the uncertainty in starting time, ending time, and duration. To allow temporal repetitions, sets of cyclical time points and cyclical time annotations can be defined.

Plans in ASBRU have a rich set of plan states. After passing the *filter* condition and (optionally) the confirmation through the user (a physician), they are *active*. Then they can either be temporarily *suspended* (under the *suspend* condition) until they are *resumed* (under the *restart* condition). If successful (determined by the *complete* condition), they are *completed*. If circumstances occur which cause the particular action to fail (as stated in the *abort* condition), the plan is *aborted*.

5.2. Medical background

We used ASBRU to model a practice guideline for the 'Management of Hyperbilirubinaemia in the Healthy Term New-born'. The version we have worked with is the guideline developed by the American Association of Paediatrics (AAP) [2]. The guideline, although intended for application in a hospital setting, is addressed to health-care professionals who are not specialist in diseases of new-born babies.

Jaundice (or hyperbilirubinemia) is a common disease in new-born babies. Under certain circumstances, elevated bilirubin levels may have detrimental neurological effects. In many cases jaundice disappears without treatment but sometimes phototherapy is needed to lower the level of total serum bilirubin (TSB), which indicates the presence and severity of jaundice. In few cases it is a sign of a more serious disease. The jaundice protocol of the AAP is intended for the

management of jaundice in healthy term newborns.

The guideline consists of an evaluation (or diagnosis) part and a treatment part. The evaluation part is charged of uncovering the possibility of a more serious disease. If this happens, the application of the guideline is interrupted without any further action. Otherwise the treatment part decides the appropriate therapy according to the level of TSB.

We chose this guideline as example because—unlike other guidelines—it contains a lot of detailed information about the treatment in it, as well as several interactions between diagnosis and treatment.

In this paper we focus on the following parts of the guideline:

- The eligibility for the guideline¹ is decided before beginning the treatment but it may be revised during treatment if contraindications occur.
- The selection of the appropriate treatment is performed whenever new information about the health state of the patient is entered.
- The modes of treatment itself are started by the triggering of the filter condition of one of the treatment options and stopped by either the complete condition or the abort condition of the plan representing the treatment step or the plans containing it.

The examples in this paper are shown in a table notation which directly reflects the XML²-elements constituting ASBRU. The latter

would be by far too space consuming to be reproduced in this paper³.

6. Results

In this section we give a few examples from the guideline described above to illustrate the modelling of diagnosis and treatment in ASBRU. We focus on the eligibility criteria, on the selection of the appropriate treatment, and on an example of treatment step.

6.1. Modelling the eligibility criteria for the guideline

Since the guideline only covers healthy term jaundiced neonates, any sign of additional diseases or risks precludes its application. Therefore, the eligibility criteria for the guideline are constituted by the negation of a set of exclusion criteria. We model them in ASBRU as different sets of questions posed to the physician. If a single one among them is answered with ‘yes’, there is the possibility of some pathologic reason for jaundice, which mandates further examination beyond the scope of the guideline, which therefore must be aborted.

This is implemented by combining the answers to the questions using the Boolean ‘or’ operator, yielding a Boolean parameter. Examples of parameters defined in this way are possibility-of-other-diseases and possibility-of-cholestatic-disease. A tabular representation of the ASBRU code for one of such parameters is given in Fig. 1. In addition, there is a more general parameter named pathologic-reason, which depends on the previous ones, amongst others. This parameter becomes true if any of the possibility-of-... ones is ‘yes’. The parameter pathologic-reason is compared in the filter condition of the overall-treatment plan (see Fig. 2) with the

¹ Note that while physicians talk about the eligibility of a patient for a certain guideline, in this paper we use the term eligibility in relation to the guideline rather than the patient [5].

² The eXtensible Markup Language is the upcoming international standard for structured information.

³ The full XML (ASBRU) version of the guideline can be found at <http://www.protocol.org/>.

parameter-group		
parameter-def	name	= possibility-of-other-diseases
		type = Boolean
	boolean-operation	type = OR
		parameter-ref name = lethargy
		parameter-ref name = apnoea
parameter-def	name	= lethargy
		type = Boolean
	raw-data-def	mode = manual
...		

Fig. 1. Definition of an abstraction used in the eligibility criteria. ‘Possibility-of-other-diseases’ holds if ‘lethargy’ or ‘apnoea’ or... are present. Each of these symptoms is a manually entered Boolean (yes or no).

plan		name = overall-treatment
filter-condition	parameter-proposition	
	parameter-name = pathologic-reason	
	value-description	type = equal
	qualitative-constant	value = false
	context:	any
plan-body	time-annotation: now	
	type = any-order	
	plan-activation	
	plan-schema name = regular-treatments	
	on-abort	
plan-activation	plan-schema name = exchange-transfusion	
	plan-schema name = exchange-transfusion	

Fig. 2. The plan overall-treatment. It is started after initial anamnesis unless ‘pathologic-reason’ is true. ‘Pathologic-reason’ is abstracted in a similar way as shown in Fig. 1. The plan ‘regular-treatments’ is started. If it fails, plan ‘exchange-transfusion’ is started. Alternatively, plan ‘exchange-transfusion’ is started immediately if its filter condition is fulfilled and that of ‘regular-treatments’ is not.

value false, which causes the treatment to start only if none of the exclusion criteria were true. This means that no pathologic reason was uncovered through the questions asking for abnormal findings and/or risk factors. Due to the nature of parameters in ASBRU, new values for any of the risk factors can be

entered at any time. On arrival of a new value, all conditions related to that value are evaluated again and appropriate changes are made.

6.2. Modelling the selection of a treatment option

After making sure the guideline is applicable, ruling out any pathologic reason, the appropriate treatment is decided according to the level of TSB. There are different limits according to the age of the baby (2, 3 days, and later). Therapeutic options are: observation and/or normal photo-therapy, normal photo-therapy, intensive photo-therapy, and exchange transfusion.

Table 1 shows the limits for the various treatment options. Note that neonates jaundiced on the first day of their life are excluded from the guideline, since early jaundice warrants special attention.

The association of the actual TSB reading with one of the treatment options is implemented using ASGAARD’s data abstraction unit. The time of measurement is entered together with the measured TSB value and forms the *abstraction context* for this value. This means that there are three different contexts (day 2, day 3, and later) for the transformation of the TSB reading (named

Table 1
TSB-levels and their associated treatments (adapted from [2])

Therapy recommendation	Age (h)		
	24–48	49–72	> 72
Observe or normal phototherapy	≥ 12	≥ 15	≥ 17
Normal phototherapy	≥ 15	≥ 18	≥ 20
Intensive phototherapy	≥ 20	≥ 25	≥ 25
Exchange transfusion	≥ 25	≥ 30	≥ 30

This table is encoded in the data abstraction used by the filter conditions for each treatment mode. It is evaluated constantly. Consequently, changes in the input lead to changes in the treatment.

TSB-raw) into the qualitative value suggesting a certain treatment (named TSB-qualitative), which are determined by the neonate's age at the time of the measurement. Another important value abstracted from the TSB reading is its change since the last measurement (TSB-change). For this parameter, a positive value indicates a decrease in TSB and a negative value indicates an increase.

The transformations of the TSB reading into the above mentioned derived values are performed automatically by the data abstraction unit with every new reading. This implies that, during the application of the guideline, the parameter TSB-qualitative will be updated according to the TSB variations considering the increased age of the patient to reflect the most appropriate therapy at any moment.

6.3. Modelling the treatment

The treatment is modelled in a hierarchy of plans. On the top-level of the guideline, some plans explicitly asking the user to enter necessary information are started before the *overall-treatment* plan shown in Fig. 2. This plan in turn starts either the *regular-treatments* plan or the *exchange-transfusion* one, depending on which of the filter conditions hold (the ASBRU any-order construct makes the execution to proceed in a non-deterministic way with the plans which have true filter conditions). In addition, if *regular-treatments* plan fails, *exchange-transfusion* is triggered. The regular treatments include the following alternatives: *observe*, *observe-or-normal-phototherapy*, *normal-phototherapy*, and *intensive-phototherapy*. The plan corresponding to the latter is presented in Fig. 3.

The qualitative abstraction of the TSB level (TSB-qualitative) directly corresponds to one of the treatment options. Therefore, each of them has at least a filter condition stating that TSB-qualitative must have the value asso-

ciated to this treatment, and an abort condition stating that this mode of treatment is stopped as soon as the TSB reading suggests another mode. Some treatments have additional conditions. So, for instance, *intensive-phototherapy* has two types of conditions: filter conditions and abort conditions. One part of the filter conditions states that this plan is activated if the value of TSB-qualitative is intensive-phototherapy. The first part of the abort conditions states that this plan is abandoned if the value of TSB-qualitative is not equal to intensive-phototherapy. The rest specifies that the plan should be aborted if the level of TSB does not decrease after 4 h or if it does but the decrease rate is too low (less than 1 mg/dl) within 4–6 h. Summarising, the abort conditions, without time annotations, can be seen as the logical expression:

((TSB-qualitative \neq phototherapy-intensive
or (TSB-change \leq 0)
or ((TSB-change $>$ 0) and
(TSB-change $<$ 1)))

For space reasons, some of these details have been omitted in Fig. 3. Lastly, the body of this plan solely consists of the invocation of the user-performed plan *perform-intensive-phototherapy*.

The *overall-treatment* plan has *pathologic-reason* (the non-eligibility criteria) in its filter condition (see Fig. 2). Besides, it has as implicit complete condition the completion of the treatments it comprises. The same happens with the abort condition, which implicitly depends on the abort of the treatments.

7. Discussion

In the above example we demonstrated various scenarios for the combination of

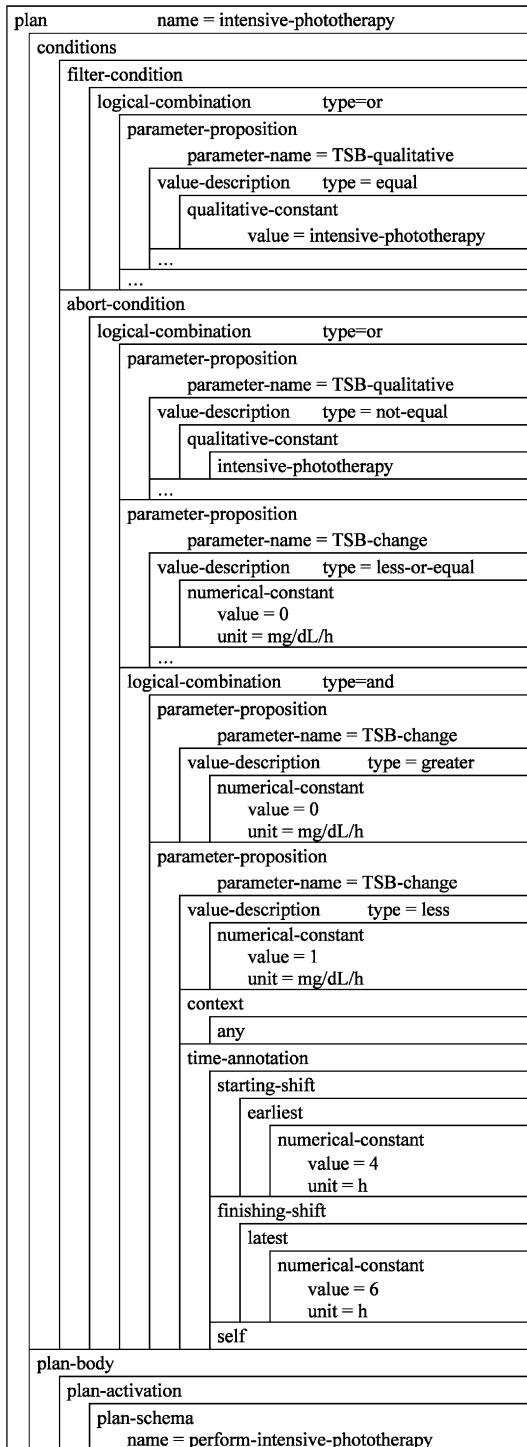


Fig. 3

diagnosis and treatment and their representation in ASBRU.

The eligibility criteria for the guideline as a whole and for a particular treatment step are represented as filter conditions of the plans representing them. The selection of the appropriate treatment is implemented by collecting the treatment options in a parent plan which executes one of them at a time according to the filter and abort conditions of each treatment plan. The actual treatment steps are modelled as user-performed plans, which contain explanatory text for the user to be displayed when the plan starts.

There are a series of other approaches to guideline modelling [5]. They differ from the ASGAARD system mostly concerning the context sensitive data abstraction, modelling of the temporal dimension, and the integration of the whole range of task-specific problem-solving methods around a uniform representation.

Comparable approaches are PROforma [6], GLIF [7], EON [8], and work by Quaglini et al. [9]. While these projects share many features with ASGAARD, they do not provide equal facilities for temporal data abstraction and continual support of both diagnosis and treatment. There are several projects dealing with temporal data abstraction with some support of diagnosis, e.g. the work by Chakravarty and Shahar [10] and Larizza et al. [11], which do not cover treatment planning.

Fig. 3. The plan intensive-phototherapy. This plan has two conditions, a filter condition and an abort condition. The filter condition states that this plan is activated if the value of 'TSB-qualitative' is 'intensive-phototherapy'. 'TSB-qualitative' is abstracted from the TSB readings and the age of the patient using Table 1. The abort condition states, that this plan is abandoned if either the value of 'TSB-qualitative' is not equal to 'intensive-phototherapy' or if 'TSB-change' is less or equal than 0 mg/dl per h or of 'TSB-change' is between 0 and 1 mg/dl per h 4 h after the start of the plan. The body of this plan solely consists of the invocation of the user performed plan 'perform-intensive-phototherapy'.

8. Conclusions

Computer support for both diagnosis and treatment leads to the improvement of health care quality and to the relief of health care staff. To be efficient, guideline modelling approaches must cope with the complexity of real-world scenarios.

Although some of the currently available guidelines and textbooks do not reveal this complexity, it comes into play when guidelines are integrated into clinical practice. The guideline representation language ASBRU provides means to model such complex scenarios. While there are more sophisticated systems for diagnosis and for treatment management alone, ASBRU joins the two allowing for close interaction between them.

Acknowledgements

We thank Frank van Harmelen, Annette ten Teije, Hugo Roomans, Geert Berger, Freia Kruizinga, and Robert Kosara for their contributions to modelling the guideline. This project is supported by ‘Fonds zur Förderung der wissenschaftlichen Forschung—FWF’ (Austrian Science Fund), P12797-INF.

References

- [1] M.E. Johnston, K.B. Langton, R.B. Haynes, A. Mathieu, Effects of computer-based clinical decision support systems on clinician performance and patient outcome: a critical appraisal of research, *Ann. Intern. Med.* 120 (1994) 135–142.
- [2] American Academy of Pediatrics, Provisional Committee for Quality Improvement and Subcommittee on Hyperbilirubinemia, Management of Hyperbilirubinemia in the Healthy Term Newborn, *Pediatrics* 94 (4) (1994) 558–565.
- [3] Y. Shahar, S. Miksch, P. Johnson, The ASGAARD project: a task-specific framework for the application and critiquing of time-oriented clinical guidelines, *Artif. Intell. Med.* 14 (1998) 29–51.
- [4] P.E. Friedland, Y. Iwasaki, The concept and implementation of Skeletal Plans, *J. Automat. Reason.* 1 (2) (1985) 161–208.
- [5] M. Peleg, S. Tu, J. Bury, P. Ciccarese, J. Fox, R.A. Greenes, R. Hall, P.D. Johnson, N. Jones, A. Kumar, S. Miksch, S. Quaglini, A. Seyfang, E.H. Shortliffe, M. Stefanelli, Comparing computer-interpretable guideline models: a case-study approach, *J. Am. Med. Inf. Assoc.* (2003), in press.
- [6] J. Fox, N. Johns, A.H. Rahmzadeh, Disseminating medical knowledge: the PROforma approach, *Artif. Intell. Med.* 14 (1–2) (1998) 157–181.
- [7] L. Ohno-Machado, J.H. Gennari, S. Murphy, N.L. Jain, S.W. Tu, D.E. Oliver, E. Pattison-Gordon, R.A. Greenes, E.H. Shortliffe, G.O. Barnett, The guideline interchange format: a model for representing guidelines, *J. Am. Med. Inf.* 5 (4) (1998) 357–372.
- [8] M.A. Musen, S.W. Tu, A.K. Das, Y. Shahar, EON: a component-based approach to automation of protocol-directed therapy, *J. Am. Med. Inf. Assoc.* 3 (6) (1996) 367–388.
- [9] S. Quaglini, L. Dazzi, L. Gatti, M. Stefanelli, C. Fassino, C. Tondini, Supporting tools for guideline development and dissemination, *Artif. Intell. Med.* 14 (1998) 119–137.
- [10] S. Chakravarty, Y. Shahar, A constraint-based specification of periodic patterns in time-oriented data, *Proceedings of the TIME-99*, IEEE Comp. Soc., 1999.
- [11] C. Larizza, R. Bellazzi, A. Riva, Temporal abstractions for diabetic patients management, in: *Proceedings of the Artificial Intelligence in Medicine (AIME-97)*, Springer, Berlin, 1997, pp. 319–330.

[1] M.E. Johnston, K.B. Langton, R.B. Haynes, A. Mathieu, Effects of computer-based clinical decision support sys-