Artificial Intelligence in Medicine (2006) **38**, 101—113

ARTIFICIAL INTELLIGENCE IN MEDICINE

http://www.intl.elsevierhealth.com/journals/aiim

POSITION PAPER

# Temporal representation and reasoning in medicine: Research directions and challenges

**Klaus-Peter Adlassnig [a], Carlo Combi [b], Amar K. Das [c], Elpida T. Keravnou [d], Giuseppe Pozzi [e,*]**

[a] *Section on Medical Expert and Knowledge-Based Systems, Core Unit for Medical Statistics and Informatics, Medical University of Vienna, Spitalgasse 23, A-1090 Vienna, Austria*
[b] *Dipartimento di Informatica, Università degli Studi di Verona, Strada le Grazie 15, I-37134 Verona, Italy*
[c] *Stanford Medical Informatics, Stanford University, 251 Campus Drive, X233 Stanford, CA 94305, USA*
[d] *Department of Computer Science, University of Cyprus, 75 Kallipoleos Str., CY-1678 Nicosia, Cyprus*
[e] *Dipartimento di Elettronica e Informazione, Politecnico di Milano, p.za L. da Vinci 32, I-20133 Milano, Italy*

Received 21 August 2006; received in revised form 5 October 2006; accepted 9 October 2006

**KEYWORDS**
Temporal representation;
Temporal reasoning;
Fuzzy sets and logic;
Health information systems;
Temporal clinical databases;
Medicine

**Summary**

*Objective:* The main aim of this paper is to propose and discuss promising directions of research in the field of temporal representation and reasoning in medicine, taking into account the recent scientific literature and challenging issues of current interest as viewed from the different research perspectives of the authors of the paper.
*Background:* Temporal representation and reasoning in medicine is a well-known field of research in the medical as well as computer science community. It encompasses several topics, such as summarizing data from temporal clinical databases, reasoning on temporal clinical data for therapeutic assessments, and modeling uncertainty in clinical knowledge and data. It is also related to several medical tasks, such as monitoring intensive care patients, providing treatments for chronic patients, as well as planning and scheduling clinical routine activities within complex healthcare organizations.
*Methodology:* The authors jointly identified significant research areas based on their importance as for temporal representation and reasoning issues; the subjects were considered to be promising topics of future activity. Every subject was addressed in detail by one or two authors and then discussed with the entire team to achieve a consensus about future fields of research.

* Corresponding author. Tel.: +39 02 2399 3649; fax: +39 02 2399 3411.
*E-mail addresses:* klaus-peter.adlassnig@meduniwien.ac.at (K.-P. Adlassnig), carlo.combi@univr.it (C. Combi), amar_das@med.stanford.edu (A.K. Das), elpida@cs.ucy.ac.cy (E.T. Keravnou), giuseppe.pozzi@polimi.it (G. Pozzi).

0933-3657/$ — see front matter © 2006 Published by Elsevier B.V.
doi:10.1016/j.artmed.2006.10.001

*Results:* We identified and focused on four research areas, namely (i) *fuzzy logic, time, and medicine*, (ii) *temporal reasoning and data mining*, (iii) *health information systems, business processes, and time*, and (iv) *temporal clinical databases*. For every area, we first highlighted a few basic notions that would permit any reader—including those who are unfamiliar with the topic—to understand the main goals. We then discuss interesting and promising directions of research, taking into account the recent literature and underlining the yet unresolved medical/clinical issues that deserve further scientific investigation. The considered research areas are by no means disjointed, because they share common theoretical and methodological features. Moreover, subjects of imminent interest in medicine are represented in many of the fields considered.

*Conclusions:* We propose and discuss promising subjects of future research that deserve investigation to develop software systems that will properly manage the multifaceted temporal aspects of information and knowledge encountered by physicians during their clinical work. As the subjects of research have resulted from merging the different perspectives of the authors involved in this study, we hope the paper will succeed in stimulating discussion and multidisciplinary work in the described fields of research.



## 1. Introduction

Temporal representation and reasoning in medicine has been investigated as a subject of research since the end of 1980s. During the 1990s a small but growing research community published its proposals and results in computer science journals as well as medical journals, and also presented them at general and specialized conferences [1–11]. Temporal representation and reasoning in medicine and, more generally, all research focusing on time-related aspects and medical information have certain peculiarities that distinguish them from other research topics in computer science applied to medicine.

Like general research on time-related topics in computer science, this research is orthogonal with respect to the required scientific and technical skills. Dealing with time to manage medical information requires that scientific results from various technical fields be shared, such as modal logic, constraint networks, data modeling and querying, deductive systems, real-time systems, fuzzy logic, data mining, and others. Time and medicine lead scientists through various technical disciplines.

A further distinguishing element common to the general field of computer science in medicine concerns the different medical tasks that must be taken into consideration: decision support for diagnosis, advice on therapy, clinical data summarization, intensive care unit (ICU) monitoring, and epidemiological studies, to mention a few.

A specific feature of this research area is that, despite changes in the considered medical applications and research topics, the overall interest in this area is slowly but consistently growing. In other words, although temporal representation and reasoning in medicine is not a trendy research topic that suddenly attracted the attention of a relatively large number of scientists, it has gained fairly solid ground. The sound state of this research is confirmed by the number of papers dealing with time and medical information, published in journals and conference proceedings of general and specialized events [2,4,10].

In view of these facts, predicting relevant issues of research in the near future is no simple task. The approach we used was to identify a few main research areas based on the expertise of the authors of the paper. A preliminary report was written for each area and the areas were then discussed with the entire team of authors. Finally, a few basic notions were formulated for each area. These have been formulated in a way that even a non-specialized reader will be able to comprehend the main features of each line of research, the relevant recent literature on the subject, and promising directions of research taking into consideration the yet unresolved challenges in the medical domain. This paper does not provide a comprehensive treatment of any aspect of research. Rather, the overall purpose is to stimulate discussion and multidisciplinary research in temporal representation and reasoning in medicine by merging different perspectives and various forms of expertise.

We identified four, possibly overlapping, research areas. The first two areas, named *fuzzy logic, time, and medicine*, and *temporal reasoning and data mining*, deal with a few well-known directions of research and discuss representing and reasoning over complex medical knowledge and data with artificial intelligence (AI) techniques. The last two areas, named *health information systems, business* *processes, and time*, and *temporal clinical databases*, take a closer look at the management of temporal clinical data and activities while considering the most recently developed techniques in the field of databases and information systems.

The paper is organized as follows. The four research areas are discussed in the particular order, as our purpose is to move from knowledge to data. Thus, Section 2 discusses *fuzzy logic, time, and medicine*. This is followed by Sections 3–5, dealing with *temporal reasoning and data mining*; *health information systems, business processes, and time*; and *temporal clinical databases*. A few concluding remarks are presented in Section 6.

## 2. Fuzzy logic, time, and medicine

If one follows Zadeh's original notion of a fuzzy set [12], one will immediately realize that temporal concepts such as "in the last few days", "more than four weeks ago", "at the same time", and many others are fuzzy by nature. What does this mean? According to ref. [12] (p. 339),

> A fuzzy set (class) $A$ in $X$ is characterized by a membership (characteristic) function $f_A(x)$ which associates with each point in $X$ a real number in the interval [0,1], with the value of $f_A(x)$ at $x$ representing the "grade of membership" of $x$ in $A$. ...When $A$ is a set in the ordinary sense of the term, its membership function can take on only the two values 0 and 1, with $f_A(x) = 1$ or 0 according as $x$ does or does not belong to $A$.

Thus, if we take "in the last few days" as a term in the ordinary sense by defining what we mean by "in the last few days"—let us assume we mean in the last three days—all events within the past 72 h fulfil this term, i.e., $f_A(x) = 1$, and all events outside this time period do not, i.e., $f_A(x) = 0$. What about 73, or 74, 75, ..., 82 h? Intuitively we feel that 73, 74, 75, ..., and even 81 h are, at least to a certain degree, compatible with this term. This is what fuzzy set theory allows us to do. It allows us to formally express (and further propagate by fuzzy logic or other fuzzy theories) the idea that 73 h are still compatible with the term "in the last few days" to a degree, say, 0.9. Furthermore, 74 h are then compatible to a degree 0.8, ..., 81 h to a degree of 0.1, until 82 h are no longer compatible with the term under consideration. Its degree of compatibility is then 0.

The same concept underlies the formalization of "more than four weeks ago". Four weeks are exactly 28 days, but what about 29 days, 30 days, and so on? Or "at the same time"—does it mean "exactly at the same time", or "nearly at the same time", or "more or less at the same time"? All of these terms are characterized by linguistic uncertainty; their semantic interpretation is not a crisp definition but a "fuzzily" defined term. This argument holds for all other temporal concepts and, in essence, for all concepts that denote classes of objects or values taken from the real world. "The only subsets of the universe that are not in principle fuzzy are the constructs of classical mathematics," says Kosko in ref. [13] (p. 268).

In medicine, concepts such as symptoms, signs, interpreted results from the laboratory or from clinical investigations, usually identified by a linguistic term, can be formalized by fuzzy sets. By doing so one captures their inherent "fuzziness", their uncertainty at the boundaries, their gradual transition to adjacent linguistic terms. For example, the transition from "... has a normal glucose level" to "... has an increased glucose level" is rather graded than crisp. The two fuzzy sets "normal" and "increased" (in the same context of glucose levels) overlap each other to a certain extent, allowing for measured borderline glucose levels being compatible with both terms, but only to a certain degree to each.

This is true for all aspects of a patient's examination. For instance, "increased body temperature (between 37 and 38 °C)" taken from a patient's history may be defined as a fuzzy set with blurred boundaries. Similarly, "limited mobility of the elbow joint", a sign observed during physical examination, or any measured laboratory test result that is mapped onto the linguistic terms "normal", "increased", "decreased", or others, can be defined by suitable fuzzy sets. Even radiographic signs such as "narrowing of joint space" may be considered to be formally represented as fuzzy sets.

Many of these medical terms that denote diagnostically and/or therapeutically relevant conditions in the patient have a temporal component. Examples are "morning joint stiffness $\geq$30 min" reported by the patient during history-taking or the physical signs "present joint swelling", "joint swelling in the last three months", or "joint swelling three months ago (and earlier)" obtained from the patient's physical examination.

In order to directly model these linguistic expressions, fuzzy sets that allow one to assign degrees of compatibility of what is observed in the patient to what the term stands for may be applied. For example, if morning joint stiffness is not very explicit or if the duration of morning joint stiffness is not definitely $\geq$30 min, this may be collectively expressed by assigning a degree of 0.8 to the fuzzy set "morning joint stiffness $\geq$30 min".

However, if we wish to express the observed finding and its time component separately, we may use two fuzzy sets and combine them by "and" to evaluate the composite term. Our example, "morning joint stiffness $\geq$30 min" may be divided into "morning joint stiffness" and "lasting $\geq$30 min". Here, the simplest "and" operation is performed by the minimum operator.

In ref. [14], fuzzy courses and fuzzy trend detection that depend on both measured signals and time are proposed for these situations. The original two-dimensional fuzzy graphs were extended to three-dimensional surfaces. This method was applied to derive the time of onset of toxoplasmosis infection from serological test results provided in a chronological sequence [15,16]. A second application followed using this fuzzy course and trend concept by defining fuzzy transitions in fuzzy automata [14]. A fuzzy automaton for clinical monitoring of patients with adult respiratory distress syndrome in ICUs was developed and tested. It models terms such as "adequate oxygenation", "hypoxemia", and similar concepts as three-dimensional fuzzy sets [17—19]. They all consist of a signal and a temporal component. These fuzzy trend templates for matching temporal patterns were also applied in ref. [20] and used to develop an intelligent anaesthesia monitor [21]. In a comparison of trend detection algorithms published in ref. [22], this fuzzy course approach to trend detection was deemed highly suitable for detecting temporal trends in real physiological signals due to its ability to identify underlying trends during high fluctuations in a signal and also in the presence of noisy information.

Pure temporal concepts such as "two to four weeks later" can be nicely defined as fuzzy sets. A knowledge-based system for interpretive analysis of toxoplasmosis serology test results was developed [23,24] that uses this type of fuzzy sets. In this system, the minimal temporal distance between antibody tests is thus expressed. If the necessary temporal distance between the tests is not fully provided the generated interpretive text is "softened". For example, a possible acute toxoplasmosis infection is indicated instead of a confirmed infection.

If we again follow Zadeh's conviction, which is best expressed in the following terms [25] (p. vii),

> . . ., it is important to recognize that any crisp theory *X* can be fuzzified — and hence generalized to fuzzy *X* — by replacing the concept of a crisp set in *X* by that of a fuzzy set. . . . What is gained through fuzzification is greater generality, higher expressive power, an enhanced ability to model real world problems and a methodology for exploiting the tolerance for imprecision and uncertainty to achieve tractability, robustness and low solution cost.

we shall see methodological attempts to fuzzify crisp formal concepts of temporal reasoning.

In this spirit, an attempt to extend and thus enhance a given crisp formal framework for medical knowledge representation by fuzziness was successfully made by developing FuzzyArden [26]. The Arden Syntax for Medical Logic Modules (MLMs) is a language for encoding medical knowledge—or parts of it—for computerized processing [27]. Each MLM contains rule-like medical knowledge (among other information) that is triggered by some clinical event to generate useful alerts and reminders, diagnostic interpretations, therapy proposals, as well as provide clinical research and administrative support. The Arden Syntax is a programming language that contains standard language elements such as operators (arithmetic, comparison, etc.), statements (assignment, if-then, loop, read, etc.), and others. It also includes several medically important constructs that involve time. Examples are "Now", "Is [Not] Within the Same Day As", "Occurs [Not] Before", and "A Time Delay After an Event" [28]. In FuzzyArden, all possible constructs that allow fuzzification were modified and extended. Fuzzy truth values, degrees of presence, and degrees of applicability that all take values between unity (completely true, present, applicable) and zero (completely false, not present, not applicable) were introduced to be able to work with graded data and knowledge. All Arden constructs involving time were also fuzzified. For example, the syntax statement for "is within the same day as" reads in FuzzyArden <n:fuzzy>: = <n:time> is within the same day as <n:date> [fuzzified by <n:duration>]. If the argument (<n:time>) does not fall exactly on the same day <n:date> but is within the "fuzzified by" argument (<n:duration>, e.g., 12 h) the evaluation assigns to <n:fuzzy> a truth value larger than zero (instead of zero as in common Arden). This value can be processed further and yields additional useful information. So far, FuzzyArden has been applied to the knowledge base of CADIAG-II [26] and to glaucoma monitoring [29]. Less systematically than in FuzzyArden, yet of the same nature, is the application of fuzzy relational operators in a hospital-information-system-oriented expert system shell for the representation and processing of temporal propositions about time series [30]. These fuzzy relational operators are internally represented as fuzzy sets. For intervals, either clear (crisp) or fuzzy boundaries are allowed. Of special interest in this report is a formula for a fuzzy count of occurrences.

Kiseliova and Wagner [31] show the process of generalizing a two-valued temporal logic to a multi-valued, fuzzy temporal logic. They even define a generalized time quantifier and a generalized compositional rule of inference suitable for approximate reasoning in temporal settings.

An early approach to fuzzy temporal reasoning was proposed by Kahn and Gorry [32]. They show a way in which knowledge over time can be incorporated into problem-solving programs. A time specialist program understands temporal specifications. They introduced a fuzzy expression that describes the uncertainty as to exactly when an event occurred, not the uncertainty of the event having occurred at all.

An extended formal attempt to represent and process fuzzy temporal knowledge was made by Dubois and Prade [33]. Zadeh's possibility theory [34] was used as a general framework for modeling temporal knowledge pervaded with imprecision or uncertainty, poorly known data, time intervals with fuzzy boundaries, fuzzy durations, and uncertain precedence relations between events.

In Dutta's event-based fuzzy temporal logic [35], uncertain events are represented as fuzzy sets, with the membership function giving the possibility of occurrence of the event in a given interval of time. This system was adapted in ref. [36] to develop a temporal model based on fuzzy set theory for determining temporal relationships.

Barro et al. [37] present a more complex model for the representation and handling of fuzzy temporal references. The concepts of date, time extent, and interval are defined again according to the formalism of possibility theory. Relations between the temporal entities, dates, and intervals are interpreted as constraints on the distance between dates, and mapped onto a fuzzy temporal constraint satisfaction network.

In ref. [38], Innocent and John show how fuzzy cognitive maps can encode fuzzy causal structures. Fuzzy temporal reasoning and constraints are then applied to support diagnostic decision-making.

A recent attempt to provide an algebraic framework for qualitative fuzzy temporal reasoning is given in ref. [39]. Fuzzy constraint satisfaction is applied to extend Allen's well-known [40] interval-based temporal framework.

Moreover, another approach to model the inherent linguistic uncertainty of temporal concepts and represent them in a fuzzy temporal knowledge base was done in Bugarín et al. [41]. As a formal descriptive tool a fuzzy Petri-net approach was used.

In summary, the following may be concluded. Although data and text mining in medicine is an upcoming research area and is generating useful practical results, explicit medical knowledge representation and processing are still required and depend very much on the acquisition of knowledge from medical experts. If—for a certain medical field or purpose—no theory or model is available in medical textbooks or in the minds of physicians, the first step is to develop a crisp theory or model. This will be extremely useful. If this crisp theory or model has been proven correct or at least applicable, the theory or model can be fuzzified and rendered more intuitive and, in many instances, provide a better mapping of the actual situation. However, in some instances a fuzzy model may be the only way to map entities and relationships between them from reality into a formal framework because it might not be possible to establish a crisp model at all; it would be too simplistic from the very start.

## 3. Temporal reasoning and data mining

Temporal reasoning in AI addresses the problem of inferring the state of a system at various points in time as it changes in response to events [42]. In the case of medical temporal reasoning, the states of interest are patient states and the events causing state changes are associated with disorder processes, therapeutic interventions, environmental/contextual factors, or even spontaneous/random events.

Time-based decision support in healthcare [10,43] addresses the tasks of diagnosis, treatment, prognosis, and monitoring of patients. The physiology and pathophysiology of human beings are complex processes, with many unknown aspects. The gaps in our knowledge increase as we proceed into the deeper genomic and molecular levels. For serious, life-threatening disorders, for which no treatment is yet known, such as various forms of cancer or degenerative disorders, this deep knowledge may be the only way to effective treatment. For example, discovering the relevant disorder mechanisms at the genomic level may lead to the discovery of appropriate drugs and treatments [44,45]. Even for disorders with known treatments, such deeper knowledge and understanding may lead to better and individualized treatments without adverse effects for the specific patient [44].

Given the fact that there may be several interacting processes—related to disease as well as therapy—in a patient, medical reasoning is complex by nature. The interactions in the specific situation (additive, subtractive, etc.), and the events causing them, need to be dynamically determined. Understanding and reasoning explicitly with time provides a more accurate representation of reality and may help to filter out impossible scenarios. In particular, the rate of evolution of disorders and the effects of treatment might differ between patients. One challenge would be to develop medical models that effectively combine causal, temporal and action knowledge, and enable the dynamic (contextual) derivation of different forms of interactions and the rates of evolution of the various operative processes.

Medical reasoning is inherently uncertain. The uncertainty is often modeled probabilistically [46,47]. Many researchers consider Bayesian networks particularly suited for capturing and reasoning with uncertainty. Prognostic Bayesian networks are a relatively recent development with applications in oncology and infectious diseases [47]. Even if temporal knowledge is not represented explicitly, prognostic Bayesian networks have a clear general temporal structure. When temporal reasoning comes into play, a new dimension of uncertainty is added. This dimension concerns uncertainty with regard to the occurrence of states, events, etc. Atemporal uncertainty deals with the uncertainty of state variables (What value does a state variable attain?) and their dependencies. When information about time is added, multiple temporal instantiations of state variables may occur. Temporal and atemporal uncertainty would therefore need to be combined and reasoned with. Mixing different time granularities is one way of controlling temporal uncertainties by reasoning at a temporal context that is conceptually appropriate for the entities involved [48]. Formalizing the different forms of uncertainty in medical reasoning and dealing with them through the effective combination of temporal-probabilistic-fuzzy knowledge for a specific medical task is another area of further research.

Clinical guidelines and protocols are (skeletal) temporal plans of actions giving different treatment options that need to be evaluated for a particular patient so that the "best" option is selected and appropriately instantiated before putting it into action. The application of the chosen regimen needs to be dynamically monitored (cf., intensive care patients) and, if necessary, modified. A plan may include choice points, and the selection of a particular choice may be based on the evolving reality. Representation and reasoning with clinical guidelines and protocols and explicit temporal constraints is a very active area of research offering a number of yet unresolved challenges [43]. The task of monitoring is also a key task in healthcare, particularly in ICUs. The purpose of monitoring is to observe how closely the evolved reality matches the projected evolution of patient states. If there is no projected evolution the task of monitoring is concerned with observing changes with respect to the patient's current state and to qualify these changes appropriately (steady, worsening, or improving situation), thus raising alarms. Bestowing medical systems, giving advice on therapy and/or monitoring patients, with the ability to change a temporal plan of action dynamically and in a timely manner by responding to relevant sensors, still needs to be fully addressed.

Temporal abstraction is a central functionality in medical temporal reasoning of relevance to all medical tasks, as it aims to close the (conceptual) gap between the knowledge level of process models (physiological processes, disorders, therapeutic processes, etc.) and the data level of patient-specific information [43]. For many medical problems, regardless of whether the data is sparse or voluminous, the gap is large. Unless the relevant temporal abstraction mechanisms are applied (vertical abstraction, horizontal abstraction, interpolation, etc.), it will not be possible to match data against knowledge—which is an essential step to make accurate suggestions and conclusions. Temporal data abstraction is not a widely encountered process in general AI work on temporal reasoning. In medical temporal reasoning it is of paramount importance. It is also a key process in temporal data mining for the discovery of temporal patterns and dependencies leading to new knowledge, as discussed, for example, in ref. [49].

Temporal data management and temporal data abstraction capabilities may be encapsulated in an intermediate mediator that is independent both of the application and the database [43]. The concept of such a data mediator facilitates the incorporation of different ontologies in different applications, which is a challenging subject [50].

Temporal reasoning may involve the solution of (temporal) constraint problems. The constraints are of different types (persistence, causal, synchronous, etc.) and temporal semantics (metric, relative, mixed, multiple granularities, etc.). Problems that need to be addressed are related to checking the consistency of a set of constraints and deciding about the satisfiability of some constraint with respect to a set of constraints that are assumed to be mutually consistent [51]. The complexity of these problems increases when mixed constraints and multiple granularities are used; this is another area of future study.

Data mining (DM) in general, and temporal data mining in particular, are acquiring increasing importance as fields of research with a considerable potential for the future. In medicine the increasing importance of the field is evidenced by the publication of special journal issues [45] and the allocation of major sessions to this subject at recent research conferences [52]. Top information technology companies are predicting that, by the turn of the present decade, information-based medicine will constitute the basis of personalized healthcare [44]. The term ''information-based medicine'' is coined to denote the deployment of artificial expert system technology and, in particular DM technology, for understanding several major diseases at the molecular level as a means of discovering drugs that can be effectively applied at a personalized level. This certainly is a fruitful and challenging field of research.

Genome-wide expression profiling of thousands of genes provides rich data sets that can be mined to extract information on the genes that best characterize the disease state. Identifying interactions between genes, based on experimentally obtained expression data in microarrays, is one of many significant research topics. The analysis of temporal patterns may reveal how the variables interact as a function over time—allowing observations and test results to be correlated, and permitting the researcher to detect previously unidentified patterns associated with disease. This will lead to a better understanding of disease mechanisms. Bayesian networks are used for the analysis of such biological time series. In general, the growing wealth of information and the advances in biology call for the development of approaches to discover new knowledge and predict drug effectiveness, which will enduringly change the methods used for prevention, diagnosis, and therapy.

## 4. Health information systems, business processes, and time

Business processes consist of the coordinated activities of an organization to achieve a goal which, possibly, is mission-critical for that organization. Clinical business processes deal with processes related to medical care and may be viewed from different perspectives. Indeed, clinical protocols and guidelines support physicians in their daily activities; guidelines and protocols help physicians to improve patient care and health outcomes for their patients by providing recommendations based on scientific evidence and expert clinical opinion. Guidelines and protocols usually consist of documents (e.g., available as pdf files) containing suggestions about actions and decisions (i.e., diagnosis and therapy) to be taken when dealing with patients who have specific problems. Such recommendations, even though provided in an informal and semi-structured way in the form of unstructured text, figures, lists of items, steps, etc. may be regarded as a description of a clinical business process the specific health organization has to execute in a specific situation. Protocols and guidelines have been considered in recent times. Several proposals have been made to support physicians and nurses in their execution and exchange, and to express them in a formal and sound way [53—60]. Moreover, from the organizational point of view, business processes in health organization have become a subject of major interest, being one of the issues to manage health budgets and to monitor the quality of services related to health [61—65].

Time is an important and underestimated aspect of modeling and managing processes within health information systems. A promising research area concerns the adoption, extension, and application of studies focusing on time-related aspects of business process modeling and workflow management to the challenging medical scenario.

Workflows are activities involving the coordinated execution of single atomic activities (named *tasks*), assigned and executed by processing entities (named *agents*) to achieve a common goal. Business processes can be modeled by workflows and enacted by suitable software systems, namely workflow management systems (WfMSs). WfMSs assign tasks to executing agents. The latter may be a human or an automatic executor such as a software system, or a combination of both, according to predefined policies. Most WfMSs use a database management system (DBMS) generally based on the relational model.

Activities are described by means of conceptual models (also named *schemas*) capturing the behavior of the process, and provide a formal process model. Instances of a schema are known as *cases*. Several types of conceptual models have been defined in the literature. We refer here to refs. [66—68]. Any WfMS in the market proposes its own model. Despite the efforts of the workflow management coalition (WfMC) [67], no widely accepted standard has yet been established for the definition of process models. Interoperable process models applicable to commercial WfMSs from different vendors are still far from being a reality.

WfMSs are suitable tools to organize long-standing activities that may be critical for the organization where the activity is enacted. In most of these activities, time management is not critical. However, when WfMSs are applied to health processes, temporal semantics cannot be neglected. This widens a new frontier for research over temporal data management and workflow systems. Indeed, health processes are processes whose temporal deadlines must be matched. Any deadline mismatch may lead to unacceptable consequences, as those we have observed in a nuclear plant management or a military defense system. As an example of a health (sub) process, let us consider (a part of) the guideline related to diabetes diagnosis [69]. There are some administrative tasks (such as making appointments for visits to nutritionists or for an eye examination) which can be executed in parallel. When the related clinical activities (the visits) are executed, a synchronization point is needed before one continues and makes an appointment with the diabetes specialist. Some deadlines and constraints must be managed, such as limiting the maximum delay between a reservation and the corresponding visit, or limiting the duration of the overall process. Without strict control over temporal constraints and deadlines for the completion of tasks, a WfMS could hardly manage health business processes.

Abnormal delays may occur during the enactment of cases within a WfMS, thus reducing the global performance of the managed business process. If no temporal information can be defined, this exceptional situation is hardly detected or requires that a complex software system has to be bundled into the WfMS [70]. The definition and management of temporal information of processes will improve the management of processes by a WfMS and widen the applicability of WfMSs to many different clinical areas, including, but not limited to, ICU process management, administrative health process management, clinical protocols support, and clinical guidelines support. In particular, some of the temporal aspects that can be considered are aimed at defining constraints over data specific to the process being managed by the WfMS (e.g., dates of visits for diabetic patients), over data of the workflow itself, and related to starting and ending time-stamps of tasks, and over data about agents involved in the processes.

The management of some temporal aspects is defined at the conceptual model level, e.g., when defining deadlines and temporal constraints between clinical tasks [71]. Moreover, histories of process schemas need to be managed: for instance, changes defined over a schema affect both cases that still have to be started, and cases that have started already but have not yet been completed [72,73]. The management of certain other temporal aspects is defined at runtime; for instance, workload balancing among agents can be performed only during case enactment. In general, considerable advantages may be achieved by the adoption of a temporal architecture for the WfMS [74]. Due to the absence of high-performance and reliable temporal database systems, no WfMS relies on a *temporal* DBMS. Besides those typical of the management of standard data, the main advantages of adopting a temporal DBMS within a WfMS are related to the execution of temporal queries as in schema version management [73], to the selection of the executing agent with load balancing over time, to exception management, and to temporal data warehouse analysis, to mention just a few.

With regard to these temporal aspects, the adoption of a *temporal* DBMS based on a data model where time and time dimensions are suitably represented and managed [75] could greatly facilitate the development of a system devoted to the automatic management of clinical workflows.

## 5. Temporal clinical databases

The developers of the earliest database system for clinical users realized the importance of maintaining and querying the temporal dimensions of clinical data [76]. This system used the time-oriented database (TOD), the first published temporal model for a database [77], which represents clinical data as a data cube with time-stamped visits as one axis. Significant efforts have been undertaken during the subsequent three decades to create database models that can capture the rich semantics of time needed for temporal querying and temporal reasoning in clinical decision support [10,78,79]. Past work [78—80] has shown that the many types of clinical data should be represented or queried as having persistence, using start and end time-stamps or a duration value. Previous research [81—83] has also taken note of the fact that the actual timing of clinical data is not exactly known; it is measured either by an interval during which the event occurred or by the use of varying time units, or granularities, for a time-stamp. Despite these advances in understanding how to effectively model the rich semantics of time within databases, we still are faced with several challenges when devising temporal database methods and techniques to be adopted by decision support, data mining, and workflow systems. The rapid growth of clinical data repositories and integration with other data types have resulted in computing applications that require new approaches for temporal representation and querying in database systems.

Although querying time-stamped data are necessary for many clinical decision support applications and other temporal reasoning applications, we lack standardized or widely available temporal database methods for such application developers to use. Many developers, in fact, adopt relational database systems—based on the relational data model, the most common data model in use for commercial database systems [84]—to store and manage time-oriented clinical data. The limitations of relational database systems in supporting the semantics of time was first noted 25 years ago [85]. Although relational database systems can readily store time values as attributes, they have two primary limitations in managing time values. First, the relational data model cannot specifically consider the semantics of time as a data element, since the relational model treats all columns equally. Second, query languages for relational databases, such as the widely used standard structured query language (SQL), have very few constructs to manipulate time values. Expressing complex temporal queries in SQL can thus be very difficult for users [86]. A decade ago the temporal database community made considerable efforts to achieve a consensus with regard to a temporal relational model and a related query language (the TSQL2 model and language) to resolve this problem [87]. Subsequently, efforts were made to involve national and international standards organizations: a new part for SQL3 (the more recent ISO standard version of the SQL language), named SQL/*Temporal*, was proposed and formally approved in July 1995 as part of the SQL3 draft standard [88]. Although extensions to SQL3 to manage temporal data were suggested, SQL/*Temporal* was withdrawn for a number of reasons [89]. As a matter of fact, most commercial relational database systems simply provide a few extensions to the relational model to represent times and dates through suitable data types, and a few extensions to SQL that allow manipulation of stored time values. However, these specifications have not been fully standardized across vendors.

As a result, developers of clinical systems are faced with difficulties when using relational database systems to manage time-oriented clinical data. First, since relational database systems provide no consistent mechanism to associate time values with non-temporal data in a table, joining multiple tables results in multiple time values in each row of the resultant table with no correspondence of time values and non-temporal data. Second, relational databases lack a standard approach to maintain the granularity of time values. This information is either lost or becomes implicit when time values are stored in a relational database. However, the exact granularity is needed to make accurate comparisons among research data that involve temporal ordering. Finally, relational database systems do not provide operations to coalesce rows of data in which non-temporal data are identical and time periods are overlapping or adjacent; these features may be needed to maintain the integrity of data as well as provide correct results about the duration of time periods. Several temporal querying methods for clinical relational databases [82,90—92] have been set up to resolve these problems in specific systems. However, the methods are not portable across systems.

Given the importance of querying time for clinical applications and the lack of standard temporal features in SQL, the specification of a standard, rich temporal relational model of clinical data still poses a challenge. Such work could start by evaluating the expressivity of TSQL2, and of extensions related to the SQL3/*Temporal* project, for clinical applications, and by adopting or extending the most suitable features. Such community-wide efforts may relieve future developers of the effort of writing customized software to provide proper storage and retrieval of time values, and thus avoid duplication of work undertaken by past researchers. Future work could also reference the semantics of a standard temporal relational model with a comprehensive temporal ontology, such as that proposed by [48]. If developers also used the latter knowledge representation to inform the time model in the clinical application, they would experience the additional benefit of not having to resolve temporal model heterogeneity within decision support systems.

Reference ontologies of time for databases may help users to manage the variety, volume, and scope of temporal data now being collected in biomedical data repositories through activities that range from ICU monitoring [93] to genomic studies [94]. Users of such repositories frequently find it difficult to compare or combine data represented with different temporal semantics. For example, the typical database of a clinical information system incorporates disparate sets of data from medical narratives, laboratory test results, and administrative codes [95]. These data must be merged together in clinical applications that present to healthcare providers the course and outcomes of a single patient's care. Such integration is necessarily time oriented, and thus requires an understanding of the temporal semantics of constituent data. Medical narratives, on the one hand, record clinical information as relative time references, such as "the patient was admitted to the hospital for angina (chest pain) that occurred four hours before a visit to the emergency room". Recent research [96] indicates that the ordering of most clinical events (e.g., angina or emergency room visit) in such text can be modeled as a directed acyclic graph (simple temporal problem). Time references associated with related administrative or diagnostic data, on the other hand, may consist of series of time-stamped calendar date values. New database methods are needed to allow comparison of relative time references with fixed time-stamped values. This would result in anchoring a constraint-based graph to a calendar date timeline, as well as providing a more specific order to the time-stamped data. Such temporal integration can also determine whether conflicting temporal information exists between the narrative report of events and the database sequence of events.

Biomedical data repositories provide other research challenges for temporal database researchers. A lot of previous work on temporal querying has focused on matching patterns for an individual patient's record. Using data repositories for clinical research, in contrast, requires population-based analyses, such as finding the average rate of change in a clinical parameter during a time window for a defined population of patients. Although such patterns may be handled by a temporal abstraction method, running such an algorithm to query thousands of cases may neither be feasible nor as optimal as it is when performed through database techniques. Large-scale data analysis is permissible in relational databases, but the standard SQL query language again is limited, as it only provides a small number of aggregate functions, such as sum, count, minimum, and maximum. Prior database research [97,98] has focused on efficient algorithms to compute aggregate functions over temporal groups, using one of two approaches: span aggregation or instant aggregation. In span aggregation, the aggregate value is computed for fixed intervals of time (such as month or year). Instant aggregation, in contrast, defines temporal groups based on computing from the temporal attribute values the intervals of time during which the aggregate value is constant. These methods are not sufficiently adequate to compute user-defined aggregate functions over clinically relevant temporal groups. For example, determining the average blood pressure measurement during periods of antihypertensive treatment cannot be determined by instant or span aggregation alone. The extended multi-feature (EMF) syntax for SQL [99] allows an arbitrarily complex aggregation to be directly specified into a database query. This method has been deemed useful to find temporal sequences of events within clinical databases that use an entity-attribute-value representation [100], as well as query a range of statistical aggregations over clinical data in a time window [101]. Further work is needed to determine and investigate the needs of clinical researchers working with large sets of temporal data, and to create efficient methods for statistical analysis with temporal representation and reasoning approaches.

## 6. Concluding remarks

In this position paper we have proposed and discussed a few promising subjects of future research on the issues that arise when one is faced with the task of managing the multifaceted temporal aspects of information and knowledge physicians encounter during their clinical activities. We identified four general research areas, presented basic features of each, discussed some related literature on the subject, and introduced research problems that are either unresolved or relatively new. Although the report does not profess to be comprehensive in terms of the cited literature or the challenging directions of research, we have tried to discuss several directions of research and merge different perspectives. We hope this report will stimulate discussion and multidisciplinary approaches in doing research on temporal representation and reasoning in medicine.

## References

[1] Keravnou ET. Medical temporal reasoning (editorial). Artif Intell Med 1991;3(6):289—90.
[2] Keravnou ET. Temporal reasoning in medicine (editorial). Artif Intell Med 1996;8(3):187—91.
[3] Goodwin SD, Hamilton HJ. It's about time: an introduction to the special issue on temporal representation and reasoning. Comput Intell 1996;12:357—8.
[4] Shahar Y, Combi C. Temporal reasoning and temporal data maintenance in medicine: issues and challenges. Time-oriented systems in medicine. Comput Biol Med 1997;27 (5):353—68 [special issue].
[5] Chittaro L, Montanari A. Editorial: temporal representation and reasoning. Ann Math Artif Intell 1998;22(1—2):1—4.
[6] Chittaro L, Montanari A. Temporal representation and reasoning in artificial intelligence: issues and approaches. Ann Math Artif Intell 2000;28(1—4):47—106.
[7] Temporal representation and reasoning. Ann Math Artif Intell 2000;30(1—4):1—2 [special issue].
[8] Bettini C, Montanari A. Temporal representation and reasoning. Data Knowl Eng 2003;44(2):139—41.
[9] Artale A, Dixon C, Fisher MD, Franconi E. Editorial: special issue on temporal representation and reasoning (time). J Log Comput 2004;14(1):1.
[10] Augusto JC. Temporal reasoning for decision support in medicine. Artif Intell Med 2005;33(1):1—24.
[11] Bouzid M, Combi C, Fisher MD, Ligozat G. Guest editorial: temporal representation and reasoning. Ann Math Artif Intell 2006;46(3):231—4.
[12] Zadeh LA. Fuzzy sets. Inf Control 1965;8:338—53.
[13] Kosko B. Neural networks and fuzzy systems. Englewood Cliffs, NJ: Prentice Hall; 1992.
[14] Steimann F. Diagnostic monitoring of clinical time series. PhD thesis. Technical University of Vienna, Vienna, Austria; 1997.
[15] Steimann F, Hayde M, Panzenböck B, Adlassnig K-P, Pollak A. Fuzzy support for serodiagnosis: the ONSET program. IEEE Eng Med Biol Mag 1994;13:705—9.
[16] Steimann F, Hayde M. A method to derive the time of onset of infection from serological findings. Methods Inform Med 1997;36:51—8.

[17] Steimann F, Adlassnig K-P. Clinical monitoring with fuzzy automata. Fuzzy Sets Syst 1994;61:37—42.
[18] Steimann F. The interpretation of time-varying data with DIAMON-1. Artif Intell Med 1996;8(4):343—57.
[19] Steimann F, Adlassnig K-P. Fuzzy medical diagnosis. In: Ruspini EH, Bonissone PP, Pedrycz W, editors. Handbook of fuzzy computation. Bristol, UK: IOP Publishing Ltd./Oxford University Press; 1998. p. G13.1:1—14.
[20] Lowe A, Harrison MJ, Jones RW. Diagnostic monitoring in anaesthesia using fuzzy trend templates for matching temporal patterns. Artif Intell Med 1999;16(2):183—99.
[21] Lowe A, Jones RW, Harrison MJ. The graphical presentation of decision support information in an intelligent anaesthesia monitor. Artif Intell Med 2001;22(2):173—91.
[22] Melek WW, Lu Z, Kapps A, Fraser WD. Comparison of trend detection algorithms in the analysis of physiological time-series data. IEEE Trans Biomed Eng 2005;52:639—51.
[23] Kopecky D, Rappelsberger A, Hayde M, Prusa A-R, Adlassnig K-P. Knowledge-based interpretation of toxoplasmosis serology test results including fuzzy temporal concepts. In: Proceedings (CD-Rom Version) of joint 9th IFSA world congress and 20th NAFIPS international conference; 2001. 2297.pdf.
[24] Anastassova N. Development and evaluation of TempToxopert—knowledge-based computer-aided decision support during prenatal screening for infection with toxoplasma gondii. MD thesis. Medical University of Vienna, Vienna, Austria; 2005.
[25] Zadeh LA. Foreword. In: Kruse R, Gebhardt J, Klawonn F, editors. Foundations of fuzzy systems. Chichester: John Wiley & Sons; 1994.
[26] Tiffe S. FuzzyArden: representation and interpretation of vague medical knowledge by fuzzified Arden Syntax. PhD thesis. Technical University of Vienna, Vienna, Austria; 2003.
[27] Hripcsak G. Writing Arden syntax medical logic modules. Comput Biol Med 1994;24:331—63.
[28] Health Level Seven. Arden Syntax for Medical Logic Systems, Version 2.5. 3300 Washtenaw Ave., Suite 227, Ann Arbor, MI 48104, USA: Health Level Seven, Inc.; 2005.
[29] Tiffe S. Defining medical concepts by linguistic variables with fuzzy Arden syntax. In: Kohane IS, editor. Proceedings of the annual symposium of the American medical informatics association AMIA 2002. Philadelphia, PA: Hanley & Belfus; 2002. p. 796—800.
[30] Qiu D, Dudeck J. Temporal propositions in a HIS-oriented expert system shell. Med Inform 1991;16:65—76.
[31] Kiseliova T, Wagner H. A generalized time quantifier approach to approximate reasoning. Fuzzy Sets Syst 2004;145(2):213—28.
[32] Kahn KM, Gorry GA. Mechanizing temporal knowledge. Artif Intell 1977;9(1):87—108.
[33] Dubois D, Prade H. Processing fuzzy temporal knowledge. IEEE Trans Syst Man Cybern 1989;19:729—44.
[34] Zadeh LA. Fuzzy sets as a basis for a theory of possibility. Fuzzy Sets Syst 1978;1:3—28.
[35] Dutta S. An event based fuzzy temporal logic. In: Proceedings of the 18th international IEEE Symposium on multiple valued logic. Los Alamitos, CA: IEEE Computer Society Press; 1988. p. 64—71.
[36] Dutta S. Temporal reasoning in medical expert systems. In: Proceedings of the IEEE symposium on the engineering of computer-based medical systems. Los Alamitos, CA: IEEE Computer Society Press; 1988. p. 118—22.
[37] Barro S, Marín R, Rodríguez-Patón A. A model and a language for the fuzzy representation and handling of time. Fuzzy Sets Syst 1994;61:153—75.
[38] Innocent PR, John RI. Computer aided fuzzy medical diagnosis. Inf Sci 2004;162(2):81—104.
[39] Badaloni S, Giacomin M. The algebra $IA^{fuz}$: a framework for qualitative fuzzy temporal reasoning. Artif Intell 2006;170: 872—908.
[40] Allen JF. Towards a general theory of action and time. Artif Intell 1984;23(2):123—54.
[41] Bugarín A, Cariñena P, Barro S. Techniques and applications of fuzzy systems based on the Petri-net formalism. Fuzzy Theory Syst Tech Appl 1999;4:1711—49.
[42] Fisher MD, Gabbay DM, Vila L, editors. Handbook of temporal reasoning in artificial intelligence. Amsterdam, Netherlands: Elsevier; 2005.
[43] Keravnou ET, Shahar Y. Temporal reasoning in medicine. In: Fisher MD, Gabbay DM, Vila L, editors. Handbook of temporal reasoning in artificial intelligence. Amsterdam, Netherlands: Elsevier; 2005. p. 587—652.
[44] IBM Business Consulting Services. Personalized healthcare 2010: are you ready for information-based medicine? [accessed: September 1, 2006]. http://www-1.ibm.com/services/us/bcs/html/bcs_health-care.html; 2003.
[45] Larrañaga P, Menasalvas E, Peña JM, Robles V. Editorial: data mining in genomics and proteomics. Artif Intell Med 2004;31(2):iii—iv.
[46] Hanks S, Madigan D. Probabilistic temporal reasoning. In: Fisher MD, Gabbay DM, Vila L, editors. Handbook of temporal reasoning in artificial intelligence. Amsterdam, Netherlands: Elsevier; 2005. p. 315—42.
[47] Lucas PJF, van der Gaag LC, Abu-Hanna A. Bayesian networks in biomedicine and health-care. Artif Intell Med 2004;30(3):201—14.
[48] Keravnou ET. A multidimensional and multigranular model of time for medical knowledge-based systems. J Intell Inform Syst 1999;13:73—120.
[49] Carrault G, Cordier M-O, Quiniou R, Wang F. Temporal abstraction and inductive logic programming for arrhythmia recognition from electrocardiograms. Artif Intell Med 2003;28(3):231—63.
[50] Boaz D, Shahar Y. A framework for distributed mediation of temporal-abstraction queries to clinical databases. Artif Intell Med 2005;34(1):3—24.
[51] Keravnou ET. Temporal constraints in clinical diagnosis. J Intell Fuzzy Syst 2002;12(1):49—67.
[52] Miksch S, Hunter J, Keravnou ET, editors. Artificial intelligence in medicine. 10th conference on artificial intelligence in medicine, AIME 2005 proceedings. Heidelberg, Germany: Springer; 2005 [vol. 3581 of lecture notes in computer science].
[53] Shahar Y, Miksch S, Johnson P. The Asgaard project: a task-specific framework for the application and critiquing of time-oriented clinical guidelines. Artif Intell Med 1998;14 (1—2):29—51.
[54] Duftschmid G, Miksch S. Knowledge-based verification of clinical guidelines by detection of anomalies. Artif Intell Med 2001;22(1):23—41.
[55] Terenziani P, Molino G, Torchio M. A modular approach for representing and executing clinical guidelines. Artif Intell Med 2001;23(3):249—76.
[56] Duftschmid G, Miksch S, Gall W. Verification of temporal scheduling constraints in clinical practice guidelines. Artif Intell Med 2002;25(2):93—121.
[57] Montani S, Terenziani P. Mapping clinical guidelines representation primitives to decision theory concepts. In: de Mántaras RL, Saitta L, editors. Proceedings of the 16th European conference on artificial intelligence (ECAI'2004). Amsterdam, Netherlands: IOS Press; 2004 . p. 1063—4.

[58] Terenziani P, Montani S, Bottrighi A, Molino G, Torchio M. Clinical guidelines adaptation: managing authoring and versioning issues. In: Miksch S, Hunter J, Keravnou ET, editors. Artificial intelligence in medicine, 10th conference on artificial intelligence in medicine, AIME 2005, proceedings. Heidelberg, Germany: Springer; 2005. p. 151—5 [vol. 3581 of lecture notes in computer science].
[59] Fuchsberger C, Hunter J, McCue P. Testing Asbru guidelines and protocols for neonatal intensive care. In: Miksch S, Hunter J, Keravnou ET, editors. Artificial intelligence in medicine, 10th conference on artificial intelligence in medicine, AIME 2005, proceedings. Heidelberg, Germany: Springer; 2005. p. 101—10 [vol. 3581 of lecture notes in computer science].
[60] Pellegrin L, Bonnardel N, Antonini F, Albanèse J, Martins C, Chaudet H. EORCA: a collaborative activities representation for building guidelines from field observations. In: Miksch S, Hunter J, Keravnou ET, editors. Artificial intelligence in medicine, 10th conference on artificial intelligence in medicine, AIME 2005, proceedings. Heidelberg, Germany: Springer; 2005. p. 111—20 [vol. 3581 of lecture notes in computer science].
[61] Kun LG. Telehealth and the global health network in the 21st century: from homecare to public health informatics. Comput Methods Prog Biomed 2001;64:155—67.
[62] Catarci T, Santucci G, Silva SF. An interactive visual exploration of medical data for evaluating health centres. J Res Pract Inform Technol 2003;35(2):99—119.
[63] Ghosh B, Scott JE. Comparing knowledge management in health-care and technical support organizations. IEEE Trans Inform Technol Biomed 2005;9(2):162—8.
[64] Berler A, Pavlopoulos S, Koutsouris D. Using key performance indicators as knowledge-management tools at a regional health-care authority level. IEEE Trans Inform Technol Biomed 2005;9(2):184—92.
[65] van Gendt M, ten Teije A, Serban R, van Harmelen F. Formalising medical quality indicators to improve guidelines. In: Miksch S, Hunter J, Keravnou ET, editors. Artificial intelligence in medicine, 10th conference on artificial intelligence in medicine, AIME 2005, proceedings. Heidelberg, Germany: Springer; 2005. p. 201—10 [vol. 3581 of lecture notes in computer science].
[66] Casati F, Ceri S, Pernici B, Pozzi G. Conceptual modelling of workflows. In: Papazoglou MP, editor. OOER'95: object-oriented and entity-relationship modelling, 14th international conference proceedings. Heidelberg, Germany: Springer; 1995. p. 341—54 [vol. 1021 of lecture notes in computer science].
[67] The Workflow Management Coalition. The workflow management coalition [accessed: September 1, 2006]. http://www.wfmc.org; 2005.
[68] van der Aalst WMP, van Hee KM. Workflow management: models, methods, and systems. Boston, MA: MIT Press; 2002.
[69] Hederman L, Smutek D, Wade V, Knape T. Representing clinical guidelines in UML: a comparative study. In: Surján G, Engelbrecht R, McNair P, editors. Medical informatics in Europe 2002. Amsterdam, Netherlands: IOS Press; 2002. p. 471—7.
[70] Casati F, Ceri S, Paraboschi S, Pozzi G. Specification and implementation of exceptions in workflow management systems. ACM Trans Database Syst 1999;24(3):405—51.
[71] Combi C, Pozzi G. Temporal conceptual modelling of workflows. In: Song I-Y, Liddle SW, Ling TW, Scheuermann P, editors. Conceptual modelling—ER 2003, 22nd international conference on conceptual modeling, proceedings. Heidelberg, Germany: Springer; 2003. p. 59—76 [vol. 2813 of lecture notes in computer science].
[72] Casati F, Ceri S, Pernici B, Pozzi G. Workflow evolution. Data Knowl Eng 1998;24(3):211—38.
[73] Combi C, Pozzi G. Towards temporal information in workflow systems. In: Spaccapietra S, March ST, Kambayashi Y, editors. Conceptual Modeling—ER 2002, 21st international conference on conceptual modeling, workshop proceedings. Heidelberg, Germany: Springer; 2002. p. 13—25 [vol. 2503 of lecture notes in computer science].
[74] Combi C, Pozzi G. Architectures for a temporal workflow management system. In: Haddad H, Omicini A, Wainwright RL, Liebrock LM, editors. Proceedings of the 2004 ACM symposium on applied computing (SAC). New York, NY: ACM; 2004. p. 659—66.
[75] Jensen CS, Snodgrass RT. Temporal data management. IEEE Trans Knowl Data Eng 1999;11(1):36—44.
[76] Weyl S, Fries JF, Wiederhold G, Germano F. A modular self-describing clinical databank system. Comput Biomed Res 1975;8(3):279—93.
[77] Wiederhold G, Fries JF, Weyl S. Structured organization of clinical data bases. In: Proceedings of AFIPS NCC. Montvale, NJ: AFIPS Press; 1975. p. 479—85.
[78] Kahn MG. Modeling time in medical decision-support programs. Med Decision Making 1991;11(4):249—64.
[79] Combi C, Cucchi G, Pinciroli F. Applying object-oriented technologies in modeling and querying temporally-oriented clinical databases dealing with temporal granularity and indeterminacy. IEEE Trans Inform Technol Biomed 1997;1:100—27.
[80] Shahar Y. A framework for knowledge-based temporal abstraction. Artif Intell 1997;90:79—133.
[81] Dolin RH. Modeling the temporal complexities of symptoms. J Am Med Inform Assoc 1995;2(5):323—31.
[82] Combi C, Pinciroli F, Pozzi G. Managing different time granularities of clinical information by an interval-based temporal data model. Methods Inform Med 1995;34:458—74.
[83] O'Connor MJ, Tu SW, Musen MA. Representation of temporal indeterminacy in clinical databases. In: Overhage JM, editor. Proceedings of annual AMIA symposium. Bethesda, MD: AMIA; 2000. p. 615—9.
[84] National Research Council Committee on Innovations in Computing and Communications. Funding a Revolution: Government Support for Computing Research. Washington, DC: National Academy Press; 1999.
[85] Jones S, Mason P, Stamper RK. LEGOL 2.0: a relational specification language for complex rules. Inf Syst 1979;4(4):293—305.
[86] Snodgrass RT. Developing time-oriented database applications in SQL. San Francisco, CA: Morgan Kaufmann Publishers; 1999.
[87] Snodgrass RT, editor. The TSQL2 temporal query language. Dordrecht, Netherlands: Kluwer Academic Publishers; 1995
[88] Eisenberg A, Melton J. SQL standardization: the next steps. SIGMOD Rec 2000;29(1):63—7.
[89] Hare KW. JCC's SQL standards page. [accessed: September 1, 2006]. http://www.jcc.com/sql.htm; February 2006.
[90] Das AK, Musen MA. A temporal query system for protocol-directed decision support. Methods Inform Med 1994;33(4):358—70.
[91] Nigrin DJ, Kohane IS. Temporal expressiveness in querying a time-stamp based clinical database. J Am Med Inform Assoc 2000;7(2):152—63.
[92] Desphande AM, Brandt C, Nadkarni PM. Temporal query of attribute-value patient data: utilizing the constraints of clinical studies. Int J Med Inform 2003;70(1):59—77.

[93] Saeed M, Lieu C, Raber G, Mark RG. MIMIC II: a massive temporal ICU patient database to support research in intelligent patient monitoring. In: Computers in cardiology. Los Alamitos, CA: IEEE Computer Society Press; 2002 p. 641—4.
[94] Rhee S-Y, Gonzales MJ, Kantor R, Betts BJ, Ravela J, Shafer RW. Human immunodeficiency virus reverse transcriptase and protease sequence database. Nucl Acids Res 2003;31 (1):298—303.
[95] Johnson SB. Generic data modeling for clinical repositories. J Am Med Inform Assoc 1996;3(5):328—39.
[96] Hripcsak G, Zhou L, Parsons S, Das AK, Johnson SB. Modeling electronic discharge summaries as a simple temporal constraint satisfaction problem. J Am Med Inform Assoc 2005;12(1):55—63.
[97] Kline N, Snodgrass RT. Computing temporal aggregates. In: Yu PS, Chen ALP, editors. Proceedings of the 11th international conference on data engineering (ICDE). Taipei, Taiwan: IEEE Computer Society; 1995. p. 222—31.
[98] Moon B, Vega López IF, Immanuel V. Efficient algorithms for large-scale temporal aggregation. IEEE Trans Knowl Data Eng 2003;15(3):744—59.
[99] Chatziantoniou D. The PanQ tool and EMF SQL for complex data management. In: Chaudhuri S, Madigan D, Fayyad UM, editors. Proceedings of the fifth ACM SIGKDD international conference on knowledge discovery and data mining. New York, NY: ACM; 1999. p. 420—4.
[100] Johnson SB, Chatziantoniou D. Extended SQL for manipulating clinical warehouse data. In: Proceedings of annual AMIA symposium. Bethesda, MD: AMIA; 1999 . p. 819—23.
[101] Rafiq MI, O'Connor MJ, Das AK. Computational method for temporal pattern discovery in biomedical genomic databases. In: Fourth international IEEE computer society computational systems bioinformatics conference (CSB 2005). Los Alamitos, CA: IEEE Computer Society; 2005. p. 362—5.