Temporal Reasoning and Temporal Data Maintenance in Medicine: Issues and Challenges

Carlo Combi[1] and Yuval Shahar[2]

[1]Dipartimento di Matematica e Informatica -Universita' degli Studi di Udine—Italy

[2]Section on Medical Informatics, School of Medicine, Stanford University, CA—U.S.A.

Abstract

We present a brief, nonexhaustive overview of research efforts in designing and developing time-oriented systems in medicine. The growing volume of research on time-oriented systems in medicine can be viewed from either an application point of view, focusing on different generic tasks (e.g., diagnosis) and clinical areas (e.g., cardiology), or from a methodological point of view, distinguishing among different theoretical approaches.

In this overview, we focus on highlighting methodological and theoretical choices, and conclude with suggestions for new research directions. Two main research directions can be noted: Temporal reasoning, which supports various temporal inference tasks (e.g., temporal abstraction, time-oriented decision support, forecasting, data validation), and temporal data maintenance, which deals with storage and retrieval of data that have heterogeneous temporal dimensions. Efforts common to both research areas include the modeling of time, of temporal entities, and of temporal queries. We suggest that tasks such as abstraction of time-oriented data and the handling of different temporal-granularity levels should provide common ground for collaboration between the two research directions and fruitful areas for future research.

Keywords: temporal reasoning, temporal maintenance, temporal databases, temporal abstraction, clinical data, medical informatics.

1. Introduction

Time is an important and pervasive concept of the real world and needs to be managed in several different ways: events occur at some time points, certain facts hold during a time period, and temporal relationships exist between facts and/or events [1]. Time has to be considered when representing information within computer-based systems [2], when querying information about temporal features of the represented real world [3], and when reasoning about time-oriented data [4].

Researchers in the medical-informatics field investigated temporal data modeling, temporal maintenance and temporal reasoning, to support both electronic medical records and medical expert systems [5, 6, 7, 8, 9, 10, 11, 12, 13, 14]. One indication of the significance of research on time-oriented systems in medicine is that the level and the amount of scientific works in this area motivated two different special issues of the journal *Artificial Intelligence in Medicine* [15, 16].

Another indication is that research focusing on time in clinical applications received attention also from the general computer science field [17, 18, 19, 20, 21, 22, 23, 24, 25, 26].

It is interesting to consider the wide variety of applications that need to deal with temporal aspects of clinical data, such as:

- management of time-oriented data stored in medical records of ambulatory or hospitalized patients [27, 28, 29, 11, 12, 30, 31, 32, 33, 34, 35, 36]
- prediction of future values of clinical data, given past trends [37, 38, 39]
- abstraction of time-oriented clinical data [8, 40, 41, 13, 42]
- time-oriented knowledge-based decision-support systems, such as systems supporting diagnosis, monitoring, or therapy planning [43, 26, 6, 44, 45, 46, 47, 14, 40, 48].

Studies of time-oriented applications have been performed in multiple clinical areas: cardiology [49, 50, 11, 30, 51, 47, 52, 53], oncology [45, 44, 8, 41, 35], psychiatry [26], internal medicine [54, 43, 34, 13, 46], intensive care [55, 45, 37, 56, 57], cardiac surgery [10], orthopedics [6], urology [34], infectious diseases [35], anesthesiology [38, 58, 56], pediatrics [52], endocrinology [59]. Various clinical tasks are supported by the systems proposed in literature: diagnosis [6, 60, 26], therapy administration and monitoring [40, 36, 53], protocol- and guideline-based therapy [44, 35, 13, 48], and patient management [27, 61, 62, 34, 29, 30].

In the paper, without being exhaustive in any way, we try to describe the main features (from a methodological point of view) that we can observe in the medical-informatics literature that deals with time-oriented systems. In order to clarify the theoretical basis for certain views we sometimes cite work from the general computer science area that is strictly related to the described topic. Several aspects related to the modeling of time-oriented clinical concepts are common to most of the approaches. We distinguish the set of temporal-reasoning subtasks from the set of temporal-maintenance tasks. We note that temporal-reasoning tasks, and in particular temporal abstraction, typically characterize medical decision-support systems, while management of temporal databases and handling of temporal granularity often characterize temporal data maintenance systems.

The rest of the paper is structured as follows: Section 2 provides a brief description of the main choices and problems related to modeling of time-oriented clinical concepts; Section 3 provides a brief overview of the research in temporal reasoning in medicine; In Section 4 we discuss temporal maintenance and databases for clinical information systems. Section 5 then presents in somewhat more depth the tasks of temporal abstraction and management of variable temporal granularity. These tasks form a potential bridge between the research community working on temporal data-maintenance systems and the one interested mainly in temporal-reasoning systems, in particular in time-oriented decision-support systems. Section 6 reflects on several open problems and suggests future research directions and challenges by merging the different research areas and methodologies. The summary in section 7, finally, concludes the paper.

2. Modeling temporal concepts

Both in temporal reasoning, in temporal abstraction of clinical data, and in modeling and managing clinical data, a common focus of effort is the definition or the adoption of a set of basic concepts that enable a description of the time-oriented clinical world in a sound and unambiguous way. Several suggestions have emerged from generic fields of computer science, such as artificial intelligence, or the knowledge and data management areas [63, 64, 65, 66, 3, 1]. Within medical informatics, this effort has progressed from an ad-hoc definition of concepts supporting a particular application to the adoption and the proposal of more generic definitions, supporting different clinical applications [45, 6, 8, 9, 10, 35, 30, 13, 14]. For example, the emphasis in the pioneering work of Fagan on the interpretation of real-time quantitative data in the intensive-care domain is on the application-dependent problems related to the support of a module that suggests the optimal ventilator therapy at a given time [45], while the work described in [14] uses a generic temporal ontology and a general, comprehensive, model of diagnostic reasoning.

Several related concepts involving time appear in the medical-informatics literature. We distinguish two related issues: modeling the concept of time and modeling entities having a temporal dimension [4, 1].

2.1. Modeling time

In modeling time for management of or reasoning about time-oriented clinical data several basic choices have to be done, depending on the needs of the domain.

2.1.1 Instants and intervals

Usually both the concepts of *time point* (or *instant*) and *time interval* have been used in the medical informatics literature to represent time [10, 13, 30, 14, 3]. These concepts are usually related to instantaneous events (e.g. myocardial infarction), or to situations lasting for a span of time (e.g. drug therapy). Care needs to taken in associating these concepts to clinical entities, such as symptoms, therapies, and pathologies: a myocardial infarction, for example, could be considered an instantaneous event, within the overall clinical history of the patient, or an interval-based concept, if observed during an ICU staying. A further distinction exists between the basic time primitives, usually instants (time points), and the time entities that can be associated with clinical concepts [35, 13, 30]. In defining basic time entities, time points (i.e., instants) are often adopted. Intervals are then represented by their upper and lower temporal bounds (start and end time points). In practice, most systems employed in medical informatics applications have used a time point based approach, similar to McDermott's *points* [67], rather than use time intervals as the basic time primitives, as proposed by Allen [64]. Several variations exist. Thus, Shahar [25] defines a set of time primitives, called *time stamps*; predicates, however, such as values of clinical parameters, can only be interpreted over *time intervals*, which are defined as ordered pairs of time stamps (including instants, which are zero-length intervals). This approach has been previously formalized in the artificial-intelligence literature by Shoham [68]. Time points usually are characterized by (possibly extended) properties of discrete numbers.

Thus, we observe two main approaches in associating time with clinical entities: the first deals both with instant-related entities and with interval-related entities [10]; the second associates clinical entities only with a certain type of time concept, usually an interval, dealing in a homogeneous way also with intervals degenerating to be a single instant [35, 30, 13].

2.1.2. Linear, branching and circular times

Different properties can associated with a time axis composed by instants: usually, both in general and clinically-oriented databases, time is *linear*: the set of time points is completely ordered [35, 30, 14]. However, for the tasks of diagnosis, projection, or forecasting (such as prediction of a clinical evolution over time), a *branching* time might be necessary. Such a representation has been found to be useful for pharmacoeconomics, and has been implemented using an object-oriented temporal model, as demonstrated in this issue [69]. *Circular* time is needed when we have to describe recurrent events, such as "administration of regular insulin every morning".

2.1.3. Relative and absolute times

The position on the time axis of an interval or of an instant can be given as an absolute position, such as the calendaric time when mapped to the time axis used [27, 8, 54, 70, 35, 13] (e.g.: "Tachycardia on November, 3 1996"). This is a common approach adopted by data models underlying temporal clinical databases. However, it also is common in medicine to reason with relative time references: "angina after a long walk" or "several episodes of headache during puberty". Incorporation of purely relative time-oriented, interval-based information (especially disjunctions, such as “the patient had vomited before or during the diarrhea episode”) within a standard temporal database is still a difficult task. More is said about that task when we discuss the issues of temporal granularity and uncertainty in Section 5).

Relevant to the topic of relative times are several proposals that employ implicit [6, 14, 52] or explicit [24] temporal contexts, which support the representation of relative or context-sensitive temporal clinical information or knowledge.

2.1.4. Modeling temporal relationships

In modeling temporal relationships, Allen's interval algebra [64] has been widely used in medical informatics [10, 35, 24, 25]. Extensions to Allen’s basic thirteen interval relationships have also been proposed [30]. Temporal relationships include two main types: qualitative (angina before headache) and quantitative (angina two hours before headache). Several general formalisms and approaches [64, 71, 72, 4] have been effectively adopted for satisfying the various needs encountered while modeling temporal relationships in clinical data.

2.2. Modeling temporal entities

A question that has been investigated in some depth in the medical informatics literature is: What are the basic medical concepts that have temporal dimension? How should time-oriented clinical data be modeled?

In general, we distinguish two different approaches in modeling temporal entities in medical applications: Addition of a temporal dimension to existing objects and creation of task-specific, time-oriented entities.

The first approach, originating from research into databases, uses simple, "atomic" temporal entities [35, 30]. This approach is similar to the one underlying the temporal extensions proposed for relational and object-oriented data models: A temporal dimension is added at the tuple/object level or at the attribute/method level [35, 73]. Combi et al. in [30] introduced the concept of temporal assertion to model in a homogeneous way both instant- and interval-based information. Das and Musen in [35] added the temporal dimension at the level of a database tuple, in a temporal extension of the relational model to support clinical databases used for decision-support application. In Das’s approach, the focus is on modeling time-oriented clinical data, to allow the DBMS to store and manage this type of data. In such approaches, complex temporal features of clinical data can be queried by suitable query languages [35, 73, 58].

One of the first applications of databases to clinical domains, explicitly addressing the time representation problem, is the Time Oriented Database (TOD) model [27], originally developed at Stanford university during the 1970s. This model has been adopted, for example, by the American Rheumatism Association Medical Information System (ARAMIS), to manage data related to the long-term clinical course of patients suffering from arthritis or, more generally, from rheumatic pathologies [74]. TOD uses a "cubic" vision of clinical data: values of data related to a particular patient visit are indexed by patient identification number, time (visit date), and clinical- parameter type. Specialized time-oriented queries enable researchers to extract, for particular patients, data values that follow certain simple temporal patterns (e.g., increase at some rate). Assignment of a temporal dimension at the tuple level is a method common to many applications of clinical databases [75, 31, 11].

The second approach, originating mostly from the area of artificial intelligence in medicine, focuses on modeling different temporal features of complex, task-specific entities. The temporal entities are defined by the needs of the relevant temporal-abstraction and, in general, temporal-reasoning tasks (see Section 3). Based on temporal entities that are stored at the database level, several types of compound (abstract) entities are introduced. For example, in the HyperLipid system [43], patient visits were modeled as instant-based objects called *events*, while administration of drugs was modeled as *therapy* objects whose attributes included a time interval. *Phases* of therapy (inspired by the clinical algorithm modeled by the system) were then introduced to model groups of heterogeneous data that is related to both visits and therapies. Events, therapies and phases were connected through a network.

Kahn and colleagues in [8] introduced formally the concept of a Temporal Network (TNET) and later extended it by the Extended TNET, or ETNET model [9]. In both models, a T-node (or an ET-node) models task-specific temporal data, such as a chemotherapy cycle, at different levels of abstraction. Each T-node is associated with a time interval during which the information represented by the T-node's data is true for a given patient.

In the M-HTP system for monitoring heart-transplant patients [10], clinical facts related to a patient are structured in a temporal network (TN) inspired by Kahn's TNET model [9]. Through this network, a physician can obtain different temporal views of the patient's clinical history. Each node of the TN represents an *event* (a *visit*) or a *significant episode* in the patient's clinical record. An event is time-point based; its temporal location can be specified by an absolute date or by the temporal distance relative to the transplantation event. A episode holds during an interval, during which a predefined property (evaluated by reasoning about several events) holds.

Keravnou and Washbrook introduce *findings*, *features*, and *events* to distinguish various types of instantaneous and interval-based information (patient-specific or general) [6].

3. Temporal Reasoning

Temporal reasoning has been used in medical domains as part of a wide variety of generic tasks [76], such as diagnosis (or, in general, abstraction and interpretation), monitoring, projection, forecasting, and planning. These tasks are often interdependent. **Projection** is the task of computing the likely consequences of a set of conditions or actions, usually given as a set of cause–effect relations. Projection is particularly relevant to the *planning task* (e.g., when we need to decide how the patient's state will be after we administer to the patient a certain drug with known side effects). **Forecasting** involves predicting particular future values for various parameters given a vector of time-stamped past and present measured values, such as anticipating changes in future hemoglobin-level values, given the values up to and including the present. **Planning** consists of producing a sequence of actions for a care provider, given an initial state of the patient and a goal state, or set of states, such that that sequence achieves one of the goal patient states. Possible actions are usually operators with predefined certain or probabilistic effects on the environment. The actions might require a set of enabling *preconditions* to be possible or effective. Achieving the goal state, as well as achieving some of the preconditions, might depend on correct *projection* of the actions up to a point, to determine whether preconditions hold when required. **Interpretation** involves abstraction of a set of time-oriented patient data, either to an intermediate level of meaningful temporal patterns, as is common in the **temporal-abstraction** task or in the **monitoring** task, or to the level of a definite diagnosis or set of diagnoses that explain a set of findings and symptoms, as is common in the **diagnosis** task. Interpretation, unlike forecasting and projection, involves reasoning about only past and present data and not about the future.

From the methodological point of view, one general criterion that can be used when classifying temporal-reasoning research that had been applied to clinical data is whether it uses a deterministic or a probabilistic approach [16].

Within the deterministic approach, different frameworks have been used. Some of them are based on well-known formalisms from the artificial-intelligence field [55, 26]; others are based on ad-hoc rules and/or ontologies [22].

The probabilistic approach typically is associated with the tasks of interpretation or forecasting of time-stamped clinical data whose values are affected by different sources of uncertainty [46, 47]. Dagum and his colleagues [37, 38] have developed the Dynamic Network Models (DNMs) methodology, a synthesis of belief-network models and classical time-series models, and have applied them with encouraging results to domains such as predicting outcomes of critically-ill intensive-care patients [37] and forecasting episodes of apnea in sleep-apnea patients [38]. Causal Probabilistic Networks (CPNs) is a graphical formalism, widely used in probabilistic systems applied to clinical problems [46]. Several studies examine the extension of CPN and other formalisms by also considering time for clinical problems [39, 37, 46]. Aliferis et al. [77] use for their analysis of temporal abstraction a framework they have defined previously, Modifiable Temporal Belief Networks (MTBNs), while Ngo et al. [78] propose a language for representing context-sensitive temporal probabilistic knowledge, based on a standard formalism for representing belief networks.

Several solutions have been proposed to the commonly occurring temporal-abstraction task, namely, the interpretation task of reasoning about high-level concepts (e.g., a pattern of bone-marrow toxicity specific to a particular chemotherapy-related context) that can be abstracted from time-oriented clinical data (e.g., a time-stamped series of chemotherapy-administration events and various hematological laboratory tests) [6, 8, 9, 10, 41, 13] (see also the discussion in Section 5). For example, in M-HTP [10] the white blood-cell (WBC) count, measured during a visit, is an instantaneous event in the knowledge base, indexed by the visit date; WBC-count decrease is an episode, spanning several days, detected by the values of WBC count. Similarly, in TNET [8] T-nodes are able to describe at different levels of abstraction data related to a patient undergoing different chemotherapy treatments. In both cases, a temporal knowledge-based reasoner that uses IF-THEN rules is applied to the system's temporal model of the patient [9, 10]. These systems are able to deal with complex temporal conditions. A typical rule is, for example, [42]:

```
IF   DURING last 10 days ARE PRESENT
         low CMV antigenemia OF TIME SPAN at least 7 days
            AND
         leukopenia OF TIME SPAN at least 5 days
  AND
     DURING last 15 days IS NOT PRESENT
```

```
        CMV infection OF TIME SPAN at least 1 day
THEN CMV infection is highly suspected
```

An example of a diagnostic-support system is the Skeletal Dysplasia Diagnosis (SDD) expert system and its temporal reasoning framework, which have described by Keravnou and Washbrook [6]. The SDD system is designed to provide a diagnostic aid in diagnosing skeletal dysplasias and syndromes. The temporal reasoner module in the SDD system has a layered architecture and is able to deal with findings and features in order to provide the user with higher-level representations of findings and dysplasia expectations for the given patient. Even if the temporal framework of the SDD system has a somewhat practical orientation, the influence of the event calculus theory [79] should be emphasized [6].

An interesting task that involves some amount of abstraction as well as a certain amount of forecasting, and which typically requires both deterministic and probabilistic techniques, is the **validation** of time-oriented data, possibly also including the suggestion of specific repairs. An example of applying several such techniques in a knowledge-based system is presented by Horn et al. [57] in this issue. Horn and his colleagues propose several time-oriented methods for validating and repairing high-frequency clinical data; they have applied these methods to the neonatal intensive-care (in particular, to the artificial-ventilation) domain.

A different methodological approach to temporal reasoning is Case Based Reasoning (CBR). CBR is a research field in which problems related to the acquisition of domain knowledge are partially avoided by using previous known and solved cases when addressing a new case [26]. The considered patient case at hand is compared to similar previous cases, on the basis of general and temporal features.

In recent years, a number of research efforts were aimed at the extension and generalization of several results and properties of temporal reasoning systems in clinical domains. We distinguish two different approaches: several researchers try to define a general framework, catering for different needs that stem from different clinical applications [22, 25, 24, 52]; others try, instead, to apply general theories from the temporal reasoning field in the artificial intelligence area, such as the event calculus, case-based reasoning, temporal constraint networks, and belief networks, to solve clinical problems [80, 60, 81, 55, 26, 56].

4. Maintenance of Time-Oriented Data: Temporal Databases

Moving on to the task of management of time-oriented clinical data, we observe that the literature has progressed from the early systems, which were mostly application dependent, to more general approaches, that, even when applied to the solution of real problems in management of time-oriented clinical data, have a more generalizable value and inherent soundness [27, 82, 8, 83, 11, 84, 35, 62, 30, 73]. Initially, systems that were designed to manage temporal clinical data were based on the flat relational model [27, 82]. These systems, such as Wiederhold's TOD and Blum's Rx, were based on time stamping the database tuples: the date of the visit was added to the specific attribute values. Kahn et al. [8, 9] proposed a specific query language, TQuery, for data that is structured by the TNET model. Even though TQuery was patient oriented and was not based on a generic data model, it was one of the first proposals for an extension of query languages so as to enable the system to retrieve complex temporal properties of stored data. Most query languages and data models used for clinical data management were application-dependent; thus, developers had to provide ad-hoc facilities for querying and manipulating specific temporal aspects of data [9].

Recent works on temporal clinical databases present a more general approach. Extensions of common data models, and in particular of relational models, are based also on the general database-field literature, in which temporal databases enjoyed special attention in the past few years [66, 85, 3, 1]. In the temporal-database field, beside the previously cited problems in modeling time and entities to which a temporal dimension is added, there is also the problem of what *kinds* of temporal dimensions need to be supported by the temporal database. Snodgrass [65, 86] had previously identified three different temporal dimensions:(1) the *transaction time*, that is, the time at which data are stored in the database (e.g., the time in which the assertion "WBC count is 7600" was entered into the patient's medical record); (2) the *valid time*, that is, the time at which the data are true for the modeled real world entity (e.g., the time in which the WBC-count was, in fact, 7600), and (3) the *user-defined time*, whose meaning is related to the application and thus is defined by the user (e.g., the time in which the WBC count was determined in the laboratory). In the Snodgrass taxonomy, four kinds of databases can be defined: (a) *snapshot databases*, based on flat, timeless data models; (b) *rollback databases*, which represent explicitly only the transaction time, (c) *historical databases*, which represent explicitly only the valid time [65, 66], and (d) what is now called *bitemporal databases*, which represent explicitly both transaction time and valid time and thus are both historical and rollback.

Another distinction useful for the characterization of much of the research in maintenance of temporal data is whether the main topic is definition of temporal data models [8, 29, 30, 23] or definition and design of temporally-oriented query languages [35, 36, 73, 58]. Both topics are discussed in the medical informatics literature, although we note a more focused interest in data modeling.

In medical informatics, attention had been paid mostly to historical databases (which emphasize valid time), extending relational or object-oriented models [83, 35, 73, 58]. Thus Das and Musen [35], for example, define four different types of relational tuples: event, start, body, stop, to define, respectively, instantaneous facts and three aspects of uncertainty about interval-based facts (i.e., uncertainty regarding the start time of the fact, uncertainty regarding the end time of the fact, and a certain period of time in which the fact held). Das proposes an extension of the relational algebra to manage temporal information and temporal relational operations, and an extension of SQL, based on the proposed algebra, called Time Line SQL (TLSQL). One of the goals of Das's representation is to facilitate the handling of temporal uncertainty and in particular the management of data represented at variable levels of temporal granularity, a task that we discuss in Section 5. Combi et al. [73] extended an object-oriented data model and the related query language to deal with temporal clinical data: Granular Clinical History - Object SQL (GCH-OSQL) was proposed as a query language for temporal clinical databases, taking into account different and mixed temporal granularities. Goralwalla and colleagues adapted an existent object database model to the management of time-oriented data, and have applied it to the modeling of pharmacoeconomic clinical trials [69]. The broad set of types supported by the adopted object data model enables, for example, a modeling of branching timelines, corresponding, for instance, to the evaluation of different pharmacological treatments.

5. Temporal Abstraction and Management of Temporal Granularity

Two commonly recurring and closely related tasks in both the temporal-reasoning and the temporal-maintenance research areas are (1) the temporal-abstraction task mentioned in Section 3, and (2) the handling of variable temporal granularity mentioned in Section 4. Since these tasks are highly relevant to both research communities, they might be viewed as one of the potential bridges between them (besides fundamental issues mentioned in Section 2, such as the time model). The two tasks have been investigated both in the medical informatics field and in the general computer-science area [4].

Temporal abstraction provides a more powerful, concise, and integrated description of a collection of time-stamped raw data. (The term “temporal abstraction” is somewhat misleading, it is the time-oriented data, and not the time itself, which are being abstracted.) In the medical-informatics field, temporal abstraction plays a central role in supplying care providers with data at a level suitable for support of decision making. Temporal abstraction in general, and in medicine in particular, has been investigated in some depth in recent years [10, 42, 17, 18, Haimovitz96, 41, 24, 13, 25, 8, 40]. An interesting examination of the utility of temporal abstraction is provided in this issue by Aliferis and colleagues [77], in a study that addresses the problem of providing and evaluating appropriate levels of temporal abstraction using a common formalism for medical decision-support systems.

Management of variable temporal granularity deals, in fact, with an abstraction of the time primitives themselves; it concerns the level of abstraction (e.g., time unit, such as a day or a month) at which the time element (instant, interval, and so on) associated with the relevant data is represented [4]. Using this definition, we note that the tasks of temporal abstraction and of handling variable temporal granularities are interconnected. When reasoning about various temporal-granularity levels, emphasis is placed on the abstraction of the representation of the *time component* of a time-oriented assertion; when performing a temporal-abstraction task, the emphasis is placed on the abstraction of the time-oriented *entity* itself.

We adopt here the a framework proposed in [30], that facilitates comparisons between systems and frameworks that handle different time granularities. Three main types of temporal granularity were identified:

• *abstraction granularity*: This granularity-management aspect is not related directly to the time axis. Abstraction granularity refers to the ability to express complex and composite temporal concepts, for example "A parameter value that is increasing during a period of three months".

• *absolute-time granularity*: This is the ability to express the temporal dimension of the data by mixing and using different absolute time references, for example "The interval began within the period from January 21, 15:23, to January 21, 16:34." Absolute-time granularity refers to the uncertainty in specifying a temporal dimension or to the use of different time-units.

• *calendar-date granularity*: This is the capacity for expressing the temporal dimension through the use of multiple time units, e.g. years, months, days.

5.1. Abstraction granularity

Medical decision-support systems often do not associate the granularity of time with the calendar time. Rather, the temporal granularity level is affected by the abstraction needed by the relevant clinical problem [9, 6, 87, 10].

Many representations of temporal data at high abstraction levels in medical expert systems were inspired by Allen’s interval-based logic [63, 64, 43, 10]. Kahn’s TNET and ETNET models (see Section 2.2) aimed at extending the TOD model (see Section 2.2), by defining suitable persistent objects [9]. TNET is composed of T-nodes: each T-node represents a time interval during which a clinical event happened. The starting and ending time instants identify the time interval. Clinical events are organized within a hierarchical structure that corresponds to a model of significant clinical contexts. ETNET adopts the same temporal structure of TNET [9, 87]; in addition, ETNET associates with each ET-node certain computational methods (rules) that lead to the conclusion of new information about the relevant events.

Shahar and Musen [41, 13, 25] proposed a general framework for abstraction of time-stamped data, and in particular of clinical data, called the Knowledge-Based Temporal-Abstraction (KBTA) Method. The KBTA framework includes a theoretical model for time and for propositions that hold over time, a general inference method, and five specific computational temporal-abstraction mechanisms that solve the five subtasks into which the KBTA method decomposes the temporal-abstraction task. The five mechanisms are *context formation, contemporaneous abstraction, temporal inference*, *temporal interpolation*, and *temporal pattern matching* [13, 25]. The output of these mechanisms includes abstractions of type *state*, *gradient*, *rate*, and *pattern* (e.g., LOW, DECREASING, and FAST abstractions for the hemoglobin-value clinical parameter, and the QUIESCENT-ONSET-CGVHD pattern abstraction in the domain of chronic graft-versus-host disease (CGVHD)). The five mechanisms require four well-defined domain-independent types of domain-specific knowledge: *Structural*, *classification* (functional), *temporal-semantic* (logical), and *temporal-dynamic* (probabilistic) knowledge [25]. The KBTA method had been implemented by the RÉSUMÉ system [41] and has been evaluated within several clinical domains, such as oncology, therapy of patients who have AIDS, monitoring of children's growth, and management of insulin-dependent diabetes [13]. The RÉSUMÉ system uses as temporal primitives time stamps at various predefined levels of granularity, typically offset from a clinically relevant time stamp, such as the time of bone-marrow transplantation, the beginning of chemotherapy, or the date of birth of the patient (e.g., for monitoring children's growth). Input data or output abstractions can hold, however, only during time intervals, defined as ordered pairs of time stamps. Highly complex patterns can be described and computed, but the set of granularity levels (and therefore also the implied temporal uncertainty) is limited to a predefined one that includes minutes, hours, etc.

5.2. Absolute-time granularity

The necessity to sometimes provide absolute-time granularity, that is, the capability to refer to the time-axis in multiple ways, not only through different time units, has been addressed by several recent works in medical informatics [60, 35, 30, 62, 22]. Two different issues have to be addressed when providing absolute-time granularity. The first is the representation of vagueness, uncertainty, or indeterminacy (the terminology varies among different scientific communities) regarding the location on the time axis of relevant time points or time intervals [3, 30, 22]. The second is the use of time units or references that include not only those associated with the Gregorian calendar, but also domain-specific ones (e.g., *weeks-from-conception*, *fetal-period*, *infancy* [14]) [30, 35].

Das and colleagues [35] proposed an extension to the relational model and to the query language SQL; they introduced the concept of Interval Of Uncertainty (IOU) to model the uncertain time intervals that include the starting and ending instants of the interval of validity of a tuple. Thus, representing a relational-database entity that is valid during an interval with indeterminate start and stop instants involves representing explicitly the uncertain start, the certain body, and the uncertain end interval [35] (see Section 4). Combi and colleagues [30, 73] have described a data model using two different formalisms, both based on an object-oriented approach, that is able to represent intervals and time points given at different and mixed absolute time-granularity, such as the interval referred to in the sentence "an atrial fibrillation episode occurred on December 14th, 1995 and lasted for three minutes".

Uncertainty in clinical domains has to be handled when relationships are computed between intervals (or time points) that have different absolute-time granularities. Several researchers have proposed an extension of the classic two-valued logic to a model that enables also a representation of *possibly true* (or *possibly false*) values to temporal relationships [30, 60].

5.3. Calendar-date granularity

Medical applications require in a natural manner systems that are able to represent and manage different time units [35, 58]. This type of temporal granularity is a common one and has been widely studied in the temporal database community [3, 88].

6. Future Directions: Merging Different Areas and Approaches

Temporal reasoning systems and temporal data-maintenance systems are often independent efforts, even though they usually contribute towards the same goal. For example, time-oriented decision-support systems often do not adopt any kind of a formal temporal data model or a temporal query language to manage stored time-oriented clinical data.

We suggest that currently, after several years of research on the topics described in Sections 2 to 5, new and more powerful solutions could be derived from a merging of different approaches.

The temporal-abstraction task and the management of temporal granularity seem to be a meeting point between research efforts originating in the artificial-intelligence and in the database communities, at least as these efforts have been applied to medical domains. Furthermore, as pointed out in Section 2, the issue of the appropriate time model is always a pertinent one. Thus, several research themes, most of which are relevant to the community of general computer scientists, will be important, in our opinion, for next-generation time-oriented systems in medicine.

- **Adoption of advanced data models**. The adoption of advanced data models, such as the object-oriented data model and the EER data model, will improve the capability of describing real world clinical entities at high abstraction levels [89, 90, 91, 92]. Thus, the focus may shift to more domain-specific inference actions.
- **Maintenance of clinical raw data and abstractions**. Several recent systems allow not only the modeling of complex clinical concepts at the database level, but also the maintenance of certain inference operations at that level. For example, active databases can store and query also derived data; these data are obtained by the execution of rules that are triggered by external events, such as the insertion of patient related data [93]. Furthermore, integrity constraints based on temporal reasoning [57] could be evaluated at the database level, for example to validate data during their acquisition.
- **Management of different temporal dimensions of clinical data**. In both the artificial-intelligence and the database research areas, as these were applied to time-oriented systems in medicine, typically only the concept of valid time (i.e., a historical database) has been considered. Storing and reasoning about also the transaction time might imply certain benefits, such as being able to restore the state of the database that was true when the physician or a decision-support system decided on a particular therapeutic action, an ability that has significance both for explanation and legal purposes. Another temporal dimension of information considered recently is the *decision-time* [1]: the decision time of a therapy, for example, could be different from both the valid time during which the therapy is administered

and from the transaction time, at which the data related to the therapy are inserted into the database. All these temporal dimensions of clinical data have theoretical and practical interest.

- **Merging the functions of temporal reasoning and temporal maintenance**. By combining these two functions within one architecture, sometimes called a *temporal mediator*, a transparent interface is created to a database, a knowledge base, or both. An example of ongoing research is the Tzolkin temporal-mediation module [94], which is being developed within the EON guideline-based–therapy system [48]. The Tzolkin module merges Shahar's temporal-abstraction system, RÉSUMÉ [41], with Das's temporal-maintenance system, Chronus [35], into a unified temporal-mediator server. The Tzolkin server answers complex temporal queries using both the time-oriented patient database and the domain-specific temporal-abstraction knowledge base, but hides the internal division of computational tasks from the user (or the from the calling process). Many questions will still have to be answered, such as how does a temporal mediator decide which computational module to use for what temporal queries, and will provide interesting issues for future research.
- **Handling deterministic versus probabilistic data**, **and absolute versus relative time references**. As pointed out in Section 2.1.3, clinical data often are most naturally expressed in pure relative temporal terms, although most implemented systems use an instance-based, absolute-time framework. In addition, as mentioned in Section 3, clinical data often involve inherent uncertainty, although they include many discrete, deterministic aspects. Ngo and colleagues [78] present one interesting hybrid system, which integrates discrete contexts with probabilistic belief networks to make reasoning more efficient. Additional models need to be created to combine the advantages of all of these various aspects of clinical data.
- **Provision of standardized, user-friendly temporal-query and temporal-visualization interfaces**. Physicians and other care providers are not database experts and should not be expected to be familiar with the internal workings of either a temporal-reasoning or a temporal-maintenance system or with its theoretical underpinnings. Thus, one challenge is to provide them with easy to use, perhaps even graphic, temporal-query interfaces that enable them to take advantage of the sophisticated architectures that are being built on top of the clinical, time-oriented electronic patient records [12]. Furthermore, many queries might be unnecessary if useful visualization interfaces exist. The semantics of these interfaces (e.g., deciding automatically which abstraction level of the same set of parameters to show and at what temporal granularity) might draw upon the domain-specific knowledge base. An early example was Cousins and Kahn's framework [54] for visualization of time-oriented clinical data. Cousins and Kahn defined a small but powerful set of domain-independent graphic operators with well defined semantics, and a domain-specific representation of reasonable temporal-granularities for a presentation of various entities in the specific clinical domain. More sophisticated interfaces might be built by taking advantage, for instance, of domain-specific temporal-reasoning knowledge [13].
- **Resolution of conflicts between temporal-reasoning and temporal-maintenance systems within hybrid architectures**. Currently, it is common to have temporal-reasoning systems working purely within a short-term, random-access memory, while the temporal-maintenance system stores and retrieves data and abstractions using a long-term storage device such as an external database. As a result, multiple conflicts might arise, especially when systems need to be accessed concurrently by multiple users. One problem is the inherent nonmonotonicity of temporal abstractions, which might be retracted when additional data arrives (whose valid time is either the present or the past). This problem is

solved, for instance, in [41], by the use of a logical truth-maintenance system (TMS). However, integrating a temporal-abstraction system with an external database (as might happen in a temporal-mediator architecture such as mentioned in this section) might create inconsistency problems: the temporal-abstraction system might update its old conclusions as newly-available data arrive; but a standard database system, not having the benefit of the dependency links and the TMS mechanism, will also keep the old, incorrect conclusions. In addition, arrival of new data to the patient database should be reported to the temporal-abstraction module. Thus, we need to investigate whether the short-term, random-access temporal-reasoning fact base and the long-term external database should be *tightly coupled* (each update is reflected immediately in the other database), *loosely coupled* (updates are sent intermittently to the other database) or not coupled at all. Several protocols for connecting and mutually updating the internal and external databases are theoretically possible. The choice among these protocols might depend on the properties of the specific medical domain, and the capabilities of the external database (e.g., object-oriented databases handle links among entities better); adding a transaction time to the patient's electronic record, while keeping the valid time (i.e., using a bitemporal database), would obviously be very helpful. In addition, the deductive capabilities of active databases might provide several advantages, similar to a TMS. In any case, the problem deserves further research.

- **Providing efficient storage protocols for hybrid architectures**. Finally, another issue, closely related to the conflict-resolution problem, is whether some, all, or none of the temporal-reasoning conclusions should be saved in the external, long-term database. Given that many abstractions are only intermediate, and that other abstractions might be changed by data arriving in the future (possibly even data with a past valid-time stamp, or data that exert some influence on the interpretation of the past), it might be advisable not to save any abstractions, due to their logically defeasible nature. However, it is obviously useful, from an efficiency point of view, to cache key conclusions for future use, either to respond to a direct query or to support another temporal-reasoning process. The caching is especially important for saving high-level abstractions, such as "nephrotic syndrome, " that have occurred in the past, are unlikely to change, and are useful for interpreting the present. Such abstractions might be available for querying by other users (including medical decision-support programs), who do not necessarily have access to the temporal-abstraction module or to the domain's full temporal-abstraction knowledge base. One option that might be worth investigating is an episodic use of "temporal checkpoints" beyond which past abstractions are cached, available for querying but not for modification.

Work on each of the new research areas we listed would contribute towards the important goal of integrating temporal data-maintenance and temporal-reasoning systems in medical domains, and thus lead to both a better understanding and to a better solution of important problems in management and reasoning about time-oriented clinical data.

7. Summary

We presented a brief, nonexhaustive overview of research efforts in designing and developing time-oriented systems in medicine. This overview is biased by two foci of emphasis: we focused on methodological and theoretical aspects, and we considered in some detail the tasks of abstraction of time-oriented data and the handling of different temporal-granularity levels.

We started by considering the modeling issues involved in the adoption of a set of basic concepts that enable a description of the time-oriented clinical world in a sound and unambiguous way.

Several suggestions have emerged from general computer-science areas, such as artificial intelligence, or knowledge and data management. Within the medical informatics area, this research effort has progressed from an ad-hoc definition of concepts supporting a particular application to the adoption and the proposal of more generic definitions, supporting different clinical applications.

We then presented research efforts specific to temporal reasoning in medical domains. Temporal reasoning is a component of a wide variety of generic tasks, such as diagnosis (or, in general, abstraction and interpretation), monitoring, projection, forecasting, and planning. These tasks are often interdependent. From a methodological point of view, one general criterion that can be adopted when classifying temporal-reasoning research as it was applied to clinical domains is whether the appraoch taken was deterministic or probabilistic.

We then moved on to the task of management of time-oriented clinical data. We observed that the literature has progressed from the early systems, which were mostly application dependent, to more general approaches, that, even when applied to the solution of real problems in management of time-oriented clinical data, have a more generalizable value and inherent soundness. An important role in this direction is that of research efforts in the general field of temporal databases, where, for example, different kinds of temporal dimensions in databases have been identified, i.e. transaction time, valid time, and user-defined time.

Two commonly recurring and closely related tasks in both the temporal-reasoning and the temporal-maintenance research areas are the temporal-abstraction task and the handling of variable temporal granularity. We examined contributions in these areas according to a framework which distinguishes abstraction granularity, absolute-time granularity, and calendar-date granularity.

Temporal reasoning systems and temporal data-maintenance systems are often independent efforts, even though they usually contribute towards the same goal. For example, time-oriented decision-support systems often do not adopt any kind of a formal temporal data model or a temporal query language to manage stored time-oriented clinical data. The temporal-abstraction task and the management of temporal granularity seem to be a meeting point between research efforts originating in the artificial-intelligence and in the database communities, at least as these efforts have been applied to medical domains. Thus, we concluded the paper by listing several important themes for future research, most of which are relevant also to the community of general computer scientists.

Acknowledgments

This work was partially supported by contributions from the Department of Mathematics and Computer Science, University of Udine, Italy, and by grants LM05708 and LM06245 from the National Library of Medicine, USA, and IRI-9528444 from the National Science Foundation, USA. Computing resources also were provided by the Stanford CAMIS project, funded under Grant No. LM05305 from the National Library of Medicine, USA.

References

[1] Ozsoyoglu G, Snodgrass RT. Temporal and real-time databases: a Survey. IEEE Transactions on Knowledge and Data Engineering 1995; 7(4): 513–532.

[2] Clifford J. Formal semantics and pragmatics for natural language querying. Cambridge tracts in theoretical computer science. Cambridge University Press. Cambridge, 1990.

[3] Snodgrass RT (ed). The TSQL2 temporal query language. Kluwer Academic Publishers, Boston, MA, 1995.

[4] Montanari A, Pernici B. Temporal Reasoning. In [66]: 534–562.

[5] Kohane IS. Temporal reasoning in medical expert systems. Technical Report 389, Laboratory of Computer Science, Massachusetts Institute of technology, Cambridge, MA, 1987.

[6] Keravnou ET, Washbrook J. A Temporal Reasoning Framework Used in the Diagnosis of Skeletal Dysplasias. Artificial Intelligence in Medicine 1990; 2: 239–265.

[7] Wiederhold G, Perreault LE. Clinical Research System. In [74]: 503-534.

[8] Kahn M, Fagan L, Tu S. TQuery: A Context-Sensitive Temporal Query Language. Computers and Biomedical Research 1991; 24: 401–419.

[9] Kahn M, Fagan L, Tu S. Extensions to the Time-Oriented Database Model to Support Temporal Reasoning in Medical Expert Systems. Methods of Information in Medicine 1991; 30: 4-14.

[10] Larizza C, Moglia A, Stefanelli M. M-HTP: A System for Monitoring Heart Transplant Patients. Artificial Intelligence in Medicine 1992; 4: 111-126.

[11] Pinciroli F, Combi C, Pozzi G, Rossi R. MS2/Cardio: Towards a Multi-Service Medical Software for Cardiology. Methods of Information in Medicine 1992; 31: 18-27.

[12] Combi C, Pinciroli F, Musazzi G, Ponti C. Managing and displaying different time granularities of clinical information. In: Ozbolt JG (ed.), 18. annual symposium on computer applications in medical care. Philadelphia, Hanley & Belfus, 1994: 954-958.

[13] Shahar Y, Musen MA. Knowledge-based temporal abstraction in clinical domains. In [16]: 267-298.

[14] Keravnou ET. Temporal diagnostic reasoning based on time-objects. In [16]: 235–266.

[15] Keravnou ET. (guest editor) Special Issue on Temporal Reasoning in Medicine. Artificial Intelligence in Medicine 1991; 3(6).

[16] Keravnou ET. (guest editor) Special Issue on Temporal Reasoning in Medicine. Artificial Intelligence in Medicine 1996; 8(3).

[17] Haimowitz I.J., Kohane, I.S. An epistemology for clinically significant trends. Proceedings of the Tenth National Conference on Artificial Intelligence. Menlo Park, CA, AAAI Press, 1993: 176-181.

[18] Haimowitz IJ, and Kohane IS. Automated trend detection with alternate temporal hypotheses. Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence, San Mateo: Morgan Kaufmann, 1993: 146-151.

[19] Boehlen MH. Temporal database system implementations. SIGMOD Record 1995; 24(4): 53–60.

[20] Vassilakis C, Georgiadis P, Sotiropoulou A. A Comparative Study of Temporal DBMS Architectures. In Wagner RR, Thoma H. (eds), Proc. 7th Int. Workshop on Database and Expert Systems Application. Los Alamitos, CA, IEEE Computer Press, 1996: 153-164.

[21] Chittaro L, Goodwin S, Hamilton H, Montanari A (eds), Third International Workshop on Temporal Representation and Reasoning (TIME '96), Los Alamitos, CA, IEEE Computer Society Press 1996.

[22] Keravnou ET. Engineering time in medical knowledge-based systems through time-axes and time-objects. In [21]: 160–167.

[23] Combi C, Pinciroli F, Pozzi G. Managing time granularity of narrative clinical information: the Temporal Data Model TIME-NESIS. In [21]: 88 - 93.

[24] Shahar Y. Dynamic temporal interpretation contexts for temporal abstraction. In [21]: 64-71.

[25] Shahar Y. A framework for knowledge-based temporal abstraction. Artificial Intelligence 1997; 90 (in press).

[26] Bichindaritz I, Conlon E. Temporal knowledge representation and organization for case-based reasoning. In [21]: 152–159.

[27] Wiederhold G. Databases for Health Care. Lecture Notes in Medical Informatics, Heidelberg, G, Springer-Verlag, 1981.

[28] Long WJ and Russ TA. A control structure for time dependent reasoning, In: Proceedings of the Eighth International Joint Conference on Artificial Intelligence, Los Altos, CA, William Kaufmann, 1983: 230–232.

[29] Pinciroli F, Combi C, Pozzi G. Object-orientated DBMS Techniques for Time-orientated Medical Record, Medical Informatics 1992; 17(4): 231-241.

[30] Combi C, Pinciroli F, Pozzi G. Managing different time granularities of clinical information by an interval-based temporal data model. Methods of information in medicine 1995; 34(5): 458-474.

[31] Dorda W. WAREL: A System for Retrieval of Clinical Data Considering the Course of Diseases. Methods of Information in Medicine 1989; 28: 133-141.

[32] Dorda W. Data-Screening and Retrieval of Medical Data by the System WAREL. Methods of Information in Medicine 1990; 29: 3-11.

[33] Sauquet D, Degoulet P, Lavril M, Aime' F. LIED: A Semantic and Temporal Data Management Language. In Timmers T, Blum BI (eds.), Software Engineering in Medical Informatics. Amsterdam, North-Holland, 1991: 267-277.

[34] Degoulet P, Devries C, Chantalou J, Klinger E, Sauquet D, Zweigenbaum P, Aime' F. LIED: A Temporal Database Management System, Medinfo 1986, North-Holland, 1986: 532-536.

[35] Das AK, Musen MA. A Temporal Query System for Protocol-Directed Decision Support, Methods of Information in Medicine 1994; 33: 358-370.

[36] Das AK, Musen MA. A comparison of the temporal expressiveness of three database query methods. In: Gardner RM (ed.), 19. annual symposium on computer applications in medical care. Philadelphia, Hanley & Belfus, 1995: 331-337.

[37] Dagum P, Galper A., Horvitz E, Seiver A. Uncertain Reasoning and Forecasting. International Journal of Forecasting 1993; 11: 73-87.

[38] Dagum P, Galper A. Time-series prediction using belief network models. International Journal of Human-Computer Studies 1995; 42: 617-632.

[39] Bellazzi R. Drug Delivery Optimization through Bayesian Networks: An application to Erythropoietin Therapy in Uremic Anemia. Computers and Biomedical Research 1993; 26:274-293.

[40] Kahn MG, Marrs KA. Creating temporal abstractions in three clinical information systems. In: Gardner RM, (ed.) 19. annual symposium on computer applications in medical care. Philadelphia, Hanley & Belfus, 1995:392-396.

[41] Shahar Y, Musen MA. RÉSUMÉ: A Temporal-Abstraction System for Patient Monitoring. Computers and Biomedical Research 1993; 26: 255-273. Reprinted in van Bemmel, J.H., and McRay, A.T. (eds) (1994), Yearbook of Medical Informatics 1994: 443-461. Stuttgart: F.K. Schattauer and The International Medical Informatics Association.

[42] Larizza C, Berzuini C, Stefanelli M. A general framework for building patient monitoring systems. In: Barahona P, Stefanelli M, Wyatt J, (eds) 5 Artificial Intelligence in Medicine, LNAI 934, Berlin Heidelberg, Springer, 1995: 91 - 102.

[43] Rucker DW, Maron DJ, Shortliffe EH. Temporal Representation of Clinical Algorithms Using Expert-System and Database Tools. Computers and Biomedical Research 1990; 23: 222 - 239.

[44] Tu SW, Kahn MG, Musen MA, Ferguson JC, Shortliffe EH, Fagan LM.Episodic skeletal plan refinement based on temporal data. Communications of ACM 1989; 32(12): 1439–1455.

[45] Fagan LM, Kunz JC, Feigenbaum EA, Osborn JJ. Extensions to the rule-based formalism for a Monitoring task, in Buchanan BG, Shortliffe EH (eds). Rule-Based Expert Systems: the MYCIN Experiments of the Stanford Heuristic Programming Project. Addison Wesley, Reading MA, 1984; 22: 397 - 423.

[46] Riva A, Bellazzi R. Learning temporal probabilistic causal models from longitudinal data. In [16]: 217-234

[47] Long W. Temporal reasoning for diagnosis in a causal probabilistic knowledge base, [16]:

193-216.

[48] Musen MA, Tu SW, Das AK, Shahar Y. EON: A component-based approach to automation of protocol-directed therapy. Journal of the American Medical Association 1996; 3(6): 367-388.

[49] Jung D, Brennecke R, Kottmeyer M, Hering R, Clas W, Erbel R, Meyer J. Tree-Base, A Query Language for a Time-Oriented Cardiology Data Base System. IEEE Computers in Cardiology 1986; 381-384.

[50] Yeh A. Automatically analyzing a steadily beating ventricle's iterative behavior over time. In [15]: 313 - 323.

[51] Tong DA, Widman LE. Model-Based Interpretation of the ECG: A Methodology for Temporal and Spatial Reasoning. In: Safran C(ed.), 17. annual symposium on computer applications in medical care. New York, McGraw-Hill, 1993: 133-139.

[52] Haimowitz IJ, Kohane IS. Managing temporal worlds for medical trend diagnosis. In [16]: 299–322.

[53] Haddawy P, Helwig JW, Ngo L, Krieger RA. Clinical simulation using context-sensitive temporal probability models. In: Gardner RM (ed.), 19. annual symposium on computer applications in medical care. Philadelphia, Hanley & Belfus, 1995: 203 - 207.

[54] Cousins SB, Kahn MG. The visual display of temporal information. In [15]: 341-357.

[55] Chittaro L, Del Rosso M, Dojat M. Modeling medical reasoning with the event calculus: an application to the management of mechanical ventilation. In: Barahona P, Stefanelli M, Wyatt J, (eds) Artificial Intelligence in Medicine, LNAI 934, Berlin Heidelberg, Springer, 1995: 79 - 90.

[56] Chitarro L, Dojat M. Using a General Theory of Time and Change in Patient Monitoring: Experiment and Evaluation. Computers in Biology and Medicine 1997 (this issue).

[57] Horn W, Miksch S, Egghart G, Popow C, Paky F. Effective Data Validation of High-Frequency Data: Time-Point, Time-Interval, and Trend-Based Methods. Computers in Biology and Medicine 1997 (this issue).

[58] Combi C, Missora L, Pinciroli F. Supporting Temporal Queries on Clinical Relational Databases: the S-WATCH-QL Language. In: Cimino JJ (ed.), 1996 AMIA Annual Fall Symposium. Philadelphia, Hanley & Belfus, 1996: 527 - 531.

[59] Aliferis CF, Cooper GF. A new formalism for temporal modeling in medical decision-support systems. In: Gardner RM (ed.), 19. annual symposium on computer applications in medical care. Philadelphia, Hanley & Belfus, 1995: 213 - 217.

[60] Console L, Torasso P. On the Co-operation between Abductive and Temporal Reasoning in Medical Diagnosis. In [15]: 291-311.

[61] Dolin RH. Expressiveness and query complexity in an electronic health record data model. In: Cimino JJ, (ed.), 1996 AMIA Annual Fall Symposium. Philadelphia, Hanley & Belfus, 1996: 522-526.

[62] Dolin RH. Modeling the temporal complexities of symptoms. JAMIA 1995; 2: 323-331.

[63] Allen J. Maintaining Knowledge about Temporal Intervals, ACM Comm. 1983; 26(11): 832-

843.

[64] Allen J. Towards a General Theory of Action and Time, Artificial Intelligence 1984; 23: 123-154.

[65] Snodgrass R. A Taxonomy of Time in Databases. ACM/SIGMOD 1985: 236-246.

[66] Tansel AU, Clifford J, Gadia S, Jajodia S, Segev A, Snodgrass R. Temporal Databases: Theory, Design, and Implementation. The Benjamin/Cummings Publishing Company, Redwood City, CA, 1993.

[67] McDermott DV, A temporal logic for reasoning about processes and plans. Cognitive Science 1982; 6(2): 101–155.

[68] Shoham Y, Temporal logics in AI: Semantical and ontological considerations. Artificial Intelligence 1987; 33(1): 89-104.

[69] Goralwalla IA, Ozsu MT, and Szafron D. Modeling Medical Trials in Pharmacoeconomics Using a Temporal Object Model. Computers in Biology and Medicine 1997 (this issue).

[70] Combi C, Pinciroli F, Pozzi G. Temporal Clinical Data Modeling and Implementation for PTCA Patients in an OODBMS Environment, Proceedings of Computers in Cardiology. Los Alamitos, IEEE Computer Society Press, 1994: 505-508.

[71] Dechter R, Meiri I, Pearl J. Temporal Constraint Networks, Artificial Intelligence 1991; 49: 61-95.

[72] Maiocchi R, Pernici B, Barbic F. Automatic Deduction of Temporal Information. ACM Transactions on Database Systems 1992; 17(4): 647-698.

[73] Combi C, Pinciroli F, Cavallaro M, Cucchi G. Querying temporal clinical databases with different time granularities: the GCH-OSQL language. In: Gardner RM (ed.), 19. annual symposium on computer applications in medical care. Philadelphia, Hanley & Belfus, 1995: 326-330.

[74] Shortliffe EH, Perreault LE(eds.), Wiederhold G, Fagan LM (ass. eds.). Medical Informatics: Computer Application in Health Care. Addison Wesley, Reading, MA, 1990.

[75] Bortolan G., Cavaggion C., DeganiR., Bressan M., Marinato P. The Role of Patient History in a Decision Support System, IEEE Computers in Cardiology 1989, pp. 357-360.

[76] Chandrasekaran B. Generic tasks in knowledge-based reasoning: High-level building blocks for expert system design. IEEE Expert 1986; (1):23–30.

[77] Aliferis CF, Cooper GF, Pollack ME, Buchanan BG, Wagner MM. Representing and Developing Temporally Abstracted Knowledge as Means towards Facilitating Time Modeling in Medical Decision-Support Systems. Computers in Biology and Medicine 1997 (this issue).

[78] Ngo L, Haddawy P, Krieger RA, Helwig J. Efficient Temporal Probabilistic Reasoning Via Context-Sensitive Model Construction. Computers in Biology and Medicine 1997 (this issue).

[79] Kowalski R, Sergot M. A logic-based calculus of events. New Generation Computing 1986; 4: 67-95.

[80] van Beek P. Temporal query processing with indefinite information. In [15]: 325 - 339.

[81] Console L, Torasso P. Temporal Constraint Satisfaction in Causal Models. Information Sciences 1993; 69: 1-32.

[82] Blum RL. Discovery and representationof causal relationships from a large time-oriented clinical database. Lecture Notes in Medical Informatics. Springer Verlag, New York.

[83] Ferrari R. Temporal Extensions to the Iris Model in a Medical Domain. Technical Report, HPL-91-132, Hewlett-Packard Software and Systems Laboratory, 1991.

[84] Kouramajian V, Fowler J. Modeling past, current, and future time in medical databases. In: Ozbolt JG(ed.), 18. annual symposium on computer applications in medical care. Philadelphia, Hanley & Belfus, 1994: 315-319.

[85] Clifford J, Tuzhilin A. (eds) Recent Advances in Temporal Databases. Heidelberg, Springer Verlag, 1995.

[86] Snodgrass R, Ahn I. Temporal Databases. IEEE Computer 1986; 19(9): 35-42.

[87] Kahn M, Fagan L, Tu S. Combining Physiologic Models and Symbolic Methods to Intepret Time-Varying Patient Data. Methods of Information in Medicine 1991; 30: 167-178.

[88] Wiederhold G, Jajodia S, Litwin W. Dealing with Granularity of Time in Temporal Databases. Proceedings of the Third Nordic Conference on Advanced Information Systems Engineering, Trondheim, Norway, 1991: 124 - 140.

[89] Wuu GTJ, Dayal U. A Uniform Model for Temporal and Versioned Object-oriented Databases. In [66]: 230–247.

[90] Elmasri R, Wuu GTJ, Kouramajian V. A Temporal Model and Query Language for EER Databases. In [66]: 212 - 229.

[91] Elmasri R., Navathe S. Fundamentals of Database Systems, 2nd Edition. The Benjamin/Cummings Publishing Company, Redwood City, CA, 1994.

[92] Sujansky W, Altman R, Towards a standard query model for sharing decision-support applications. In: Ozbolt JG(ed.), 18. annual symposium on computer applications in medical care. Philadelphia, Hanley & Belfus, 1994: 325-331.

[93] Chomicki J, Toman D. Implementing temporal integrity constraints using an active database. IEEE Transactions on Knowledge and Data Engineering 1995; 7(4): 566-582.

[94] Das AK, Shahar Y, Tu SW, Musen MA. A temporal-abstraction mediator for protocol-based decision support. In: Ozbolt JG (ed.), 18. annual symposium on computer applications in medical care. Philadelphia, Hanley & Belfus 1994: 320-324.