Temporal Reasoning for Decision Support in Medicine

Juan Carlos Augusto

School of Computing and Mathematics
University of Ulster at Jordanstown
BT37 0QB Newtownabbey, Co. Antrim, UK
Phone: +44 (0)28 9036 612
Fax: +44 (0)28 9036 6068
jc.augusto@ulster.ac.uk

Abstract

Time-related concepts handling is essential in medicine. During diagnosis it can make a substantial difference to know the temporal order in which some symptoms occurred or for how long they lasted. During prognosis the potential evolutions of a disease is conceived as a description of events unfolding in time. In therapy planning the different steps of treatments must be applied in a precise order, with a given frequency and for a certain span of time in order to be effective.

This article offers a survey on the use of temporal reasoning for decision support-related tasks in medicine. Key areas are highlighted and used to organize the latest contributions. The survey of previous research is followed by an analysis of what can still be improved and what is needed to make the next generation of decision support systems for medicine more effective.

Key words: temporal reasoning, decision support for medicine, diagnosis, prognosis, therapy planning

Contents

1 Introduction

Reasoning about the possible orders of activities or about deadlines is certainly an everyday matter in our lives. Accordingly, research in Computer Science about time-related issues has been growing steadily for the last two decades. For example, the last decade the TIME [1–4] series has brought together researchers with different perspectives, aims and backgrounds conducting research on the fundamentals of dynamic systems and its applications.

In this article we are especially interested in temporal reasoning as related with decision support-related tasks in medicine. Many researchers, especially from the areas of *Artificial Intelligence* (AI) [5] and *Medical Informatics* (MI) [6], have contributed to this area in the last decades and a good amount of research has been published. We are not aiming to be exhaustive. Instead we address key areas and point to the main publications that could provide a good starting point for the reader. When several options are available or when several related works were produced by the same author we cite their latest publication, from which the interested reader will be referred to previous related publications.

The remaining sections are organized as follows. First (section 2) we introduce some basic concepts in relation with temporal reasoning. Then (section 3) we explain why time is important in Medicine. Next (section 4) we provide a non-exhaustive but comprehensive description of the work that have been done in those key topics identified in the previous section. Time granularity is considered separately as clearly is a subject does not belong to any particular aspect of the decision process but also can have major relevance in all them. We conclude this article (sections 5 and 6) by providing an analysis of what still can be improved and what is needed to make the next generation of decision support systems for medicine more flexible and effective.

2 Time Representation and Use in Artificial Intelligence: Some Basic Concepts

The formalization of temporal reasoning proved to be a great task as the literature in logic, philosophy, linguistics, AI and Computer Science testifies. There are then many perspectives and many options to consider at the stage when one is trying to define what the basic layer of a temporal system should be. We are not aiming here at covering all the aspects and will instead focus very shortly in a subset of them from the perspective of medicine. The reader is addressed to [7] for a lengthier description of the basic concepts as well as a description of the main trends of temporal reasoning in AI. We will first consider some concepts related to the ways we can define a structure for time-related references. Then we describe some concepts that are usually associated to time references. The final step is to relate the previous discussions about time representations and temporal concepts to characterize one of the key notions for any dynamic system, e.g. change.

2.1 Defining the underlying structure for temporal references

Characterizing the temporal scenario to be assumed strongly influences the theory to be proposed and this, in turn, the tools to be produced. Recently Computer Science and AI researchers (building on previous work from philosophers, linguistics and logicians) have considered temporal notions as part of the program to characterize the notion of *intelligence* and as a first step to different goals. See for example [7–12].

One aspect to be considered is what basic “shape” and set of properties the underlying temporal structure is assumed to have. Time is usually conceived as a line where temporal references can be aligned. This conception of time proved to be popular since the Newtonian physics paradigm was adopted and provides the simplest conception and way to arrange temporal references. See [13] for example. Another popular alternative in Computer Science (CS) and artificial intelligence (AI) is a future-branching structure representing past as linear and the present as a distinguished point where the future opens as a bunch of possibilities. The adoption of a future-branching structure could be motivated in several ways. Usually it is the possibility of representing the capability of an intelligent entity to choose between alternatives or a way to provide hypothetical reasoning which is behind its adoption. [14] A past-branching structure could provide a good framework for abductive reasoning and other related tasks as it provides natural representation for the different possible ways of arriving to a present state of affairs. Circular time could be useful when analyzing repetitive patterns, cyclic processes. See [15] for a philosophical background and [16] for a computational perspective. In a distributed system (like a multiagent system or a distributed database) it could also be useful or necessary to consider that each or some of the participating nodes have their own account of time organized in a specific way. For example such a system can be conceived as an arrangement of parallel lines. In this scenario, coordinating mechanisms and specific protocols should be provided in order to reach agreement on the operations between the interacting parts of the system.

Some tasks may demand consideration of time up to a point. For example, when analyzing symptoms that lead to a particular disease, facts that happened beyond some day will not be of interest. In this context time can be conceived as having a final moment. Reciprocally, when considering a treatment there may only be an interest on facts occurring beyond a specific date. There are other conceptions of time where no specific limits are imposed or to the contrary that conceives time as limited structure with a beginning and end [17].

Time can be also considered from a “topological” perspective, e.g. discrete, dense or continuous. This led to the so called *topological time* because temporal structures could be analyzed under the light of topology as known in mathematics. Sometimes it suffices to define the time references as a discrete succession, e.g. isomorphic to $\mathbb{N}$ or $\mathbb{Z}$. This characterization is particularly relevant in relation with the notion of time granularity, a topic we will consider in more detail later on. For example, tasks like clinical management, diagnosis, prognosis, and treatment can well benefit of a calendric view where references to years, months, days and hours can be represented. However a realistic representation and modelling of subtle physical, biological or chemical processes may demand a dense or continuous conception of time, e.g. isomorphic to $\mathbb{Q}$ or $\mathbb{R}$. However, these steps could not be given without a price to pay. There exists for example the so-called *dividing instant problem* [18] which warns us about some difficulties in continuous change representation.

The problem of deciding which kind of reference must be considered the basic one has been subject to intense debate. Literature about the philosophy of time provides us with several articles from people sustaining an instant-based view of time [19] while others support a period-based approach ([20,21]). Names vary with authors but usually *instants* and *time points* are used to refer to punctual occurrences while *periods*, and *intervals* are used to talk about durative temporal references. Recently some proposals have explored the benefits to allow both kinds of references in the same level of importance ([22–26]). See [18] for an analysis of the three alternatives, i.e. to consider instants, periods or both. Both kind of references are certainly useful in medicine where we need to refer to punctual as well as to durative events and happenings of interest. For example, some symptoms are described as occurring at a particular day, like “The symptoms started on 4th of July” while others have a duration associated like in “He already had fever for three days”. The first type of references are sometimes called *anchored* while the later are called *unanchored*.

Usually intervals are assumed to be periods with known beginning and end. This assumption simplifies matters for computation but it could be an unrealistic assumption to assume that knowledge will be always available. Another common assumption in the literature for temporal reasoning in AI is that they are *convex*, that is uninterrupted spans. It is interesting to see that both kinds of references, punctual and durative, can be defined in terms of each other. For example, periods could be seen as sets of instants or the duration denoted by two instants acting as beginning and ending points. Also instants could be defined as the meeting point between two periods. See [11] for a more detailed account of this complementary conceptions of time.

In the medical domain is also very useful to provide some way to handle what some researchers call *semi-intervals*, [27] i.e. intervals for which either the beginning or the ending is not directly known e.g. “Started yesterday evening and stopped at some point during the night”. It is also important to provide some account for repetitive processes, e.g. “It has headache each time that goes to music class” and frequencies of occurrence, e.g. “It has been taking this medicine three times each day” or “It has been coughing frequently during the night”. Rich calendric references should be handled, like seasons, in order to discover potential causes of disease, e.g. allergies.

It is useful to bear in mind that the above considered set of possibilities for defining different aspects of a temporal structure are independent from each other. For example, the decision if the structure is linear or branching does not rule out considering if it is bounded or not neither to consider if it is discrete, dense or continuous. So there are plenty of choices at design time. Each one will have an effect in terms of the balance between expressiveness and the computational complexity needed to handle it.

2.2 Time-Related Concepts

Another important issue in all temporal theory is to decide what sort of information is subject to change or, in another way, what kind of concepts are considered in the theory beyond time itself. While the information to be associated with temporal concepts can vary from scenario to scenario, some concepts appear repeatedly when we examine temporal reasoning-related literature.

Because temporal reasoning involves solving problems in a changing world, there is a need to represent what *properties* the objects of that world can have or do have at each meaningful temporal reference. These characterize aspects like size, weights, temperature measurements and other distinctive features of each object considered in the intended scenarios. The set of objects and their properties define a *state* of the world. A given state of the world is changed by the occurrence of *events*. These are strongly tied to the notion of time because it is natural to think about time as a mechanism around which we organize reasoning about change in the world. *Actions* are identified with the agents' capabilities of interaction with the world. They are considered event-producers but events could be the by-product of just other event(s) like machine failures, power cuts. Usually the representation of causal relationships will be closely tied to the level of conceptual granularity considered.

It is usual to assume that properties are *homogeneous* which means that if a property holds in an interval then it holds in each part of the interval. For example, if a patient was diagnosed to have been under anesthetic effects from 11AM until 12AM s/he has been on that state at every minute and second between 11AM and 12AM. That is, the property can be transferred “downwards” to smaller time references. Instead, events are assumed not to be homogeneous. If an event, lets say e occurs during an interval it is supposed not to occur as such in any part of that interval. It is possible that another instance of the same type of event occurs but not exactly e, which is unique. For example, if a patient had surgery from 11AM until 12AM that surgery conducted at that place for that patient is a unique event and although s/he was in surgery from 11:45AM until 11:46AM that surgical intervention as a whole event did not lasted from 11:45AM until 11:46AM but for a complete hour.

A sketch of event classification was brought by Allen and Ferguson [28] where they grouped occurrences depending on how predictable they are. They classified events into *triggered*, *definite* and *spontaneous*. The first group identifies events provoked by the system being represented and its consequences are supposed to be known as in “When temperature raised to a critical level the alarm was triggered”. Definite events are not provoked by the system but are known by it and could be predicted like in “Dr. Harrison arrives everyday at 9:00 AM”. Spontaneous events are unexpected and the system could not predict where and when they could happen, e.g. “A power cut started at 7PM”. Although this is not the only possible classification, grouping events into categories may be useful to prepare the system to react appropriately when critical events occurred.

Actions are a key concept in the formalization of an autonomous system. They provide a way to formalize how a system could interact with its environment. As was said above, actions typically produce events and in that way play a key role in a dynamic environment. They are usually attached to *agents* in a broad sense, e.g. persons, robots, machines and other autonomous or semi-autonomous devices. The reader is invited to see [29] for further analysis on the many aspects involved in the analysis of actions.

Another technical term that usually appears in the bibliography is that of *processes* which denotes repetition or regularly sustained activity as in “the patient has been taking a medicine for the last week”, “the patient has been doing exercise the last year” and “the patient has been on dialysis for the last six months”. This does not mean that the activity has been done continuously throughout the referred period of time, for example for the first case does not mean necessarily that the patient has been taking the medicine each minute of that week. Some authors describe them as a “state of change” as a way to differentiate them from events which could be defined as a “change of state” [11]. If they can be considered as primitive objects in a theory of change or not is still a matter of debate. Allen [13] assumed they must be considered from the very beginning. He characterized processes as those occurring at least in a part of an interval but not necessarily in the entire interval. These assumptions are criticized in [24]. In general there is no consensus about what sort of concepts this word really involves and some researchers prefer not to include them at the same level as the aforementioned notions [18].

There has been many debates about different conceptions of time modelling and temporal knowledge representation in the last decades. Here we considered some basic notions and highlighted that there are several alternative views. See [24,11,30] for some alternatives on how all these concepts can be related with each other. As a summary we can sketch a picture of how these concepts interact with each other and how they can be put together to provide an explanation of the evolution of a dynamic world. Given some conditions, some actions can be applied by agents and this will produce events to occur which in turn will change some properties of the objects affected. This change in the properties of objects will lead to the identification of state changes. This change in the state of the world will eventually allow a possible different set of allowed actions to be applied and so on. Repetitions of this loop may lead to the formation of higher level entities like processes.

2.3 Modeling Change

We can distinguish two main ways to talk about change and to make temporal references. One way is by using absolute time references like "Symptoms started on January 1st, 2003". Another common way of making time-related constructions do not make explicit mention of calendric references like in "surgery is advised as soon as possible", "two tablets have to be taken after each meal until further notice" and "the cream has to be applied over the area until the irritation has completely disappeared". In these last examples, time is mentioned as constructions relative to the notion of *present*. This is not to say that each kind of temporal references cannot be recreated in the other approach but instead that they are favored in each case.

The first case is closer to what is called the Newtonian paradigm, favored also in the past by famous philosophers and logicians like B. Russell and W. Quine as the preferred way to talk about temporal notions, where time is conceived as an unbounded line of instants. Time is conceived to have existence from the very beginning and the concepts of states and change are derived from it. Classic first order logic, or variations of it, is usually adopted as the tool to handle temporal references of this kind.

The other conception of time relies heavily on the notion of change, which is directly associated to events. They are assumed to exist from the very beginning and temporal concepts are built from this foundation. A new class of temporal logics based on this conception of time where references are relative to the events occurred and their relative order were offered by the middle of this century (see [31] for a first landmark and [9] for a more modern account). Because of its adequacy to represent the kind of temporal notions and references used in linguistics, these logics may be favored for natural language oriented applications.

Historically these conceptions have been considered as competitors originating a debate lasting decades since Quine's and Prior's work. As the reader will see in the following sections it is still unclear if one must be preferred to the other and in a CS context it probably will be the case that each one will have areas of applications where one will bring advantages over the other but both will coexist. See [11] for more details on this debate. We follow this source on the following explanation on how these different views are related to each other.

Consider two time references i_1 and i_2 and two events associated with each one we could say which occurred before or if they were simultaneous (see figure 1).

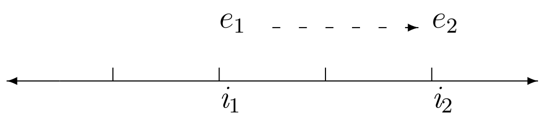

Fig. 1. Deriving event order from time order.

Two states of the world associated with two different instants are said to be different if a property holds in one of them that does not hold in the second one. Given the situation represented in figure 2 we say there was a change between i_1 and i_2 and the world evolved from a state S_{i1} to another S_{i2}. Because there was a change between i_1 and i_2 we could suppose that an event occurred.

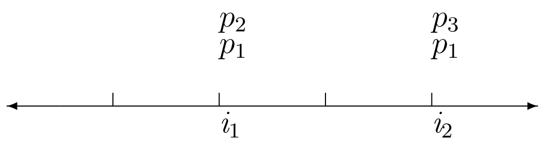

Fig. 2. Discovering a change.

In this proposal, change is not directly encoded but inferred as a by-product of time. Temporal references are usually numeric and it is easy to represent quantitative temporal relations but also qualitative relations are possible, e.g. if an instant is before another or if two periods have something in common. A state of the world is conceived as a blackboard with an infinite set of labels to which information is attached. In this sense information could be said to be *referentially neutral*, in the sense it is allowed to be used without care of its relative position with the "present".

In the Leibnizian paradigm, change is considered the fundamental concept and the concept of time must be built from it. In this case from an ordered succession of events the precedence of two moments of time can be inferred (see figure 3).

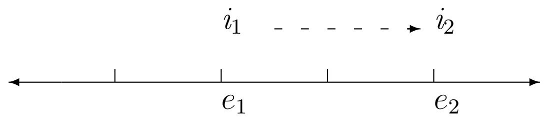

Fig. 3. Deriving the order of time references from events occurrences.

The usual relations to be used in this framework are *simultaneity* and *precedence* although others can be also useful such as “to be between two given events”. In this case also occurrence is usually considered instantaneous and the structure is assumed to be unbounded and continuous. Durative events could also be considered ([13], [32], [26]). Change is represented directly in this approach and not as a by-product of comparing the knowledge in two states of the world.

3 Time in Medicine

Time is so common in our daily thinking and our activities that become of fundamental importance in our lives. So basic that sometimes all this time-based reasoning pass unnoticed. The same can be observed when we consider decision making activities in medicine (see for example [33] for a quick account). It is enough to examine step by step any of the processes carried out during medical diagnosis, prognosis and treatment to quickly discover how some decisions are determined by previous decisions and facts and how important is to project possible scenarios and related treatments.

During the last decade the role of time in medicine has been brought to the surface and its relevance has been reflected directly in the intense activity of the area. Many articles addressing various topics are now available and in the last years we also witnessed the publication of four major special issues dedicated to the topic ([34–37]) and the dedication of a special session in major conferences of the field (e.g. in AIME the European Conference of Artificial Intelligence in Medicine). Both are signals of the recognition that this area is gaining within the field.

A long list of activities were early identified [38] as key initial steps on the process to provide explicit temporal awareness for systems related to medicine: mapping occurrences across temporal contexts, determining bounds for absolute occurrences, persistence derivation, inconsistency detection and clipping of uncertainty, deriving new occurrences from other occurrences, deriving temporal relations between occurrences, deriving the truth status of queried occurrences and deriving the state of the world at a particular time.

As already acknowledged [39], computer-assisted medicine feeds from the multidisciplinary work of many Computer Science-related disciplines. The articles published in [36] reflected that emergent complexity within the particular time-related tasks. An interesting summary was provided in [40] restricted to methodological and theoretical choices. The article highlights the importance of two main research directions: *temporal reasoning* (associated to, for example, temporal abstraction, time-oriented decision support, forecasting and data validation) and *temporal data maintenance* (as related to storage and retrieval of time-related data). A rich agenda was devised: adoption of advanced data models, maintenance of clinical raw data and abstractions, management of different temporal dimensions of clinical data, merging the functions of temporal reasoning and temporal maintenance, handling deterministic versus probabilistic data, handling absolute versus relative time references, provision of standardized and user-friendly temporal-query and temporal-visualization interfaces, resolution of conflicts between temporal-reasoning and temporal-maintenance systems within hybrid architectures and providing efficient storage protocols for hybrid architectures. This is consistent with [41] where two main trends were identified: temporal reasoning (traditionally linked with the AI community) and temporal data abstraction (more identified with the Database community). The interaction of these two areas is highlighted as an important activity for the research agenda as it has the crucial role of bringing together knowledge representation and reasoning. The special issue of [37] presented a range of increasingly sophisticated systems dealing with various subtle topics addressing this integration issues. Still, as the editors made explicit, those were only the first steps to tackle the formidable task of providing flexible and rich integration between time representation and use in medicine.

We share that view and the purpose of this article is then to build on top of that previous work and to unravel further challenges which should be addressed in order to obtain the next generation of intelligent assistants for medicine. That will demand further multidisciplinary interaction between different groups whose knowledge on time-oriented systems may find in the area of medicine a challenging but highly rewarding field of application. Researchers and practitioners from MI, temporal reasoning in AI, temporal databases, active databases, real-time databases, visualization of dynamic systems and real-time systems should have some knowledge and experience to share in this fertile area. In the following sections we highlight some critical activities in medicine, the automated support available up to now and some of the areas demanding further improving or exploration.

4 Decision Support for Diagnosis, Prognosis and Treatment

In this section we provide a survey of the latest developments and trends in some of the main areas explored by researchers from AI and MI concerning time-aware proposals for decision support. We start by mentioning some issues about time-oriented data abstraction, representation, storage and retrieval. Then we summarize the main and more recent contributions that have been made in connection with the role of time in diagnosis and therapy planning.

4.1 Representing and retrieving time-oriented data

There is a real need to store time dependent data in medicine applications and of course, the purpose is to store this information with the ultimate goal of being able to recover that information later in a useful context. The literature in CS reporting contributions to this problem from a general perspective is rich (see for example, [42,43,1–4]).

There have been other contributions to the topic from within the MI community as early of as in the '70s with the *time-oriented paradigm* which consisted on adding a time-stamp to the typical parameters of patient identification and clinical parameters.

More recently, [44] reported on an attempt to provide services tailored to protocol-directed decision support in medicine which was better prepared to cope with different granularities and with interval-based temporal references in the context of HIV treatment. The proposal was to extend SQL language in three main directions: a) allowing time-stamps to store the temporal dimension of both instant- and interval-based data b) providing a set of operations on time points and intervals to manipulate time-stamped data c) modifying the relational query language SQL so that its underlying algebra supports the specified operations on time-stamps in relational tables.

For the case of modelling medical trials in pharmaco-economics [45] proposed a more radical departure from the relational paradigm for databases including an object-oriented database system where time is considered as one of the basic components. Some issues addressed in this proposal are branching time, granularity and temporal indeterminacy. This proposal was related to the *Tigukat* object model, [46] which aimed to accommodate multiple applications that have different type semantics requiring various notions of time.

The *Tzolkin* system [47], aimed at integrating a general method for temporal-data maintenance with a general method for temporal reasoning capable to allow: (1) the expression of abstract, time-oriented queries, (2) the retrieval of data that satisfy a given set of time-oriented data-selection criteria, and (3) presentation of the retrieved data at the appropriate level of abstraction. This proposal aimed to increase reusability and for that the authors proposed a so called *temporal-database mediator*. A key issue on achieving a higher degree of reusability was the distinction between task-specific knowledge (how to solve a generic temporal reasoning problem) and domain-specific knowledge (e.g. how to infer a particular time-sensitive situation in medicine). The underlying representation for data is the Historical Database model [48] and other related systems connected with its functionality are *Chronus* [49] and *Résumé* [50].

The capability of providing a high level view of that data by means of abstractions is a valuable step in recovering and processing data depicting the evolution of a patient. By abstraction we refer to summarization of patient data. Clinicians and care providers in general need to know that some symptoms were recurrent with some particular temporal pattern in some specific context in order to diagnose correctly instead of having the rough data that, after possibly long consideration, would lead to the discovery of such condition. The reader may found in the literature quite a few terms related to this activity, e.g. *Data Analysis*, *Temporal Abstraction* or *Data Interpretation*. Also more general terms denoting areas of Computer Science related to discovery and learning, e.g. data mining, may be included as well. Many systems were proposed to provide different levels of abstraction capabilities for time related information. They extended the first attempts for providing time-aware databases by providing not just retrieval capabilities but also temporal reasoning and other processes that allowed to gather several pieces of data in the form of trends and high level descriptions of complex conditions. One of the most comprehensive proposals perhaps being the *Knowledge-Based Temporal-Abstraction* method [50] and much of the exposition that follows is based on the research related to this project. Usually the process of time-related data abstraction is divided in the following stages. First the raw data is filtered, then some abstract episodes are identified from the data available and finally by interpreting these episodes at higher level some of the patient's original data regarding some particular aspect of her/his evolution can be summarized as synthetic knowledge. The importance of the task and the problems likely to be faced in the attempt were clearly summarized in [51]:

> *"Intelligent Data Analysis (IDA) refers to all methods that are devoted to support the transformation of data into information exploiting the knowledge available on the domain. In bio-medical data analysis IDA may play a crucial role, since the value of each datum may be higher than in other contexts: experiments can be costly, due to personnel and instrumentation involvement and/or to the patients discomfort; the data set can be small, reporting non-reproducible situations. Nevertheless, the data may also be affected by several sources of uncertainty, from measurements errors to missing data, or from data coding errors to information buried into textual reports."*

One system that is focused on this kind of service is Résumé [52] where several subtasks are considered:

(1) *temporal-context restriction*: different contexts require different assumptions, different information and can guide significantly the inferential procedure, e.g. reasoning about intensive care demands different level of assurance than in a routinely consultation visit.
(2) *vertical temporal inference*: obtaining higher level information about an episode based on more specific data obtained from the patient. For example, from separate vital signs sensing can be inferred the patient is in a particular stage of a disease.
(3) *horizontal temporal inference*: provides sound simplification of information referring to the same fact in a complementary way. For example, coalescing two consecutive measurements readings or two consistent pieces of information coming from different sources.
(4) *temporal interpolation*: bridges the gap between non-contiguous, but related, pieces of information like two registered periods of the same illness that commonsense will dictate are in fact the extremes of a continuum process. This task is more involved in the sense that demands more hypothetical and inferential work for each step.
(5) *temporal-pattern matching*: allows the detection of a higher level condition based on other data previously discovered or inferred.

Temporal interpolation is a complex subtask by which a new temporal instance is obtained from other two, obtaining for example an interval-based temporal reference from two point-based ones. The procedure can also be applied to interpolate a point and an interval or two intervals. In this two later cases a discrete notion of compatibility is associated, ranking the operation and having then an effect on the kind of inferences allowed over that abstraction.

The possibility of association is a very context-dependent one as for example, two weeks of pregnancy separated by two years cannot be joined but two weeks within the same month are compatible and, given some other constraints like age, may be joined. Certainly contexts are also given relevance in Résumé as explained in [53]. In fact in that system an ontology of *interpretation contexts* is considered as a hierarchical arrangement where each instant of a context can be associated to an interval where it is active. Many aspects of the process being evaluated can change the context of reasoning, say by turning a previous context into a more specific one. For example, in an intensive care unit, a vital sign approaching a critical threshold can trigger a special protocol of operation. These contexts are not considered necessarily as disjunctive, they can be also combined. Contexts not are only useful for more accurate and realistic inference but also allows consideration of a wide range of situations as the intervals over which they have an effect can relate with each other in, for example, all the different ways two intervals can in Hamblin-Allen's sense

([20], [13]). Although illustrations of *Résumé* are usually made by referring to its applications on the diabetes therapy domain, the proposal is presented as a general problem-solving technique.

There are plenty of other interesting issues to consider and contexts that may turn the task of intelligent temporal data analysis difficult. Due to space limitations we just name a few of them below. Applications of these techniques considering both historical and real-time data are considered in [54] for general domains and illustrated over data related to central venous pressure. *RASTA* (RASTA: A System for Temporal Abstraction) [55] enhances *Résumé* with distributed capabilities for problem solving in order to scale-up the system to more complex scenarios. Sometimes data in clinical domains can be vague and imprecise and that includes the time were some critical events happened, e.g. symptoms occurrence. Specially tailored techniques like those presented in [56] to approximate data by using fuzzy categories and applied to anaesthesia monitoring are then meaningful. Some abstraction techniques are better tailored towards a specific area in Medicine. For example, in [57] an application of intelligent data analysis is reported over the Diabetes Mellitus domain. The reader may find in [58] a description of other techniques and applications that however rather out of date now, may offer a good survey of the basic techniques, aims and potential benefits from the area. A more up to date and very comprehensive account on the field is given in [59].

Early developments soon recognized the value of being able to retrieve time-related data and to present the results in a more comprehensible way. For example, in [60] a system based on a temporal structure called *time-line* and a set of related operations to recover and visualize information from different perspectives. At that stage was also recognized the importance of having a flexible system to handle temporal granularity issues, as different events usually have different levels of granularity where they most usually occurs. Some of them may correspond with calendar units, some may not.

It is fundamental for all the basic steps of diagnosis, prognosis and therapy planning to be able to identify the symptoms of a disease and the chronology associated with them.

4.2 Diagnosis

Diagnosis can be defined as *the determination of the nature of a case of disease* and, as pointed out in [61], there is a rich range of temporal references in the description of symptoms which may be presented associated to explicit time references (e.g. last Tuesday) and sometimes as a collection of inter-related events without explicit mention to time (e.g. “first I felt sick and after that I have severe headache and rush in my stomach"). Usually the degree of uncertainty can be high as precisely when or in which order some of them occurred, e.g. at some moment during the night. Sometimes time references, e.g. "seasons" or "weekends", may not be very frequent in the technical literature of the associated technologies, e.g. databases. All these features posses interesting challenges for the research in AI and MI.

Diagnosis has been also an active area in artificial intelligence where the interest were focused on finding explanation for behavior anomalies in different artifacts. Some landmarks coming from that effort are [62], [63] and [64] while other early contributions which are more relevant to temporal reasoning in medicine are [65] and [66].

A good concise account of model-based proposals for diagnosis in AI and medicine is offered in [67]. There, the importance of considering the evolution of some patient-related indicators over time is recognized as one of the important aspects of diagnosis that realistic systems should address.

TrendDx [68] is a trend diagnosis system that has special mechanisms addressing the computational complexity arising from the possible worlds consideration made by the system. The number of possible worlds can grow polynomially and to handle the combinatorial explosion two techniques are used: beam search based on regression scores, and temporal granularity in the trend template definitions. The system will match trend templates to trend-related observations and using a temporal consistency criteria relativized by one of the possible contexts of interpretation will determine if the observations matches a particular generic trend. A component of uncertainty is considered on the temporal occurrence so a relative accuracy criterion is used to discriminate between competing trends.

A proposal connected with the model-based tradition but focused in medical diagnosis was reported in [69]. There an interval-based algebra is used to model possible qualitative relationships of symptoms for Hepatitis B. First-Order logic with special dedicated predicates for interval-based time references is used to model the causal relations of the described disease. For example for the case of Hepatitis B, four rules are used to link the description of causes with the temporal relationships that should hold for the effects and for the relations between the effects and the causes. Other rules focus on how the possible patterns of occurrence relate to each other. For example if we want to express that all possible forms of manifestation are mutually exclusive. In this approach abnormal observations are explained abductively, and it is required that an explanation is consistent with both normal and abnormal observations. An important feature in this proposal is that they provide abstraction mechanisms that may increase significantly the efficiency of the system depending on the nature of the input data.

In [70] a theory of diagnosis is presented which allows to complement causal associations with time intervals to denote the durations associated with the manifestations observed. This in turns allow this causal associations to be related by a precedence order between them. The causal relationships are represented by a directed acyclic graph, i.e. without recurring events. The temporal distance between manifestations and their durations are represented as functions associated to nodes. Another interesting feature of the proposal is that both necessary and possible causal connections between disorders and manifestations can be represented. The proposal is exemplified with its application to the food-born disease domain.

Déjá Vu [71] provides scenario recognition facilities during a patient's examination so that when a session is in progress, it can be related to the expected structure of the process. This capability allows anticipating events and detecting possible deviations so that the appropriate actions can be taken to prevent any undesirable condition. Sessions are modeled by using *Constraint Satisfaction Networks* [72] and the task is reduced to comparing two networks: the general description of the process, called "a scenario", and the one describing a session, i.e. the proposed particular instance of the general process. Other standard techniques from graph theory are used as the basis for graph comparison. The system is based in a temporal representation of the sessions that may have precise and imprecise temporal information. Recognitions can be done in two modalities: *on the fly* or *a posteriori*. The proposal is illustrated on the authors' experience with a simplified ventilation management unit.

Fuzzy sets theory is used in [73] to model the lack of precision on the temporal occurrence of symptoms, e.g. started around midday, or adjectives used on their description, e.g. intense, in order to match a set of observations developing in time with a disorder model. The proposal illustrated to diagnose intoxication by ingestion of poisonous mushrooms cases. In [74] the use of Bayesian probabilistic networks is used to represent time durations of cause and effects associated with heart diseases diagnosis.

A more recent contribution exemplifies how temporal reasoning can be used for on-line monitoring and detecting trends of behaviour, in this case applied to an intensive care unit [75]. Here a system is proposed as an intelligent alarm system. Evolution trends of the data being monitored, e.g. blood pressure, are classified into nine possible temporal shapes: "steady", "increasing", "decreasing", "positive step", negative step", "positive negative step + slope", "negative step + slope", "concave transient", and "convex transient". These are in turn broadly grouped into three major cases: "steady", "decreasing" or "increasing", in a spirit close to Résumé's trends detection. According to the trends detected and the type of information that is being monitored the system will provide real-time detection of potentially dangerous developments.

In [76] a research project that explores the definition of methods and tools for the assessment of clinical performance of a hemodialysis service is summarized. One of the main time-related tasks developed in the system is detection of temporal patterns. Basic and complex temporal abstractions are considered. By *basic* temporal abstractions it is meant numeric or symbolic uni-dimensional time series while *complex* temporal abstractions refers to specific temporal relationships between basic temporal abstractions, in Hamblin's[20]-Allen's[13] sense.

[77] describes *Idan* a distributed temporal-abstraction "mediator" for medical databases which can answer abstract, time-oriented queries by adequately handling the queries to the various key modules in a distributed system. An added feature to previous similar approaches like *RASTA* and *Chronus* is that Idan is capable to handle temporal constraints in a uniform way between the system level and the interface level.

In [78] a domain-based temporal abstraction mechanism relying on a simple qualitative and heuristic approach to handle temporal indeterminacy on period bounds. Hamblin's[20]-Allen's[13] interval relationships are considered when some information is missing, in the sense of Freksa's [27] work, and according to the particular combination of "semi-intervals" being considered, concepts like uninterrupted continuation, possible disruption, possible simultaneousness are inferred. The algorithms are strongly influenced by heuristic knowledge obtained from the domain application in order to simplify the potential cases to be considered.

4.3 Prognosis

By *prognosis* we refer to those activities related to the forecasting of the probable evolution and outcome of a disease. Clearly the benefit obtained is that by this mean it may be possible to anticipate important events and hence to avoid undesirable situations as well as increasing the chances of success of different health care-related tasks. For example, if a patient is being treated we can use the information obtained at this stage to ensure that the drugs, equipment and personnel s/he needs will be available at the time they are required. It has not been until relatively recent years that research on this subject started to emerge as a special area of consideration. Prognosis-oriented activities can be inter-related with either the diagnosis or the therapy stages. However the difficulties offered by the subject and the importance of this step to potentiate the value of the knowledge gathered during diagnosis and to influence positively the following steps during treatment, makes it worth considering on its own.

An important step given in this direction was the work presented at a workshop dedicated to prognosis in medicine [79], which later on gave place to a special issue on the topic [80]. Some of this proposals considered time explicitly as an important part of the prognostic task.

Some of this work resort to statistics, evidence theory, naive Bayes classifier, genetic algorithms, case based reasoning and fuzzy logic to formulate predictive models about specific diseases and conditions (cardiovascular risk, breast cancer, colorectal cancer, hip arthroplasty, acute trauma, intensive care, diabetes and anaesthetics). Although all of this models are considering the trends that some processes are expected to develop in time, mostly none of them considers time explicitly.

An exception is [81] where a prognostic model is offered that focuses on predicting the long-term outcome after femoral neck fracture with implantation of hip endoprosthesis. Domain knowledge is encoded as a hierarchical decision model which mix inferred knowledge with experts provided knowledge. As there is strong evidence suggesting that the patient's condition stabilizes after approximately one year, the system has to be aware of the development of the patient's condition through time in order to predict the continuation of that condition beyond that period. Amongst the temporal features given consideration were "timing of the operation", "time from injury", and "hospitalization time". However due to reasons intrinsic to the theoretical model underlying the proposal the authors focused on the first of them discarding the other two.

More recently, a report of a prognosis system to anticipate epidemics of influenza, *TeCoMed*, was given in [82]. Temporal abstraction techniques are combined with case-based reasoning to match the current pattern of evolution with previous patterns were early warnings were successful. Another report from the same volume is about the *NEONATE* project [83]. Sub-optimal decision procedures are used there to alert clinical and nursing staff about potentially concerning situations in a neonatal intensive care unit.

Many contributions to the general field of AI can be identified with the task of prognosis however there the concept is most usually referred with the label of "temporal projection" (see for example [84] for an important landmark and [5] for a more recent account). It is expected that as a longer-term interaction between the fields develop, this two approaches will be better integrated and the benefits of the mutual experiences will lead to a fully exploitation of the potential that explicit temporal reasoning can bring to prognosis.

4.4 Therapy/Treatment

Therapy can be defined as *the treatment of a disease* and also refers to predefined general courses of action to be applied to a patient on the process of treating a disease. Even when having a general pattern to be followed, there are many possibilities and combinations of events that can occur and the possible plans, or variations, to be followed are greater than expected. However planning in advance may help to identify possible unforeseen courses of development.

From the AI perspective planning can be defined generally as devising the course of action to be taken towards achieving a goal. Part of the classical approach to planning in AI will include considering: i) a description of the world as it is at the moment of the plan being applied, ii) a description of the available actions as well as the conditions in which they can be applied and the effects they would provoke in the world once applied, iii) a goal to be achieved. The description of the world may also include the resources available as they can play a decisive role in a realistic plan.

Time in the context of therapy planning can be of paramount importance. The community has been working actively on this area, see for example a good summary in [85]. Work in the area has been usually organized around the concepts of the so-called *care-protocols* or *critical guidelines* that are specialized to particular domains and can be instantiated accordingly to the context. These descriptions may involve description of actions to be applied, resources needed, activities, intentions, roles, and more. Some of this work have already addressed the importance of time as one of the important resources to be considered in therapy planning.

For example *TraumaTIQ*, [86] and [87], a real-time critiquing interface, that allows to improve plans given a problem scenario, for trauma care considered the importance of planning ahead the availability of critical resources and preventing staff to carry out wrong actions when the context changes.

The work presented in [88] focuses on the use dynamic decision problems using *Influence Views* [89] a graphical framework based on Markov Decision Processes [90], in order to solve decision problems in which the optimal choice has to be revised periodically in accordance to the evolution of the patient's conditions. The proposed methodology is applied to the plan of the prophylaxis in patients affected by a Hereditary Spherocytosis.

An interesting combination of temporal reasoning capabilities have been already incorporated in the *Asgaard* project [91]. On a broad scale this work considered two different levels of time: *design time*, encompassing tasks such as plan generation, plan verification and plan validation, and *execution time*, including tasks such as plan selection, plan adaptation and plan execution. On a more specific level, the language *Asbru* which is used as the basis for the planning process incorporates a variety of temporal features. For example, it supports the use of different granularities and reference points to represent multiple time lines. Considers durative actions, and time associated to events, actions, plans and world states with uncertainty in their appearances. Asbru also provides some general temporal relations between sub-plans like "do-some-any-order", "do-some-together". It is also worth to be emphasized that a distinctive feature of the *Asgaard* project is that they explicitly aim at providing a framework where diagnosis and treatment are linked [92].

In [93] a mixed-initiative approach is described that allows user and system to interact to solve typical problems in clinical management as conflicting resources and change of constraints over time. The medical domain is characterized in this proposal as a constraint satisfaction problem [94]. A central hypothesis in this work is that the medical domain is over-constrained and as usual may happen, constraints found in this area of application are usually opposed: e.g. patient's safety, staff satisfaction, use of limited places and material. The temporal representation of the scenario is made by using constraint networks corresponding to *Simple Temporal Problems* [72]. To resolve this tension on the system they classify constraints into *relaxable* and *non-relaxable*. Non-relaxable constraints are explicitly visualized and the user is offered the possibility to interact with the system to solve particular critical conflicts, to do that the user can ask from the system descriptions of the causes for the conflicting resources and consult a solution repository to reorganize the protocol in a non-conflicting way.

Planning and scheduling patient's requests for examination tests minimizing patient's stay in hospital and maximizing laboratory resources utilization is considered from a planning perspective in [95]. Sequential, parallel and periodic activities can be considered over sharable, non-shareable and consumable resources. The planning activities is carried out by the TRL-Planner and the temporal reasoning activities by the specialized module TRLi [96] that can consider temporal points, temporal intervals, temporal instances, and uncertain temporal intervals.

The importance of temporal scheduling constraints are recognized as important for therapy management also in [97]. There a method is proposed with the following characteristics: (1) it checks whether temporal scheduling constraints are consistent with scheduling constraints contained in a guideline (2) it provides suggestions for an equivalent but more explicit representation of non-minimal constraints (3) it can be used to assemble feasible time intervals for the execution of the prescribed activities. The language used to write the guidelines is *Asbru*. The implemented algorithms run in polynomial time.

CG_KRM (Clinical Guidelines Knowledge Representation Manager), see [98] and [99], also considers an interesting range of actions-centered concepts. They distinguish between *atomic* actions and *composite* actions. The qualitative relationships between these actions can be: sequential, concurrent and alternative (e.g. when if-then-else like decisions are made). Temporal reasoning is made more sophisticated for sequential actions by the introduction of boundaries for the time elapsed between the ending of an action and the next one in the sequence. Another useful temporal feature for the area of application is the possibility to express the frequency of actions to be repeated within a time window, useful for example when a drug has to be applied a given number of times within a definite period of time. The proposal is illustrated with an application to the scenario of a treatment applied to a patient with Non-Hodgkin lymphoma.

Although we focused on the latest publications there is more interesting research aiming to provide time-related support in guidelines centered languages which is worth to look at. Other previous proposals with explicit inclusion of temporal constraints, are: the previously cited *Asbru* language, Sherman et al.'s proposal [100] based on the Arden syntax, DILEMMA [101] and its successor the PRESTIGE project [102], the EON approach [103], the GuideLine Interchange Format (GLIF) [104], and PROforma [105].

Despite the work that has been done in planning and in particular in the medicine-related applications, much progress has yet to be done in order to successfully incorporate this technique to the everyday process of medical care. As medical procedures evolve by adapting existing treatments to new findings in the specific area or finding treatment procedures for new diseases, new challenges will be added to the already exiting ones. This will certainly promise to provide a rich source of challenges for the AI community.

4.5 Time Granularity

As mentioned in a previous section the concept of granularity [106] is ubiquitous to any temporal system and decisions taken at such level will affect the functioning of the system. It is a decision taken at the temporal structure level but it will also be present in the reasoning mechanisms. Just to mention a simple example, it can be inferred that fever was present as a symptom one day of the week if we previously knew it was present during the whole week but sound inference cannot be drawn the other way round. Hence, a system allowing different level of granularities must supplement that with an accordingly granularity-sensitive inferential system.

Using different time units is regular practice in medicine. Either in management, diagnosis, prognosis or treatment, there is always a need to relate processes that develop at a given granularity to those that may happen at a different one. For example, sometimes symptoms are identified if a given pattern has been occurring each day for a week or so. Other times an occurring pattern that was present during weeks can be associated to a season, e.g. when diagnosing hay fever. As it can be seen, the problem of handling granularity flexibly and effectively in medicine is more than calendric-related unit conversion.

Allowing different levels of granularity can also have a computational effect that has to be balanced. On one side some resources should be focused on doing all the granularity-related inferences and that may introduce further computational complexity to the system. On the other hand all that work can be compensated by the system being able to focus the inferential activities at a more abstract or higher level of granularity. Guiding the inferences by using structured knowledge may bring important computational savings.

In [107] an integral proposal for the handling of time granularity in medicine related systems is given. Time-objects are the concepts that have a time-related existence in the system. Time-objects can be related either by temporal, structural and causal relationships. Temporal references can be either point-based, that will be assumed as atomic at a given granularity level, or with duration, that can be further decomposable. Different relationships are defined for these two classes. Another distinction made in this proposal is that one of generic and case-specific time-related knowledge which become useful when distinguishing therapeutic guidelines from patient histories. Occurrences in this framework are expressed relative to a conceptual level associated with a particular granularity. Along with the time-objects centered issues of granularity a set of related functions is provided to achieve different higher level tasks like detecting conflicts, clipping or mapping occurrences from one level of granularity to another. The possible applications of the system is illustrated trough the consideration of a scenario of normal evolution of the ossification of cervical spine.

A more recent contribution [108] provides support for granularity handling of both anchored and unanchored temporal references. The proposal is illustrated within the context of diagnosing and monitoring of patients with unstable angina in the broader context of cardiology but it is clear that the kind of temporal information they handle with their proposal is useful in many different scenarios within medical applications. The proposal focuses on providing the basis for a Databases Management System capable to support the following basic activities: a) representing and storing time instants with different granularities and time spans with different and mixed granularities, b) handling granularity mismatches in operations between temporal primitives with different granularities c) converting a temporal primitive from one granularity to another, and d) considering different interpretations for time labels.

Meanwhile interesting contributions has been provided from AI and Databases in relation with granularity handling. For example, [109] is centered in a flexible language and a formalization that allows the characterization of temporal objects capable to encompass notions of structured temporal references e.g. calendars, provided they have a finite number of nested structures. Some concepts of this proposal has been already connected with medical scenarios [110]. [111] also provides a formalization which is then illustrated by applications to different areas like temporal databases, AI and data mining.

Some previous versions of these contributions to the topic has been compared before however it would be an interesting step to have an up to date comparison in terms of flexibility to create new temporal references and in terms of the computational cost needed, both important factors for the successful application of these technologies in medical applications.

5 Future Work and Proposed Agenda

Although some advances have been made, it is clear that many things still can be done to make the collaboration human-machine even more successful than what have been up to now for this particular area of application. A simplified historical evolution of some of the systems and proposals reviewed is given in table 1. A summary of the particular fields of medicine where time-related techniques have been applied to is provided in table 2.

In the classical sequence of steps: diagnosis, prognosis and therapy planning, research into *prognosis*-related topics, still has not generated as much interest within the AI in Medicine research community as diagnosis and therapy planning did. Although some work has been done still much can be investigated on the consideration of time-related concepts, for example on using AI symbolic deductive approaches. However it seems this step in the diagnosis process deserves further consideration as it is a very important link between diagnosis and therapy as temporal knowledge gathered during diagnosis can be fruitfully exploited when forecasting the evolution of a patient. Equally useful would be to export these already foreseen outcomes to the therapy planning stage. Having this information at the initial steps of the planning process can influence and focus the planning process in many beneficial ways, for example providing valuable information for the initial steps and focusing the task on the most critical issues. It could be argued that some prognosis activities are embedded in the therapy planning process but, being planning a task demanding a great deal of computational efforts, this should not be the case. For example, as a result of the prognosis activities planning can be discarded altogether. If it is not, the prognosis stage, as a result of the interaction between the clinicians and the system, can deliver a better informed description about the initial assumptions, previous cases to consider, and the goals to be achieved.

Causality is a very important notion to AI and Medicine and certainly can be one of those concepts to be considered during prognosis. It is expected that recent proposals coming from AI to automate causality-based simulations, see for example [112], will allow to predict hidden effects in complex interactions.

Natural language is a very challenging area as a general problem in AI. Suitable restrictions given by the requirements of an area can turn the problem feasible at the time it retains its advantages, namely making the communication between humans and assistant machines more flexible, realistic and useful. Given some appropriate restrictions allowing computational tractability this may turn into a very fertile area of research where temporal issues are very important. Implicit references to time-based activities, a usually neglected way of temporal reference in the proposals of the area can be useful here. Contrary to initial suggestions that it would not be an adequate vehicle to bring information to a system, it was illustrated in [113] that it can be positively incorporated to the abstraction process with computational gains.

Argumentation can be very useful as a way to structure different lines of reasoning and providing consequent explanations for each of the considered alternatives. Some work has been already done in relation to structuring explanations in medicine from an atemporal perspective [114]. Allowing these explanations to include the temporal hypothesis that underlies the reasoning tasks [115] to identify a disease or to support a therapy plan may be an important step towards making this explanations more precise and realistic. An interesting proposal has been made by [116] on these line. Still, a good deal of research and development is needed to successfully ground this proposals into real life use.

Active databases a standard tool that has been extensively researched in CS and can be used for monitoring the occurrence of events and reacting in a context sensitive way. Still little or almost no attention has been given to the topic in the literature on how artificial intelligence can contribute to make the detection of complex event detection [117] in medicine-centered contexts soundly and efficiently. This research into complex event specification and languages allowing flexible constructions on event-based references may be also a useful vehicle to characterize the specific patterns of evolution that are seek during intelligent data analysis, other stages of diagnosis or prognosis, as well as to specify goals and conditions during therapy planning.

Finally, *verification and validation of protocols* (in the generic sense of the word that denotes a set of rules enabling communication between two or more entities) has been extensively studied in Computer Science [118]. Any specification of steps to be followed which can be specified in some formal notation can be analyzed to check if the possible models captured by the specification are the expected ones. Although it is not an application of temporal reasoning in the typical sense of the term from an AI perspective, it is certainly centered on reasoning about the possible evolutions in time of a system. However, verification in medical protocols has not yet received as much attention. An interesting report is made in [119] on a protocol verification framework specifically related to the guidelines specification language *Abru*. Still seems that much work should be done on this line. For example, it seems that no systematic work has been done on verification of dynamic behavioral properties, in the software engineering sense. However, because of the importance of the application area and the safety-critical character of the activities involved, verification and validation issues should be in the list of important projects of the area.

6 Conclusions

As research in the previous decades already identified several key issues and gave the important step of providing better access to time-oriented data in medicine, the field seems to have reached maturity as to allow other goals in relation with temporal reasoning as studied in the AI field.

On the other hand time has been also extensively considered as a general topic in AI for the past two decades. Research in temporal representation and reasoning also has their own well-established events (see for example, [1–4]). Although there is awareness throughout the research community that have actively participated in this events, still many developments has not been transferred/adapted from the more general field to the more specific one. Also the feedback from the more specialized field to the more general one is not either fluent or high in quantitative terms. It seems that an interesting topic for a future agenda would be to increase and foster this interaction as it certainly will be beneficial for both communities.

One step has been given in this article by summarizing previous contributions in the field and highlighting some possible continuations for the community to explore. Certainly, there is a valuable experience gained in the last decade both within the temporal reasoning and the MI communities in this interaction. This should provide a firm departing point for the next generation of intelligent and time-aware assistants in medicine.

Acknowledgments

The author wish to express his gratitude towards the members of the Medical Informatics Research group and the Northern Ireland Knowledge Engineering Laboratory, both from the University of Ulster at Jordanstown, for their support in the realization of this work.

References

[1] S. Goodwin, A. Trudel (Eds.), Proceedings of the Seventh International Workshop on Temporal Representation and Reasoning, IEEE Computer Society Press, Cape Breton, Canada, 2000.

[2] C. Bettini, A. Montanari (Eds.), Proceedings of the Eigth International Workshop on Temporal Representation and Reasoning, IEEE Computer Society Press, Cividale del Friuli, Italy, 2001.

[3] B. T. A. Artale, M. Fisher (Ed.), Proceedings of the Ninth International Workshop on Temporal Representation and Reasoning, IEEE Computer Society Press, Manchester, UK, 2002.

[4] M. Reynolds, A. Sattar (Eds.), Proceedings of the Tenth International Workshop on Temporal Representation and Reasoning and Fourth International Conference on Temporal Reasoning, IEEE Computer Society Press, Cairns, Queensland, Australia, 2003.

[5] S. J. Russell, P. Norvig, Artificial Intelligence: A Modern Approach (Second Edition), Prentice Hall, 2003.

[6] J. van Bemmel, M. A. Musen (Eds.), Medical Informatics, Springer Verlag, 1997.

[7] J. C. Augusto, The logical approach to temporal reasoning, Artificial Intelligence Review 16 (4) (2001) 301–333.

[8] J. Allen, H. Kautz, R. Pelavin, J. Tenenberg, Reasoning About Plans, Morgan Kaufmann, San Mateo, California, 1991.

[9] D. Gabbay, I. Hodkinson, M. Reynolds, Temporal Logic (Mathematical Foundations and Computational Aspects), Clarendon Press, Oxford, 1994.

[10] E. Sandewall, Features and Fluents, Oxford University Press, Oxford, 1994.

[11] A. Galton, Time and change, in: D. Gabbay, C. Hogger, J. Robinson (Eds.), Handbook of Logic in Artificial Intelligence and Logic Programming - Vol. 4 (Epistemic and Temporal Reasoning), Clarendon Press, 1995, pp. 175–240.

[12] J. F. A. K. van Benthem, Exploring Logical Dynamics, CSLI Publications, Stanford, 1996.

[13] J. Allen, Towards a general theory of action and time, Artificial Intelligence 23 (1984) 123–154.

[14] J. Pinto, R. Reiter, Reasoning about time in the situation calculus, Annals of Mathematics and Artificial Intelligence 14 (2-4) (1995) 251–268.

[15] W. H. Newton-Smith, The Structure of Time, Routledge and Kegan Paul, London, 1980.

[16] D. Cuckierman, J. Delgrande, Towards a formal characterization of temporal repetition with closed time, in: Proceedings of TIME98, IEEE Computer Society Press, 1998, pp. 140–147.

[17] R. Turner, Logic for Artificial Intelligence, John Wiley & Sons., 1984.

[18] J. F. A. K. van Benthem, The Logic of Time (second edition), Reidel, 1991.

[19] Y. Shoham, Ten requirements for a theory of change, New Generation Computing 3 (1985) 467–477.

[20] C. Hamblin, Instants and intervals, in: F. H. J. Fraser, G. Muller (Eds.), The Study of Time, Springer Verlag, New York, 1972, pp. 324–328.

[21] J. Allen, P. Hayes, A common-sense theory of time, in: Proceedings IJCAI-85, Vol. 1, 1985, pp. 528–531.

[22] A. Bochman, Concerted instant-interval temporal semantics i: Temporal ontologies, Notre Dame Journal of Formal Logic 31 (3) (1990) 403–414.

[23] A. Bochman, Concerted instant-interval temporal semantics ii: Temporal valuations and logics of change, Notre Dame Journal of Formal Logic 31 (4) (1990) 581–601.

[24] A. Galton, A critical examination of Allen's theory of action and time, Artificial Intelligence 42 (1990) 159–188.

[25] L. Vila, Ip: An instant-period based theory of time, in: R. Rodriguez (Ed.), Proceedings of the Workshop on Spatial and Temporal Reasoning in ECAI 94, 1994.

[26] J. C. Augusto, A general framework for reasoning about change, New Generation Computing 21 (3) (2003) 209–247.

[27] C. Freksa, Temporal reasoning based on semi-intervals, Artificial Intelligence 54 (1992) 199–227.

[28] J. Allen, G. Ferguson, Actions and events in interval temporal logic, Journal of Logic and Computation 4 (1994) 531–579.

[29] A. Goldman, A Theory of Human Action, Princeton University Press, Princeton, New York, 1970.

[30] A. Galton, Qualitative Spatial Change, Oxford University Press, 2000.

[31] A. Prior, Past, Present and Future, Clarendon Press, 1967.

[32] A. Galton, The logic of occurrence, in: A. Galton (Ed.), Temporal Logic and their Applications, Academic Press, 1987, pp. 169–196.

[33] Y. Shahar, Timing is everything: Temporal reasoning and temporal data maintenance in medicine, in: Proceedings of Seventh Joint European Conference on Artificial Intelligence in Medicine and Medical Decision Making (AIMDM'99), Springer Verlag, 1999, pp. 30–46, Aalborg, Denmark.

[34] E. T. Keravnou (Ed.), Medical Temporal Reasoning, Elsevier, Special Issue from Artificial Intelligence in Medicine, Vol. 3, no. 6, 1991.

[35] E. T. Keravnou (Ed.), Temporal Reasoning in Medicine, Elsevier, Special Issue from Artificial Intelligence in Medicine, Volume 8, Issue 3, 1996.

[36] C. Combi, Y. Shahar (Eds.), Time-Oriented Systems in Medicine, Elsevier, Special Issue from Computers in Biology and Medicine, Volume 27, Issue 5, 1997.

[37] Y. Shahar, C. Combi (Eds.), Intelligent Temporal Information Systems in Medicine, Kluwer Academic, Special Issue from Journal of Intelligent Information Systems, Volume 13, Numbers 1-2, 1999.

[38] E. T. Keravnou, Introduction to the special issue on Temporal Reasoning in Medicine, [35] (1996) 187–191.

[39] C. Combi, Y. Shahar, Editorial to Time-Oriented Systems in Medicine, [36] (1997) 349–351.

[40] C. Combi, Y. Shahar, Temporal reasoning and temporal data maintenance in medicine: issues and challenges, [36] (1997) 353–368.

[41] Y. Shahar, C. Combi, Editors' foreword for special issue on Intelligent Temporal Information Systems in Medicine, [37] (1999) 5–8.

[42] A. Tansel, J. Clifford, S. Gadia, S. Jajodia, A. Segev, R. Snodgrass, Temporal Data Bases (Theory, Design and Implementation), The Benjamin Cummings Pub. Co., California, 1993.

[43] O. Etzioni, S. Jajodia, S. Sripada (Eds.), Temporal Databases: Research and Practice, Springer-Verlag, 1998.

[44] A. K. Das, S. Tu, G. Purcell, M. Musen, An extended sql for temporal data management in clinical decision-support systems, in: Proceedings of the Sixteenth Annual Symposium on Computer Applications in Medical Care, 1992, pp. 128–32.

[45] I. Goralwalla, M. Ozsu, D. Szafron, Modeling medical trials in pharmacoeconomics using a temporal object model, Computers in Biology and Medicine 27 (5) 369–387, special Issue on TimeOriented Systems in Medicine, 1997.

[46] M. Ozsu, R. Peters, D. Szafron, B. Irani, A. Lipka, , A. Muñoz, Tigukat: A uniform behavioral objectbase management system, The VLDB Journal 4 (1995) 100–147.

[47] J. H. Nguyen, Y. Shahar, S. W. Tu, A. K. Das, M. A. Musen, Integration of temporal reasoning and temporal-data maintenance into a reusable database mediator to answer abstract, time-oriented queries: The tzolkin system, [37] (1999) 121–145.

[48] R. Snodgrass, I. Ahn, A taxonomy of time in databases, in: Proceedings of ACM SIGMOD International Conference on Management of Data, Austin, Texas, 1985, pp. 236–246.

[49] A. K. Das, M. A. Musen, A temporal query system for protocol-directed decision-support, Methods of Information in Medicine 33 (1994) 358–370.

[50] Y. Shahar, M. Musen, Knowledge-based temporal abstraction in clinical domains, Artificial Intelligence in Medicine 8 (1996) 267–298.

[51] R. Bellazzi, B. Shahar, Intelligent data analysis in medicine and pharmacology (panel summary), held during AMIA'99 Annual Symposium, November 6-10 - Washington, DC, U.S.A., 1999.

[52] Y. Shahar, Knowledge-based temporal interpolation, Journal of Experimental and Theoretical Artificial Intelligence 11 (1999) 123–144.

[53] Y. Shahar, Dynamic temporal interpretation contexts for temporal abstraction, Annals of Mathematics and Artificial Intelligence 22 (1998) 159–192.

[54] A. Salatian, J. Hunter, Deriving trends in historical and real-time continuously sampled medical data, [37] (1999) 47–71.

[55] M. J. O'Connor, W. E. Grosso, S. W. Tu, M. A. Musen, Rasta: A distributed temporal abstraction system to facilitate knowledge-driven monitoring of clinical databases, in: Proceedings of MedInfo2001, 2001.

[56] A. Lowe, R. W. Jones, M. J. Harrison, Temporal pattern matching using fuzzy templates, [37] (1999) 27–45.

[57] R. Bellazzi, C. Larizza, P. Magni, S. Montani, M. Stefanelli, Intelligent analysis of clinical time series: an application in the diabetes mellitus domain, Artificial Intelligence in Medicine 20 (2000) 37–57.

[58] N. Lavrac, E. Keravnou, I. Mozetic, Second generation knowledge acquisition methods and their application to medicine, in: E. Keravnou (Ed.), Deep Models for Medical Knowledge Engineering, Elsevier, New York, 1992, pp. 177–198.

[59] N. Lavrac, E. Keravnou, B. Zupan, Intelligent data analysis in medicine, in: K. et al. (Ed.), Encyclopedia of Computer Science and Technology, Vol. 42, Dekker, New York, 2000, pp. 113–157.

[60] S. Cousins, M. Kahn, The visual display of temporal information, Artificial Intelligence In Medicine (1991) 341–357.

[61] R. H.Dolin, Modeling the temporal complexities of symptoms, Journal of the American Medical Informatics Association 2 (5) (1995) 323–331.

[62] R. Reiter, A theory of diagnosis from first principles, Artificial Intelligence 32 (1987) 57–95.

[63] J. de Kleer, B. Williams, Diagnosing multiple faults, Artificial Intelligence 32 (1987) 97–130.

[64] L. Console, P. Torasso, A spectrum of logical definitions of model-based diagnosis, Computational Intelligence 7 (1991) 133–141.

[65] L. Console, P. Torasso, On the co-operation between abductive and temporal reasoning in medical diagnosis, Artificial Intelligence in Medicine 3 (1991) 291–311.

[66] L. Console, P. Torasso, Temporal constraint satisfaction on causal models, Information Science 68 (1991) l–32.

[67] P. J. F. Lucas, Model-based diagnosis in medicine, Artificial Intelligence in Medicine 10 (3) (1997) 201–208.

[68] I. J. Haimowitz, I. S. Kohane, Managing temporal worlds for medical trend diagnosis, [35] 8 (3) (1996) 299–321.

[69] J. Gamper, W. Nejdl, Abstract temporal diagnosis in medical domains, Artificial Intelligence in Medicine 10 (3) (1997) 209–234.

[70] J. Wainer, A. de Melo Rezende, A temporal extension to the parsimonious covering theory, Artificial Intelligence in Medicine 10 (3) (1997) 235–255.

[71] M. Dojat, N. Ramaux, D. Fontaine, Scenario recognition for temporal reasoning in medical domains, Artificial Intelligence in Medicine 14 (1998) 139–155.

[72] R. Dechter, I. Meiri, J. Pearl, Temporal constraint networks, Artificial Intelligence 49 (1991) 61–95.

[73] J. Wainer, S. Sandri, Fuzzy temporal/categorical information in diagnosis, [37] (1999) 9–26.

[74] W. Long, Temporal reasoning for diagnosis in a causal probabilistic knowledge base, [35] (1996) 193–215.

[75] S. Charbonnier, On-line extraction of successive temporal sequences from icu high-frequency data for decision support information, in: M. Dojat, E. T. Keravnou, P. Barahona (Eds.), Proceedings of 9th Conference on Artificial Intelligence in Medicine in Europe (AIME 2003), Protaras, Cyprus, October 18-22, 2003, Vol. 2780 of Lecture Notes in Computer Science, Springer Verlag, 2003, pp. 1–10.

[76] R. Bellazzi, C. Larizza, P. Magni, R. Bellazzi, Quality assessment of hemodialysis services through temporal data mining, in: M. Dojat, E. T. Keravnou, P. Barahona (Eds.), Proceedings of 9th Conference on Artificial Intelligence in Medicine in Europe (AIME 2003), Protaras, Cyprus, October 18-22, 2003, Vol. 2780 of Lecture Notes in Computer Science, Springer Verlag, 2003, pp. 11–20.

[77] D. Boaz, Y. Shahar, Idan: A distributed temporal-abstraction mediator for medical databases, in: M. Dojat, E. T. Keravnou, P. Barahona (Eds.), Proceedings of 9th Conference on Artificial Intelligence in Medicine in Europe (AIME 2003), Protaras, Cyprus, October 18-22, 2003, Vol. 2780 of Lecture Notes in Computer Science, Springer Verlag, 2003, pp. 21–30.

[78] J. Bouaud, B. Sroussi, B. Touzet, Abstracting the patient therapeutic history through a heuristic-based qualtitative handling of temporal indeterminacy, in: M. Dojat, E. T. Keravnou, P. Barahona (Eds.), Proceedings of 9th Conference on Artificial Intelligence in Medicine in Europe (AIME 2003), Protaras, Cyprus, October 18-22, 2003, Vol. 2780 of Lecture Notes in Computer Science, Springer Verlag, 2003, pp. 46–50.

[79] A. Abu-Hanna, P. Lucas (Eds.), Proceedings of the Workshop on Prognostic Models in Medicine (Artificial Intelligence and Decision Analytic Approaches), held during AIMDM'99 in Aalborg, Denmark, 20th - 24th June, 1999.

[80] A. Abu-Hanna, P. J. F. Lucas, Prognostic models in medicine (editorial), [120] (2001) 1–5.

[81] B. Zupan, J. Demsar, D. Smrke, K. Bozikov, V. Stankovski, I. Bratko, J. R. Beck, Predicting patient's long-term clinical status after hip arthroplasty using hierarchical decision modelling and data mining, [120] (2001) 25–31.

[82] R. Schmidt, L. Gierl, Prognosis of approaching infectious diseases, in: M. Dojat, E. T. Keravnou, P. Barahona (Eds.), Proceedings of 9th Conference on Artificial Intelligence in Medicine in Europe (AIME 2003), Protaras, Cyprus, October 18-22, 2003, Vol. 2780 of Lecture Notes in Computer Science, Springer Verlag, 2003, pp. 31–35.

[83] J. Hunter, G. Ewing, Y. Freer, F. Logie, P. McCue, N. McIntosh, Neonate: Decision support in the neonatal intensive care unit - a preliminary report, in: M. Dojat, E. T. Keravnou, P. Barahona (Eds.), Proceedings of 9th Conference on Artificial Intelligence in Medicine in Europe (AIME 2003), Protaras, Cyprus, October 18-22, 2003, Vol. 2780 of Lecture Notes in Computer Science, Springer Verlag, 2003, pp. 41–45.

[84] S. Hanks, D. McDermott, Nonmonotonic logic and temporal projection, Artificial Intelligence 33 (3) (1987) 379–412.

[85] C. Spyropoulos, Ai planning and scheduling in the medical hospital environment, Artificial Intelligence in Medicine 20 (2) (2000) 101–111.

[86] A. Gertner, Responding to users' informational needs in time-critical situations, in: Proceedings of Fourth International Conference on User Modeling, 1994, pp. 127–132.

[87] A. Gertner, Critiquing: Effective decision support in time-critical domains, Ph.D. thesis, University of Pennsylvania (1995).

[88] P. Magni, A new approach to optimal dynamic therapy planning, in: Proceedings of AMIA'98, 1998, pp. 936–940.

[89] D. K. Owens, R. D. Shachter, R. F. Nease, Representation and analysis of medical decision problems with influence diagrams, Medical Decision Making 17 (3) (1997) 241–262.

[90] T. Dean, M. Wellmann, Planning and Control, Morgan Kaufmann, San Mateo, California, U.S.A., 1991.

[91] S. Miksch, Plan management in the medical domain, AI Communications 12 (4) (1999) 209–235.

[92] A. Seyfang, S. Miksch, M. Marcos, Combining diagnosis and treatment using asbru, International Journal of Medical Informatics 68 (1-3) (2002) 49–57.

[93] A. Oddi, A. Cesta, Toward interactive scheduling systems for managing medical resources, Artificial Intelligence in Medicine 20 (2) (2000) 113–138.

[94] U. Montanari, Networks of constraints: fundamental properties and applications to picture processing, Information Science 7 (1974) 95–132.

[95] C. Marinagi, C. Spyropoulos, C. Papatheodorou, S. Kokkotos, Continual planning and scheduling for managing patient tests in hospital laboratories, Artificial Intelligence in Medicine 20 (2) (2000) 139–154.

[96] T. Panayiotopoulos, M. Gergatsoulis, A prolog-like temporal reasoning system, in: Proceedings of the 13th IASTED International Conference on Applied Informatics, 1995, pp. 123–126.

[97] G. Duftschmid, S. Miksch, W. Gall, Verification of temporal scheduling constraints in clinical practice guidelines, Artificial Intelligence in Medicine 25 (2) (2002) 93–121.

[98] A. Guarnero, M. Marzuoli, G. Molino, P. Terenziani, M. Torchio, K. Vanni, Contextual and temporal clinical guidelines, in: Proceedings of AMIA98, 1998, pp. 683–687.

[99] P. Terenziani, F. Mastromonaco, G. Molino, M. Torchio, Executing clinical guidelines: temporal issues, in: Proceedings of AMIA2000, 2000, pp. 848–852.

[100] E. Sherman, G. Hripcsak, J. Starren, R. Jender, P. Clayton, Using intermediate states to improve the ability of the arden syntax to implement care plans and reuse knowledge, in: Proceedings of the Annual Symposium on Computer Applications in Medical Care (SCAMC95), 1995, pp. 238–242.

[101] S. Herbert, C. Gordon, A. Jackson-Smale, S. Renaud, Protocols for clinical care, Computer Methods and Programs in Biomedicine 48 (1995) 21–26.

[102] C. Gordon, I. Herbert, P. Johnson, Knowledge representation and clinical practice guidelines: the dilemma and prestige projects, in: Proceedings of Medical Informatics Europe96, R. Brompton Hospital, London., 1996.

[103] M. Musen, S. Tu, A. Das, Y. Shahar, Eon: a component-based approach to automation of protocol-directed therapy, Journal of the American Medical Informatics Association (JAMIA) (1996) 367388.

[104] L. Ohno-Machado, J. Gennari, S. Murphy, N. Jain, S. Tu, D. Oliver, E. Pattison-Gordon, The guideline interchange format: a model for representing guidelines, Journal of the American Medical Association (JAMA) 5 (1998) 357372.

[105] J. Fox, N. Johns, A. Rahmanzadeh, Disseminating medical knowledge: the proforma approach, Artificial Intelligence in Medicine 14 (1998) 157–81.

[106] J. Hobbs, Granularity, in: Proceedings of the IJCAI 85, IJCAI, 1985, pp. 432–435.

[107] E. T. Keravnou, A multidimensional and multigranular model of time for medical knowledge-based systems, [37] (1999) 73–120.

[108] I. Goralwalla, Y. Leontiev, M. Ozsu, D. Szafron, C. Combi, Temporal granularity: Completing the puzzle, Journal of Intelligent Information Systems 16 (1) (2001) 41–63.

[109] D. Cukierman, J. P. Delgrande, A formalization of structured temporal objects and repetition, in: [1], 2000, pp. 13–20.

[110] S. Chakravarty, Y. Shahar, A constraint-based specification of periodic patterns in time-oriented data, in: Proceedings of the Sixth International Workshop on Temporal Representation and Reasoning (TIME'99), 1999, pp. 29–40.

[111] C. Bettini, S. Jajodia, S. X. Wang, Time Granularities in Databases, Data Mining, and Temporal Reasoning, Springer-Verlag, Berlin, 2000.

[112] A. Galton, J. C. Augusto, Stratified causal theories for reasoning about deterministic devices and protocols, in: Proceedings of 9th International Symposium on Temporal Representation and Reasoning (Time-2002), 2002, pp. 52–54.

[113] C. F. Aliferis, C. F. Cooper, M. E. Pollack, B. G. Buchanan, M. M. Wagner, Representing and developing temporally abstracted knowledge as a means towards facilitating time modeling in medical decision-support systems, [36] (1997) 411–434.

[114] F. de Rosis, F. Grasso, D. C. Berry, Strengthening argumentation in medical explanations by text plan revision, in: E. Keravnou, C. Garbay, J. W. R. Baud (Eds.), Proceedings of Proceedings of 6th Conference on Artificial Intelligence in Medicine Europe (AIME'97), 1997, pp. 349–360, grenoble, France, vol. 1211 of Lecture Notes in Artificial Intelligence, Springer-Verlag.

[115] J. C. Augusto, G. R. Simari, Temporal defeasible reasoning, Knowledge and Information Systems 3 (3) (2001) 287–318.

[116] J. Fox, S. Das, Safe and Sound, Artificial Intelligence in Hazardous Applications, AAAI Press/MIT Press, 2000.

[117] A. Galton, Eventualities, in *The Handbook of Time and Temporal Reasoning in Artificial Intelligence*, Vila, van Beek, Boddy, Fisher, Gabbay, Galton and Morris (eds.), MIT Press, to be published.

[118] G. Holzmann, The SPIN Model Checker, Addison Wesley, 2003.

[119] G. Duftschmid, S. Miksch, Y. Shahar, P. Johnson, Multi-level verification of clinical protocols, in: Proceedings of Workshop on Validation & Verification of Knowledge-Based Systems (V&V'98), 1998, trento, Italy.

[120] A. Abu-Hanna, P. J. F. Lucas (Eds.), Prognostic Models in Medicine, Schattauer GmbH, Special Issue of Journal Methods for Informatics in Medicine, Volume 40, Issue 1, 2001.

Table 1
Some of the latest contributions to time-aware decision support systems

Year	*Diagnosis*	*Prognosis*	*Therapy/Treatment*
1991	(Cousins et al.) [60]		
1992			T-Helper [44]
1993			
1994	Chronus [49]		TraumatIQ [86]
1995	Tigukat [46]		DILEMA [101] (Sherman et al.) [100]
1996	TrendDx [68] HDP [74] Résumé [50]		EON [103] Prestige [102]
1997	t-PCT [70] (Gamper et al.) [69] (Goralwalla et al.) [45]		
1998	DejaVu [71]		CG-KRM [98] GLIF [104] PROFORMA [105] (Magni) [88]
1999	Tzolkin [47] (Wainer et al.) [73] (Lowe et al.) [56]		Abru-Asgaard [91]
2000	(Bellazzi et al.) [57]		Hostess [95] (Oddi et al.) [93] (Terenziani et al.) [99]
2001	RASTA [55]	(Zupan et al.) [81]	
2002			(Seyfang et al.) [92] (Duftschmid et al.) [97]
2003	(Charbonnier) [75] (Bellazzi) [76] *Idan* ([77])	TeCoMed([82]) NEONATE ([83])	

Table 2
Some areas of medicine with time-related techniques mentioned in this article

Area of Application	*Related Work*
HIV	[44]
Medical trials in pharmaco-economics	[45]
Diabetes therapy	[52]
Central venous pressure	[54]
Anaesthesia monitoring	[56]
Diabetes mellitus domain	[57]
Hepatitis B	[69]
Food-born disease domain	[70]
Intoxication by ingestion	[73]
Heart diseases diagnosis	[74]
Ventilation management unit	[71]
Assesment of clinical performance of a hemodialysis service	[76]
Femoral neck fracture with implantation of hip endoprosthesis	[81]
Intensive care unit	[75]
Influenza epidemics	[82]
Neonatal intensive care unit	[83]
Trauma care	[87]
Hereditary spherocytosis	[88]
Management of hyperbilirubinemia in the healthy term new-born	[92]
Clinical management	[93]
Management of patient's requests for examination tests	[95]
Therapy management	[97]
Non-Hodgkin lymphoma	[99]