

Extending Temporal Databases to Deal with Telic/Atelic Medical Data

Paolo Terenziani, Richard T. Snodgrass, Alessio Bottrighi, Mauro Torchio, and
Gianpaolo Molino

May 10, 2006

TR-85

A TIMECENTER Technical Report

Title	Extending Temporal Databases to Deal with Telic/Atelic Medical Data
	Copyright © 2006 Paolo Terenziani, Richard T. Snodgrass, Alessio Bottrighi, Mauro Torchio, and Gianpaolo Molino. All rights reserved.
Author(s)	Paolo Terenziani, Richard T. Snodgrass, Alessio Bottrighi, Mauro Torchio, and Gianpaolo Molino
Publication History	2005. An earlier version of the paper was presented in the Conference on Artificial Intelligence in Medicine (AIME 05). May 2006. A TIMECENTER Technical Report.

TIMECENTER Participants

Aalborg University, Denmark

Christian S. Jensen (codirector), Simonas Šaltenis, Janne Skyt, Giedrius Slivinskas, Kristian Torp

University of Arizona, USA

Richard T. Snodgrass (codirector), Bongki Moon, Sudha Ram

Individual participants

Yun Ae Ahn, Chungbuk National University, Korea; Michael H. Böhlen, Free University of Bolzano, Italy; Curtis E. Dyreson, Washington State University, USA; Dengfeng Gao, Indiana University South Bend, USA; Fabio Grandi, University of Bologna, Italy; Heidi Gregersen, Aarhus School of Business, Denmark; Vijay Khatri, Indiana University, USA; Nick Kline, Microsoft, USA; Gerhard Knolmayer, University of Bern, Switzerland; Carme Martín, Technical University of Catalonia, Spain; Thomas Myrach, University of Bern, Switzerland; Kwang W. Nam, Chungbuk National University, Korea; Mario A. Nascimento, University of Alberta, Canada; John F. Roddick, Flinders University, Australia; Keun H. Ryu, Chungbuk National University, Korea; Dennis Shasha, New York University, USA; Michael D. Soo, amazon.com, USA; Andreas Steiner, TimeConsult, Switzerland; Paolo Terenziani, University of Torino, Italy; Vassilis Tsotras, University of California, Riverside, USA; Fusheng Wang, Siemens, USA; Jef Wijsen, University of Mons-Hainaut, Belgium; and Carlo Zaniolo, University of California, Los Angeles, USA

For additional information, see The TIMECENTER Homepage:

URL: <<http://www.cs.aau.dk/TimeCenter>>

Any software made available via TIMECENTER is provided “as is” and without any express or implied warranties, including, without limitation, the implied warranty of merchantability and fitness for a particular purpose.

The TIMECENTER icon on the cover combines two “arrows.” These “arrows” are letters in the so-called *Rune* alphabet used one millennium ago by the Vikings, as well as by their predecessors and successors. The Rune alphabet (second phase) has 16 letters, all of which have angular shapes and lack horizontal lines because the primary storage medium was wood. Runes may also be found on jewelry, tools, and weapons and were perceived by many as having magic, hidden powers.

The two Rune arrows in the icon denote “T” and “C,” respectively.

Abstract

Objective. In this paper, we aim at defining a general-purpose data model and query language coping with both “telic” and “atelic” medical data.

Background. In the area of Medical Informatics, there is an increasing realization that temporal information plays a crucial role, so that suitable database models and query languages are needed to store and support it. However, despite the wide range of approaches in the area, in this paper we show that a relevant class of medical data cannot be properly dealt with.

Methodology. We first show that data models based on the “point-based” semantics, which is (implicitly or explicitly) assumed by the totality of temporal Database approaches, have several limitations when dealing with “telic” data. We then propose a new model (based on the “interval-based” semantics) to cope with such data, and extend the query language accordingly.

Results. We propose a new three-sorted model and a query language to properly deal with both “telic” and “atelic” medical data (as well as nontemporal data). Our query language is flexible, since it allows one to switch from “atelic” to “telic” data, and vice versa.

1 Introduction

In the area of medicine, an explicit management of the time when symptoms took place and clinical actions were taken is needed to model the patients’ state (e.g., for diagnostic or therapeutic purposes [1]). Thus several data models used to capture clinical data provide suitable supports to explicitly deal with time (consider, e.g., [2], [3], [4], [5]). Over the last two decades, the database community has devised many different approaches to model the validity time of data (i.e., the time when the data holds [6]). In particular, many temporal extensions to the standard relational model were developed, and more than 2000 papers on temporal databases have been published (see the cumulative bibliography in [7] and recent surveys [8], [9], [10], [11]). Recently, the TSQL2 approach has consolidated many years of results into a single “consensus” approach [12], which (in revision as SQL/Temporal [13]) has been proposed to the ISO and ANSI standardization committees. Such database approaches are domain-independent, so that they can be profitably exploited also to model temporal data in medical applications. However, recently, some papers pointed out that the lack of specific supports makes the task of managing medical temporal data quite complex. For instance, O’Connor et al. implemented Chronus II, a temporal extension of the standard relational model and query language with specific features to make the treatment of clinical data more natural and efficient [14].

In this paper, we focus on temporal relational models and query languages, showing that current approaches have some limitations, so that relevant temporal phenomena in the medical field cannot be adequately modeled. In Section 2, we explicitly define the “point-based” semantics, and show that it is suitable to cope with a wide range of medical temporal data. However, a model based on a point-based interpretation of temporal data has severe expressive limitations when dealing with an important class of data, namely “telic” facts. In Section 3, we will substantiate this claim by considering an example in the medical field. We take TSQL2 as a representative example of data model and query language, but analogous problems arise in the other temporal relational approaches in the DB literature (since all these approaches assume a point-based semantics). We then generalize from the example, and settle the problem of dealing with telic (vs “atelic”) facts in a more general environment. Sections 4 and 5 describe our solution to such a general problem. In particular, in Section 4 we propose a new data model and semantics (namely, “interval-based” semantics) to cope with telic facts, and extend the query language (TSQL2) accordingly. In Section 5 we argue that both “standard” (i.e., based on “point-based” semantics) and “telic” (i.e., based on “interval-based” semantics) models are needed, and that the query language has to be extended in order to allow flexible “casting” operations to switch from one model to the other one, and vice versa. Finally, in Section 6 we describe alternative solutions and related works, and in Section 7, we draw some conclusions. While the telic data model has been already presented in [19], where we also proposed an extended three-sorted temporal algebra coping with both telic and atelic relations, in this paper we have widely explored the impact of the telic/atelic distinction on medical data, and we have extended the TSQL2 query language to cope with it.

2 Data Models and Data Semantics

As mentioned, many different database approaches have been devised in order to provide specific support to the treatment of time. Although there are remarkable differences between the alternative approaches, basically all of

them adopt the same data semantics: the data in a temporal relation is interpreted as a sequence of states indexed by points in time (see, e.g., the discussions in [15], [16], [8], [13], [17]). We will call such a semantics *point-based*, in accordance with the terminology adopted in artificial intelligence, linguistics and mathematical logic (but not in the database area, where “point-based semantics” has a different interpretation [17], [18] and is often used in relation to the semantics of the *query language* [19]).

It is important to clarify that in this paper we focus on data semantics, and we sharply distinguish between semantics and representation language; our distinction is analogous to the distinction between concrete and abstract databases emphasized by [9]. For instance, in many approaches, such as SQL/Temporal, TSQL2, TSQL, HQL, and TQuel, and Gadia’s Homogeneous Relational Model, a *temporal element* (a set of time intervals) is associated with each temporal tuple (or attribute), but this is only a matter of representation language, while the semantics they adopt is point based [17].

2.1 “Standard” point-based approaches to TDB

In this Section, we exemplify the semantics vs representation issue, and the point-based semantics, starting from a medical example.

Definition. Point-based semantics for data: The data in a temporal relation is interpreted as a sequence of states (with each state a conventional relation: a set of tuples) indexed by points in time. Each state is independent of every other state.

As an example, let us consider a database representation of the following situation, involving three different patients:

Example 1

Patient 1. Patient with chronic obstructive lung disease and hypoxemia, presenting with episodes of atrial fibrillation (AFI+), submitted to telemetry monitoring. At time 10:39 an episode of AFI+ is observed; after 5 minutes the episode is still observed, and also after 10 minutes the AFI+ persists, for 5 minutes.

Patient 2. Patient with chronic ischemic hearth disease, presenting with episodes of atrial flutter (AFL+), submitted to telemetry monitoring. At time 11:00 an episode of AFL+ is observed, after 5 minutes the monitor shows sinus rhythm, and after 10 minutes a new episode of AFL+ is present, lasting 4 minutes.

Patient 3. Patient with Wolf Parkinson White disease (WPW) and recurrent episodes of paroxysmal supraventricular tachycardia (PSVT+), submitted to telemetry monitoring. At time 11:15 an episode of PSVT+ is observed, after 5 minutes the sinus rhythm is restored.

The above data can be represented by relation $HRD1^A$ in the following, using the “interval-based” encoding of the validity time proposed by TSQL2 and taking minutes as the basic granularity [10]:

P_CODE	Type	VT
#1	AFI+	{[10:39-10:43], [10:44-10:48], [10:49-10:53]}
#2	AFL+	{[11:00-11:04], [11:10-11:13]}
#3	PSVT+	{[11:15-11:19]}

Figure 1: Relation $HRD1^A$

Alternatively, if a “point-based” encoding of the validity time is used (as, e.g., in BCDM [16], [12]), the same data can be represented as shown in relation $HRD2^A$.

The key point is that, if the “standard” point-based semantics is used, relations $HRD1^A$ and $HRD2^A$ represent (in different ways) exactly the same semantic content, i.e., the one shown in Fig.3. Notice that the point-based semantics in Fig.3 correctly cope with the temporal properties of the episodes in the example. In particular, *downward* and *upward inheritance* [20] are extremely relevant in this context. Roughly speaking, a property (or fact) P has downward inheritance if, from the fact that P holds through an interval of time I, one can correctly infer that P holds over all time points/intervals contained I; P has upward inheritance if, given any two overlapping or meeting time intervals I and J, the fact that P holds through I and J implies the fact it holds through $I \cup J$ (i.e., the union of

P_CODE	Type	VT
#1	AFI+	{10:39,10:40, ...,10:43,10:44, ...,10:53}
#2	AFL+	{11:00, ...,11:04,11:10, ...,11:13}
#3	PSVT+	{11:15, ...,11:19}

Figure 2: Relation $HRD2^A$

```

10:39 → < #1, AFI+ >
10:40 → < #1, AFI+ >
.....
10:53 → < #1, AFI+ >
11:00 → < #2, AFL+ >
.....

```

Figure 3: Point-based semantics of the relations $HRD1^A$ and $HRD2^A$

the two intervals). For instance, it is part of the semantics of episodes of “atrial fibrillation” that if patient #1 had AFI+ continuously from 10:39 to 10:43, from 10:44 to 10:48, and from 10:49 to 10:53 (with no interruptions):

- (i) #1 had AFI+ at time 10:40 (that is, a particular minute; thus, *downward inheritance* [20] holds);
- (ii) #1 had a 15-minute long episode of AFI+ (thus, *upward inheritance* [20] holds);
- (iii) X had just one episode of AFI+.

This fact clearly emerges if temporal queries about inheritance and countability are asked to relation $HRD1^A$ (or, alternatively, $HRD2^A$).

(1) Downward inheritance.

(Q1) *Who had AFI+ at 10:40?*

```

SELECT P.P_CODE
FROM HRD1^A AS P
WHERE P.Type = 'AFI+' AND
      VALID(P) OVERLAP '10:40'

```

Answer 1: {< #1|{10 : 39, ..., 10 : 53} >}

(2) Upward inheritance.

(Q2) *Who had one episode of AFI+ lasting more than 10 minutes?*

```

SELECT P.P_CODE
FROM HRD1^A (PERIOD) AS P
WHERE P.Type = 'AFI+' AND
      CAST(VALID(P) AS INTERVAL MINUTE) > INTERVAL '10' MINUTE

```

Answer 2: {< #1|{10 : 39, ..., 10 : 53} >}

(3) Countability.

(Q3) *How many episodes of AFI+ did #1 have?*

```

SELECT COUNT(P)
FROM HRD1^A (PERIOD) AS P
WHERE P.P_CODE=#1'
      AND P.Type = 'AFI+'

```

Answer 3: 1

Abstracting from the above example (see also the discussion in subsection 3.4), facts for which the upward and downward inheritance (roughly corresponding to “atelic facts” in Aristotle’s categorization) perfectly fit with the standard “point-based” semantics adopted by database approaches.

3 Limitations of Data Models Grounded on the Point-Based Semantics

In this Section, we show that “standard” database approaches based on point-based semantics are not adequate to deal with telic (medical) facts.

3.1 Telic facts

As above, we illustrate the basic issue with an example, but we stress that the same problems arise whenever data models using the point-based semantics are utilized to model a whole class of data (namely, Aristotle’s class of telic data [21]; see the discussion below).

Let us consider, e.g., drug intravenous infusion (henceforth, “i.v.” for short). In some cases, the administration event might stop suddenly (e.g., if the i.v. line falls out) and be resumed immediately. In other cases, two successive i.v. (to be distinguished) of a given drug may be prescribed to the same patient, with no time gap between them. In both cases, the two different i.v. infusions (again, with no temporal gap between them) must be recorded, since the process of restoring a phlebotomy requires medical staff intervention and is costly. On the other hand, notice that the biological effect of the drug is only slightly (if at all) influenced by a short interruption of the i.v. (except in few well-known cases). This is the reason why, at the level of granularity of minutes (that we choose for the whole clinical sample database), we model interruptions as instantaneous events (i.e., the interruption is simply expressed by stating that the i.v. ends on a time granule and re-starts in the next one).

From a technical point of view, if a patient X had two i.v. infusions of drug Y, one starting at 10:00 and ending at 10:50, and the other from 10:51 to 11:30 (all extremes included), we cannot say that:

(i) X had a (complete) i.v. at time 10:31 (that is, a particular minute; thus, *downward inheritance* [20] does not hold);

(ii) X had a one-hour-and-a-half-long i.v. of drug X (thus, *upward inheritance* [20] does not hold);

(iii) X had just one i.v. (i.e., i.v. events are *countable*, and must be kept distinct one from another).

Definition. In accordance with Aristotle [21] and with the linguistic literature, we term *telic* facts those facts that have an *intrinsic goal* or *culmination*, so that the three above properties do **not** hold, and **atelic** facts (e.g., “*patient X having AFI+*” facts for which all the three implications (i)-(iii) above hold [22].

In particular notice that upward and downward inheritance holds for atelic facts, but *not* for telic ones.

The importance of properly dealing with telic facts have been widely recognized in many different areas, spanning from artificial intelligence to philosophy, from cognitive science to linguistics [19] (see also the discussion in subsection 3.4).

3.2 Limitations of point-based approaches

Now, let us use a standard (i.e., point-based) temporal DB model to deal with i.v. infusion. For concreteness, we use the bitemporal conceptual data model (BCDM) [16] (which is the model upon which TSQL2 is based [12]), in which the validity time is denoted with sets of time points. (As an aside, even if we chose to use time intervals in the *representation* language, as in Fig. 1, the problem discussed below would still occur, due to the point-based semantics, as we’ll discuss shortly. Hence, the use of BCDM is not restrictive: the same problem arises for any data model that is based on the point-based semantics.)

For example, let us model the afore-mentioned patient X, who has code #4.

Example 2

Consider the following temporal relation $PHLEBO^A$, modeling also the facts that patient #4 had an i.v. of drug Z from 17:05 to 17:34, that patient #5 had two i.v. infusions of Z, one from 10:40 to 10:55 and the other from 10:56 to 11:34, and finally that patient #6 had an i.v. infusions of Z from 10:53 to 11:32.

This relation captures, among other facts, the fact that drug Y was given by i.v. to patient #4 from 10:30 to 11:30. Formally, this semantics can be modeled as a function from time points to the tuples holding over such time points (see Fig. 5).

On the other hand, this relation (its semantics) does not capture other relevant information, namely, the fact that there were two distinct i.v. infusions, one ending at 10:50 and another starting at 10:51. Such a loss of information becomes clear and explicit if temporal queries are considered, since, needless to say, answers must be provided on the basis of the *data semantics* (and independently of the *data representation*).

P_CODE	Drug	VT
#4	Y	{10:00,10:01,...,10:50,10:51,...,11:30}
#4	Z	{17:05,17:06,...,17:34}
#5	Z	{10:40,...,10:55,10:56,...,11:34}
#6	Z	{10:53,...,11:32}

Figure 4: Relation $PHLEBO^A$

10:00 → < #4, Y >
10:01 → < #4, Y >
.....
10:50 → < #4, Y >
10:51 → < #4, Y >
.....

Figure 5: Point-based semantics of the relation $PHLEBO^A$ in Figure 1.

3.3 Making semantic limitations explicit: queries

The most important problems arise, in our opinion, in case of queries involving *downward* and *upward inheritance* and *countability* of tuples. Again, we will use the TSQL2 query language, just to be concrete, but we stress that such a choice is not restrictive.

(1) **Downward inheritance** trivially holds on all data models based on point-based semantics, since the semantics implies the validity of tuples over each point in the validity time. Consider the following query over the relation $PHLEBO^A$ in Fig.4, where a relational table based on point-based semantics (see Fig.5) is used to model telic facts.

(Q4) *Who had one i.v. of Y at 10:10?*

```
SELECT P.P_CODE
FROM PHLEBOA AS P
WHERE P.Drug = 'Y' AND
      VALID(P) OVERLAP '10:10'
```

Answer 4: {< #4 | {10 : 00, ..., 11 : 30} >}

Notice, however, that although #4 *was having* and i.v. at time 10:10, it is not correct to infer that he *had* a (complete) i.v. at that time: the i.v. started at 10:00, and ended at 10:50.

(2) **Upward inheritance** holds on all data models based on point-based semantics. Since the semantics implies the validity of tuples over each point in the validity time, it implies the validity on the whole time interval covering all of them¹. This is not correct when dealing with telic facts such as i.v. infusion.

(Q5) *Who had one i.v. of Y lasting more than 60 minutes?*

```
SELECT P.P_CODE
FROM PHLEBOA (PERIOD) AS P
WHERE P.Drug = 'Y' AND
      CAST(VALID(P) AS INTERVAL MINUTE) > INTERVAL '60' MINUTE
```

Answer 5: {< #4 | {10 : 00, ..., 11 : 30} >}

Since patient #4's two i.v. infusions of Y cannot be distinguished at the semantic level, their validity time is "merged together", so that the above tuple is reported as output. Analogous problems arise when considering qualitative temporal constraints between validity times, such as, e.g., the "after" predicate in (Q6).

(Q6) *Who had an i.v. starting after one of the i.v. infusions of Y to patient #4?*

```
SELECT P2.P_CODE
FROM PHLEBOA (PERIOD) AS P, P2
WHERE P.P_CODE=#4' AND P.Drug='Y' AND
```

¹From the technical point of view, within temporal databases approaches, upward inheritance is obtained by performing *temporal coalescing* [23] over value-equivalent tuples

VALID(P) PRECEDES VALID(P2)

Answer 6: $\langle \#4|\{17:05, \dots, 17:34\} \rangle$

Notice that the tuples $\langle \#4|\{[10:51, \dots, 11:30]\} \rangle$, $\langle \#5|\{[10:56, \dots, 11:34]\} \rangle$ and $\langle \#6|\{[10:53, \dots, 11:32]\} \rangle$ are not reported as output, even if they follow one of the infusions of patient #4 (the one which ended at 10:50).

(3) Countability. Since there is no way to distinguish, at the semantic level, temporally contiguous value-equivalent tuples, contiguous telic facts are “merged together”, and one loses the correct count. Consider the following query.

(Q7) *How many i.v. did patient #4 have?*

```
SELECT COUNT(P)
FROM PHLEBOA (PERIOD) AS P
WHERE P.P_CODE='#4'
```

Answer 7: 2

In fact, in the point-based semantics, the validity time of interval of the first tuple of relation $PHLEBO^A$ is interpreted as the set of points $\{10:00, \dots, 11:30\}$.

3.4 Generalizing the problem

It is important to notice that these problems are not related to the *representation* language, but to the underlying (point-based) semantics. Indeed, several alternative *representations* are possible, each maintaining the same (point-based) semantics [12]. For instance, in TSQL2, an “interval-based” representation is used, as shown in Fig.6

P_CODE	Drug	VT
#4	Y	{[10:00-10:50],[10:51-11:30]}
#4	Z	{[17:05-17:34]}
#5	Z	{[10:40-10:55],[10:56-11:34]}
#6	Z	{[10:53-11:32]}

Figure 6: Alternative *representation* of the relation $PHLEBO^A$: the relation $PHLEBO2^A$

It is worth noticing that, even in case the relation $PHLEBO2^A$ were used instead of $PHLEBO^A$, the same results as above would be obtained. The same consideration also concerns the adoption of first normal form [12], in which each timestamp is restricted to be a period, with timestamps associated with tuples. As long as the underlying semantics is point-based, *each possible representation* of the (telic) event that patient #4 had two i.v. infusions of Y, one from 10:00 to 10:50 and the other from 10:51 to 11:30, is equivalent to the first tuple in $PHLEBO^A$, and conveys the same content shown in Fig.5, i.e., that patient #4 had an i.v. of Y in each time point within the whole span of time starting at 10:00 and ending at 11:30.

Moreover, it is worth remarking that, although until now we have showed the impact of neglecting the telic/atelic distinction on a specific medical example, problems such as the ones discussed above arise whenever *value-equivalent* tuples (i.e., tuples which are equal in their data part) concerning telic data have temporal extents that meet or intersect in time. This phenomena can occur in *primitive* relations, such as $PHLEBO^A$ and $PHLEBO2^A$ in Fig.4 and Fig.6, but also, and more frequently, in *derived* relations. For example, *projection* of a relation on a subset of its attributes (e.g., projecting the $PHLEBO^A$ relation over the P_CODE attribute only) usually generates several value-equivalent tuples, with possibly overlapping validity times. Consider, for instance, the relation in Fig.7 (obtained by query (Q8)), and remember that, as long as point-based semantics is used, the boundaries of meeting or overlapping validity times cannot be maintained by the semantics.

(Q8)

```
SELECT P.P_CODE
FROM PHLEBOA AS P
```

Finally, notice also that removing some of the attributes is not the only way of obtaining value-equivalent tuples (with, possibly, overlapping validity times) in derived relations. For instance, if multiple calendars and granularities are coped with [10], [24] switching from a finer to a coarser *temporal granularity* in the validity

P_CODE	VT
#4	{10:00,10:01,...,10:50,10:51,...,11:30,17:05,17:06,...,17:34}
#5	{10:40,...,10:55,10:56,...,11:34,}
#6	{10:53,...,11:32}

Figure 7: Relation $PHLEBO^A$ without the "Drug" attribute

time (e.g., from minutes to hours, or days; consider, e.g., [24]) can originate temporal overlaps that were not present in the primitive data.

3.5 An inter-disciplinary perspective of the telic/atelic dichotomy

In this subsection, we aim at looking at the telic/atelic dichotomy in a wider context, sketching some related issues in the fields of philosophy, linguistic and artificial intelligence, in order to demonstrate the generality of the problem. This subsection may be skipped by non-interested readers.

The distinction between "telic" and "atelic" facts dates back to Aristotle, and has been faced within different areas. In particular the subtle interplay between the telic vs atelic dichotomy and the point-based vs interval-based semantics has been studied in the area of linguistics and computational linguistics.

Within the linguistic community, it is commonly agreed that natural language sentences can be classified within different *aktionsart* classes (e.g., *activities*, *accomplishment*, *achievements and states* in [25]; also called aspectual classes [26]) depending on their linguistic behaviour or their semantic properties. These semantic properties demonstrate that the semantics of the association of facts to time depends on the classes of facts being considered. For example, [22] proposed the following semantic criteria to distinguish between states and accomplishments.

1. A sentence φ is stative iff it follows from the truth of φ at an interval I that φ is true at all subintervals of I (e.g., if John was asleep from 1:00 to 2:00 PM, then he was asleep at all subintervals of this interval: be asleep is a stative).
2. A sentence φ is an accomplishment/achievement (or kinesis) iff it follows from the truth of φ at an interval I that φ is false at all subintervals of I (e.g., if John built a house in exactly the interval from September 1 until June 1, then it is false that he built a house in any subinterval of this interval: build a house is an accomplishment/achievement) [22].

The property (a) for states has been often called downward inheritance in the TDB and AI literature (e.g., [24], [20]). Notice that also upward inheritance holds over states: if John was asleep from 1:00 to 2:00 and from 2:00 to 3:00, then he was asleep from 1:00 to 3:00. Obviously, the *aktionsart* distinctions above have a deep impact on the semantic framework one has to adopt to model the meaning of sentences and of the facts they describe. Point-based semantics evaluate the truth of sentences over time points (see also the subsection 2.1). This semantics perfectly works on stative facts: "John was asleep" in item (a) above is true exactly for all time points within 1:00 and 2:00 PM. On the other hand, point-based semantics seems to be inadequate to deal with accomplishments. For instance, given (b), there is no specific time point p such that "John built a house" is true in p . "John built a house" is true (or, in other words, occurred) exactly in the time interval from September 1 to June 1. This and analogous observations led most researcher in Linguistics, starting from the pioneering works in [27], [28], [22], to criticize point-based semantics, which is not adequate to deal with the semantics of accomplishments (while it works well for states and activities), for which an interval-based semantics is needed.

Different authors used different terminologies and models to deal with this phenomenon. For instance, [29] based their explanation on the fact that accomplishments are telic (from the Greek: "telos" meaning "goal") in the sense that they are characterized by the fact that they reach a culmination (goal, or telos), while states (and activities) are atelic (from the Greek: 'a' as a prefix indicates negation), i.e., do not have an intrinsic culmination. However, it is important to notice that Steedman emphasized that the about distinctions are not about verbs or verb groups, nor even about things that exist in the world, but rather about descriptions of the world [30]. Thus, these distinctions "...are conceptual tools of great usefulness in the philosophy of action, the philosophy of mind, in ontology generally, as well as in linguistics ..." [31]. Since "one of the most crucial problems in any computer system that involves representing the world is the representation of time" [32], this issue has had a significant impact on the

recent AI literature. In AI, many different techniques have been used in order to model the association of facts to time, such as, e.g., reification (see, e.g., [33], [34] and the criticism in [35]) episodic variables/ontological promiscuity (see, e.g., [36], [37]) and modal temporal logics (see, e.g., the survey in [38]). For example, Schubert and Hwang [36] introduced an episodic constant in order to represent explicitly any fact (termed episode) in the world. For instance, in Schubert's approach a fact such as "X had an i.v." could be represented as $\text{phlebo}(e1, X)$, where $e1$ is a constant that uniquely identify the fact. Thus, in such an approach, one can easily distinguish between facts of the same type (and with the very same description), even if they occur in meeting, overlapping or equal intervals of time.

Coming back to the core distinction between telic and atelic facts, it was first taken into account within the Philosophical community, dating back to Aristotle, from whom we derived the terminology. Going forward to recent philosophical approaches, Bach [39] pointed out that telic and atelic facts are somehow two complementary ways of representing reality. In particular, Bach showed that the dichotomy between atelic based view and the telic based view of the facts in the world is just a counterpart of the mass-nouns versus class-nouns dichotomy. In the same way as one can say that an object is composed by pieces of material (in turn, each piece of material could be conceived as a smaller object, at another level of granularity), a telic fact is composed by atelic ones.

These complementary ways of representing reality have also had a substantial impact on the AI community, where there is a long and still ongoing debate on whether it is better to model reality as a sequence of different states (atelic based representation)², or as a sequence of different events (telic based representation). For instance, McCarthy's Situation Calculus [40] is a typical example of the state based representation, while the Event Calculus [41] an example of the event based representation. The discussion of the relative merits of the two approaches would lead us far away the main goals of this paper. However, it is important to notice that, also in the AI field, many researchers have stressed the fact that the state based and the event based ways to represent reality are complementary, and in many cases one needs a flexible approach in which both ways can be adopted (consider, e.g., [33], [42], [34], [43], [44]). For example, in his seminal approach, Allen [33] distinguished among states, activities (termed processes) and accomplishments (termed events). In his first-order reified logic, Allen introduced three different predicates to associate facts to times, and used an axiomatic approach to model the downward inheritance property of states and the fact that accomplishments can be decomposed into activities. It is also important to remark that, in [45], [33], the truth of facts (represented by logical predicates) is evaluated over time intervals, and not over time points (i.e., an interval-based semantics is adopted). Following Allen's influential approach, many AI approaches chose to adopt time intervals as basic temporal primitives (cf., e.g., the surveys in [32], [46]).

In [20], Shoham has identified different temporal properties of facts, including upward and downward inheritance. It is worth noticing that the telic/atelic distinction is also present in the research about time oriented systems in medicine since several year, even though the term "telic" is not used and the distinction is not underlined in the data modeling and querying, but is considered at the application level (see, e.g [47] [48] [3] [49]).

Moreover, in the last years, the increasing need of sharing knowledge has motivated the appearance of approaches proposing high-level domain-independent ontologies (cf., e.g., the discussion in [50]). Many of these approaches included (at least) the above distinction between telic and atelic facts. A relevant example is the ontology devised within the CYC project, a project at MCC in Austin and Palo Alto started in 1989, which aims at encoding "the hundreds of millions of facts and heuristics that comprise human consensus reality" [51]. In such an ontology, they distinguish between processes (atelic facts) and events (telic facts) and model the fact that "Process is to Events as Stuff is to Individual Objects" [51].

4 Dealing with telic data

Despite the generality of the issue, the telic/atelic distinction has not had a specific treatment within the temporal database field yet. In the rest of the paper we propose a solution to the problem.

²Notice that, as discussed in [16], "the natural extension of a conventional relation to a temporal relation encodes states instead that events". In fact, using the point-based semantics (as in BCDM) the database collects a set of snapshots of the mini-world [16] it represents. In other words, the mini-world has a state-based representation, since it is represented as a set of states, one for each temporal snapshot (time point) in the database.

4.1 Telic data model

It is important to notice that the problems described in subsections 3.2-3.4 appear whenever a telic event (roughly speaking: an event which behaves as described by points (i)-(iii) in subsection 3.1: it has *no downward* and *upward inheritance* properties and it is *countable*) is modeled through a DB data model and query language which are based on the *point-based semantics* [52]. In order to deal with telic events (which respect the particular intervals, even if adjacent), a new data model and query language are needed, based on *interval-based semantics*³.

Definition. *Interval-based semantics for data:* each tuple in a temporal relation is associated with a set of time intervals, which are the temporal extents in which the fact described by the tuple occur. In this semantics the index is a time interval. Time intervals are atomic primitive entities, in the sense that they cannot be decomposed. Note, however, that time intervals can overlap; there is no total order on time intervals, unlike time points.

In our data model we introduce telic relations.

Definition. *Telic relation* A telic relation is a relation that must be interpreted using an interval-based semantics.

As an example, the relation in Fig.8 shows a telic relation $PHLEBO^T$ modeling our i.v. example.

Notice that, from the **representation** point of view, the relation $PHLEBO^T$ is identical to $PHLEBO2^A$ in

P_CODE	Drug	VT
#4	Y	{[10:00-10:50],[10:51-11:30]}
#4	Z	{[17:05-17:34]}
#5	Z	{[10:40-10:55],[10:56-11:34]}
#6	Z	{[10:53-11:32]}

Figure 8: Relation $PHLEBO^T$

Fig.6; the difference between $PHLEBO2^A$ and $PHLEBO^T$ is not one of syntax, but rather one of semantics. If an interval-based semantics is adopted, each interval is interpreted as an atomic (indivisible) one (see Fig.9).

It is important to stress that the crucial point is not the chosen **representation** (actually, relations $PHLEBO2^A$

$$\begin{aligned}
 [10 : 00 - 10 : 50] &\rightarrow \langle \#4, Y \rangle \\
 [10 : 51 - 11 : 30] &\rightarrow \langle \#4, Y \rangle \\
 &\dots\dots
 \end{aligned}$$

Figure 9: Interval-based semantics of the relation $PHLEBO^T$ in Figure 8.

and $PHLEBO^T$ are identical from the representation point of view), but the chosen semantics.

If a point-based semantics is used for a telic facts (such as i.v. infusion) relevant pieces of information are lost (e.g., the distinction between the i.v. of patient #4 ending at 10:50 and the one starting at 10:51; see Fig.5) while such information is preserved when switching to an interval-based semantics (see Fig.9, where the distinction between the two i.v. is maintained)

4.2 Extensions to the query language

The preceding subsection focused on extensions to a temporal model to add support for telic relations. We now show how these concepts can be added to an SQL-based temporal query language. As we'll see, only a few new constructs are needed. The specifics (such as using TSQL2) are not as important; the core message is that incorporating the distinction between telic and atelic data into a user-oriented query language is not difficult.

The first change is to support the definition of telic relations (the default is designated as atelic). This can be done with an "AS TELIC" clause in the TSQL2 CREATE TABLE statement. For example, in our extended TSQL2, the telic relation $PHLEBO^T$ can be defined as follows:

```
CREATE TABLE  $PHLEBO^T$  AS TELIC(MINUTE)
```

³This point, risen by [27], is now generality accepted within the linguistic and the AI communities (see e.g. [22], and discussion in subsection 3.5)

where minute is the granularity desired for the timestamp.

For telic queries, we propose the keyword “TELIC”. For example, the four queries (Q4), (Q5), (Q6) and (Q7) could all be correctly written as TELIC SELECT ... As an example, consider (Q5’) in the following.

(Q5’) *Who had one i.v. of Y lasting more than 60 minutes?*

```
TELIC SELECT P.P_CODE
FROM PHLEBOT (PERIOD) AS P
WHERE P.Drug = 'Y' AND
      CAST(VALID(P) AS INTERVAL MINUTE) > INTERVAL '60' MINUTE
```

Answer 5’: {}

Since $PHLEBO^T$ is a telic relation, an “interval-based” semantics is used. Therefore, the two i.v. infusions of Y of patient #4 are distinguished at the semantic level (their validity time is *not* “merged together”), so that no tuple is reported in output (since no i.v. episode is more than one-hour long).

Analogously, all the “telic versions” of (Q4), (Q6) and (Q7) provide the correct results. Specifically, (Q4) would return an empty result (since no complete i.v. episode occurred at 10:10); (Q6) would return two infusions for patient #4, one starting at 10:51 and one starting at 17:05, as well as one i.v. for patient #5 (starting at 10:56) and one for patient #6, starting at 10:53. (Q7) query would return a count of 3.

5 A three sorted model and its query language

While in Section 4 we have described our treatment of telic facts, in this section we describe our overall approach considering also atelic (and non-temporal) facts.

5.1 Need for atelic data model

Unfortunately, the telic model and query language in Section 4, taken in isolation, are not powerful enough to deal with all types of facts, and in particular, atelic facts. In Section 3, we argued that a telic data model is needed. Here we argue the reverse, that an atelic data model is also needed. In fact, both kinds of data must be expressible in a temporal model.

Using a solely telic model (and query language) to deal with atelic facts such as earning a given salary, owning a house, and so on, generate exactly the dual of the problems discussed in Section 3. Both downward and upward inheritance properties hold for atelic facts; not considering them causes loss of information. Consider, for instance, the atelic relation $HRD1^A$ in Section 2, and suppose that the same data were represented using a corresponding telic relation $HRD1^T$, i.e., through an operation of:

```
CREATE TABLE HRD1T AS TELIC(MINUTE)
```

In such a way, although the same syntax (e.g., the one used in Fig.1) can be used, an “interval-based” semantics is applied. Notice that, for instance, it is part of the intended meaning of atrial fibrillation (“AFI+”) that stating that patient #1 had AFI+ (without any interrupt) from 10:39 to 10:43, from 10:44 to 10:48 and from 10:49 to 10:53, implies that #1 had a 15-minute long episode of AFI+.

Such a semantic assumption (as well as those concerning downward inheritance and non-countability) are automatically captured if the data about $HRD1$ are represented by an atelic relation (i.e., by a relation based on a point-based semantics for data). On the other hand, such assumptions do no longer hold in case a telic relation (i.e., a relation based on interval-based semantics for data) such as $HRD1^T$ is used to represent the same data. This loss of information becomes even more evident if we ask queries on $HRD1^T$.

For instance, the query (Q2’) below gets no tuple in the answer, since, due to the interval-based semantics underlying the interpretation of telic relations (such as $HRD1^T$), the validity times from 10:39 to 10:43, from 10:44 to 10:48 and from 10:49 to 10:53, although continuous, cannot be merged together.

(Q2’) *Who had one episode of AFI+ lasting more than 10 minutes?*

```
TELIC SELECT P.P_CODE
FROM HRD1T (PERIOD) AS P
```

WHERE P.Type = 'AFI+' AND

CAST(VALID(P) AS INTERVAL MINUTE) > INTERVAL '10' MINUTE

Answer 2': {}

In summary, our data model supports both **telic relations** (to properly deal with telic facts) and **atelic relations** (to properly deal with atelic facts), as well as “standard” **atemporal relations** (to deal with non-temporal facts).

5.2 Need for flexibility: coercion functions

Furthermore, in the queries, coercion functions are useful in order to convert relations of the different sorts.

(Q9) *Who had one (complete) i.v., while patient #4 was having an i.v. of Y?*

As shown in Section 3, i.v. should be regarded as telic facts. However, when stating “*while patient #4 was having an i.v. of Y*” we look inside the fact, coercing it into an atelic one. Thus, this query involves two different ways of looking at the tuples in relation $PHLEBO^T$. First, the i.v. infusions of patient #4 must be interpreted as atelic facts, since we are not looking for i.v. infusions that occurred during one of patient #4’s infusions, but, more generally, *while patient #4 was having an i.v.* (i.e., we are interested in i.v. infusions occurred during [10:00-11:30] or during [17:05-17:34]). On the other hand, the i.v. infusions we are asking for must be interpreted as telic facts, since we look for *each complete occurrence* of them which is fully contained in [10:00-11:30]. For example, we want patient #6 in our output, since patient #6 had an i.v. from 10:40 to 10:55, regardless of the fact that patient #6 also had another i.v. from 10:56 to 11:34. We thus need more flexibility: although each base relation must be declared as telic or atelic, we need coercion functions (TELIC and ATELIC) to allow switch from one interpretation to the other at query time. Thus, in our extended TSQL2 query (Q9) can be expressed as follows:

```
TELIC SELECT P2.P_CODE
FROM PHLEBOT (ATELIC PERIOD) AS P,
     PHLEBOT AS P2
WHERE P.P_CODE=#4' AND P.Drug='Y' AND
      VALID(P) CONTAINS VALID(P2)
```

It is important to notice that, although the syntactic changes to TSQL2 are very limited, their semantic impact is very relevant. In particular, the adoption of coercion functions, in addition to the possibility of declaring relations both as “telic” and as “atelic”, involves a flexible approach, in which one can switch from “point-based” to “interval-based” (and vice versa) semantics at query time.⁴ As an example, in Fig.10 we describe how the result is obtained, step-by-step, from the query (Q9)

5.3 Examples

In the following, we present other examples of queries, to further substantiate the need of a flexible approach in which both telic and atelic models are used, and coercions in the queries are used for switching between the two at query time. Suppose we have a medical database containing, among others, the atelic relation $HRD2^A$ in Fig.2 (or $HRD1^A$, what has a different representation, but the same semantic content) and the telic relation $PHLEBO^T$ in Fig.8.

The query (Q10) involves both a telic and an atelic relation and can be expressed in our extended TSQL2 as follows:

```
(Q10) Who had an episode of AFI+ lasting more than 10 minutes, and that occurred before an i.v. of #5?
SELECT P.P_CODE
```

⁴In the area of Linguistics, it is widely recognized that although sentences, in their “neutral” form, can be classified as telic or atelic, natural languages provide different ways for switching from a telic to an atelic view of a sentence, or vice versa [53], [29]. For instance, Moens and Steedman [29] proposed a compositional approach to determine the telic vs. atelic aktionsart of a sentence on the basis of its verb, aspect and temporal adverbials. For example, in Moens and Steedman’s model, progressive form naturally applies to atelic facts (termed processes in their terminology). Whenever it applies to telic ones (culminated process), it coerces them to an atelic one, by stripping out their culmination. Thus, for instance, “Roger ran a mile” denotes a telic fact, and in “Roger was running a mile”, the progressive form coerces it into an atelic one. Analogously, the in adverbial naturally applies to telic facts (“John ate an apple in 2 minutes”). When applied to atelic facts, it adds them a culmination, turning them into telic ones. For instance, Moens and Steedman noticed that “John ran in four minutes” is a correct English sentence (denoting an accomplishment) in a context where John habitually runs a particular distance, such as a measured mile.

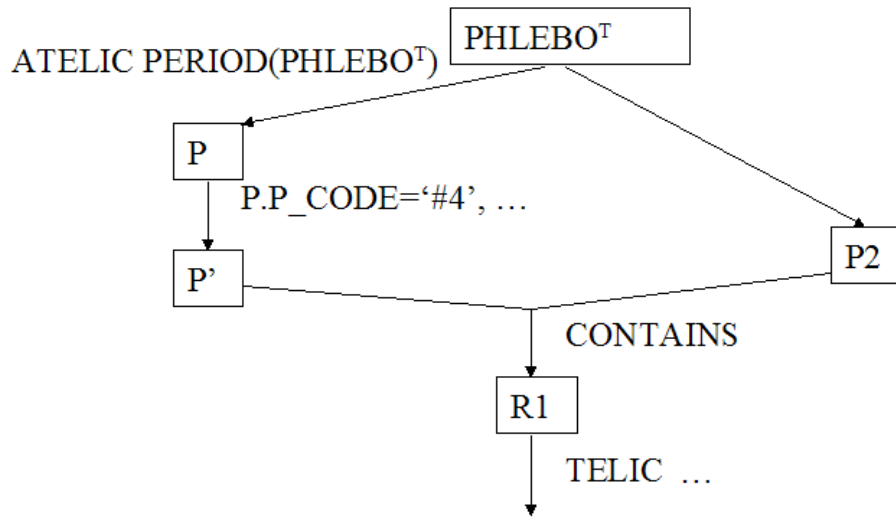


Figure 10: How the result is obtained, step-by-step, from the query (Q9)

```

FROM HRD2A (PERIOD) AS P,
  (TELIC SELECT P2.P_CODE
   FROM PHLEBOT (PERIOD) AS P2
   WHERE P2.P_CODE = '#5') AS P1
WHERE P.Type = 'AFI+' AND
      AND VALID(P) PRECEDES VALID(P1) AND
      CAST(VALID(P) AS INTERVAL MINUTE) > INTERVAL '10' MINUTE
Answer 10: { < #1 | { 10 : 39, ..., 10 : 53 } > }

```

This example shows the importance of having both atelic relations (in this case: $HRD2^A$), with temporal coercion, and telic relations (in this case: $PHLEBO^T$), on which temporal coercion must not be performed. In fact, being atelic, $HRD2^A$ is interpreted using the point-based semantics, so that a unique 15-minute long validity time is considered (from 10:39 to 10:53), as intended. Analogously, being telic, relation $PHLEBO^T$ is interpreted through the interval-based semantics, so that the two different #5's i.v.'s (from 10:40 to 10:55, and from 10:56 to 11:34) are not merged together. Therefore, patient #1 is in the result, since [10:39 - 10:53] precedes [10:56 - 11:34].

As a further example, let us consider the query (Q11):

```

(Q11) "Who had an episode of AFI+ during the time when patient #4 was having an i.v.?"
SELECT P.P_CODE
FROM HRD2A (PERIOD) AS P,
  (SELECT P2.P_CODE
   FROM PHLEBOT (ATELIC PERIOD) AS P2
   WHERE P2.P_CODE = '#4')(PERIOD) AS P1
WHERE P.Type = 'AFI+' AND
      VALID(P) DURING VALID(P1)
Answer 11: { < #1 | { 10 : 39, ..., 10 : 53 } > }

```

This query exemplifies the need of coercion from telic to atelic. In fact, we have an inner view ("was having an i.v.") of a telic relation ($PHLEBO^T$).

The progressive form in (Q11) tell us that we are not interested in an AFI+ contained in one of the i.v. episodes, but we want to look at i.v.'s as a "continuous" atelic facts. In other words, we do not want to maintain the distinction between the contiguous episodes (from 10 to 10:50 and from 10:51 to 11:30), but we want to coalesce them together. Such a coalescing is obtained through the coercion forced by the *ATELIC PERIOD* part of the nested query, which changes the telic table $PHLEBO^T$ into an atelic one. As a consequence, the two consecutive i.v.

episodes (from 10 to 10:50 and from 10:51 to 11:30) are merged together by the nested query, so that the AFI+ episode of patient #1 (starting at 10:39 and ending at 10:53) is fully contained into the resulting interval, and is reported in the output of the query.

On the other hand, if we were interested in AFI+ episodes which are during *one* of the i.v. episodes, no coercion would be needed in the nested query; in such a case, the result would be empty, since AFI+ is not contained in any of the three episodes of i.v. of patient #4 in relation *PHLEBOT*.

The central point is that the user needs to be aware of whether a relation is telic or atelic, and also how that data is to be manipulated. If the query language does not support this distinction explicitly, all manner of problems arise, as discussed in Sections 3 and 5.

6 Alternative solutions and related works

In subsection 3.5, we have sketched how the telic vs atelic dichotomy has affected the research in the area of linguistic and Artificial Intelligence. Although, to the best of our knowledge, the problem of properly modeling telic vs atelic (medical) data has not been specifically faced by any approach in the database area, in the following we sketch some alternative approach, that might appear to provide at least a partial solution to the problem.

First of all, it is worth mentioning that any approach trying to solve the problem on the basis of the *representation* formalism (and not on the *semantic* ground) cannot work. For example, since most of the problems discussed in Section 3 derive from the fact that, in an atelic model, overlapping or meeting validity times are “merged” (i.e. *coalesced*) together, one might try to solve the problem imposing a temporal first normal form (1NF) [15] as, e.g., in TSQL and in HQuel, so that just one time interval is associated with each tuple, instead of a set of intervals (i.e., a *temporal element* [15]). However, as long as one adopts the point-based semantics for data, this transformation alone does not solve the problem, since the coalescing of validity times of value-equivalent tuples is not evident at the representation level, but is present in the underlying semantics (and, of course, queries must be answered on the basis of the semantics of data, not of their representation syntax).

On the other hand, approaches that use 1NF as above, and never perform coalescing of value-equivalent tuples, as in SQL/Temporal, exhibit the same kind of problems discussed in subsection 5.1, since upward and downward inheritance would never hold. Basically, any “homogeneous” approach in which upward and downward inheritance hold on all relations (as in TSQL2), or do not hold in any relation (as in SQL/Temporal), will not be satisfactory⁵ Neither will approaches that have to fix a priori on which relations/attributes coalescing has to be performed and on which not, with no possibility of changing this property at query time (cf., e.g., [24]).

Also temporal interpolation techniques, that derive information for times for which no information is stored on the basis of related information holding at different times [16], could be useful. For example, Bettini et al. [24] proposed to explicitly associate with each table a specification of the assumptions on the semantics of temporal attributes (e.g., persistence of data), expressed in a formal language. At query time, such specifications are automatically merged with the user’s query in order to provide the correct results. Bettini also considered interval assumptions, including upward and downward inheritance, which however, are only studied in the context of evaluating the values of attributes whose validity time is expressed at different time granularities.

Finally, Chen and Zaniolo [54] use aggregate functions, such as length and contains, to perform telic operations on data assumed to be atelic. They also define the aggregate function *coales* to explicitly force upward inheritance on data assumed to be atelic.

Using temporal interpolation facilities or aggregate functions, it may be possible to deal with the telic/atelic distinction, but requires significant effort to fit this distinction into a formalism not designed with this purpose in mind. Instead, we feel that the atelic/telic distinction is so central that it should be given first-class state in both the data model and query language.

⁵For instance, an approach in which all relations are treated as telic (so that downward/upward inheritance is not directly supported) and queries are expressed in such a way that such inheritance is forced whenever atelic facts have to be considered (e.g., coalescing is performed, to deal with upward inheritance) is technically feasible, but extremely complex, and user-demanding.

An illuminating example of the additional complexity of expressing coalescing in standard SQL queries is given in [12].

7 Conclusions

In this paper, we have argued that current database approaches have some limitations, so that an important class of temporal medical data (i.e., *telic* data) cannot be properly represented, and we have proposed a new three-sorted model and a query language that overcome such limitations.

As regards future work, we envision the possibility of extending also the conceptual level (e.g., the entity-relationship model) to properly cope with *telic* (and *atelic*) facts. Moreover, we want to implement our approach and apply in GLARE (GuideLine Acquisition, Representation and Execution), a manager of clinical guidelines which strictly interacts with different databases [55], [56].

As matter of fact, the treatment of clinical guidelines constitutes a very interesting and challenging application for our approach. In fact, besides descriptions of the state of patients and of diagnoses (which typically atelic facts) clinical guidelines usually contain also a relevant number of (therapeutic or diagnostic) actions to be performed. In general, most of such actions are goal-oriented. Thus, they have an (explicit or implicit) goal or culmination, so that they behave like telic facts. As a consequence, both atelic and telic facts have to be appropriately managed by the databases supporting clinical guidelines applications.

Finally, although up to now we have only focused on the telic/atelic dichotomy, since it seems to us so prominent in the human way of perceiving and describing reality (see again Section 3.5), we also would like to consider the possibility of extending our approach to deal with other temporal proprieties, such as the ones pointed out by [20], temporal persistence, and interpolation functions [24].

References

- [1] Keravnou, E.T.: Special issue: Temporal Reasoning in Medicine. *Artificial Intelligence* 8(3) (1996)
- [2] Shahar, Y., Musen, M.: Knowledge-based temporal abstraction in clinical domains. *Artificial Intelligence in Medicine* 8(3) (1996) 267-298
- [3] Keravnou, E.T.: A Multidimensional and Multigranular Model of Time for Medical Knowledge-Based Systems. *J. Intell. Inf. Syst.* 13(1-2) (1999)73-120
- [4] Nguyen, J. H., Shahar, Y., Tu, S. W., Das, A. K., Musen, M. A.: Integration of Temporal Reasoning and Temporal-Data Maintenance into a Reusable Database Mediator to Answer Abstract, Time-Oriented Queries: The Tzolkin System. *J. Intell. Inf. Syst.* 13(1-2) (1999) 121-145
- [5] Combi, C.: Modeling Temporal Aspects of Visual and Textual Objects in Multimedia Databases. *Proc. TIME'00* (2000) 59-68
- [6] Snodgrass, R. T., Ahn, I.: Temporal databases. *IEEE Computer* 19(9), 35-42, (1986)
- [7] Wu, Y., Jajodia, S., Wang, X.S.: Temporal Database Bibliography Update. 338-366, in O. Etzion, S. Jajodia, and S. Sripada (eds.): *Temporal Databases: research and practice*, Lecture Notes in Computer Science, Vol. 1399, Springer Verlag (1998).
- [8] Roddick, J. F., Patrick, J. D.: Temporal semantics in information systems: a survey. *Information Systems* 17(3) (1992) 249-267
- [9] Chomicki, J.: Temporal Query Languages: A Survey. *Proc. of the International Conference on Temporal Logic*. Springer-Verlag (LNAI 827) (1994) 506-534
- [10] Özsoyoglu, G., Snodgrass, R. T.: Temporal and real time databases: a survey. *IEEE Transaction On Data and Knowledge Engineering* 7(4) (1995) 513-532
- [11] Jensen, C. S., Snodgrass, R. T.: Temporal Data Management. *IEEE Transactions on Knowledge and Data Engineering* 11(1) (1999) 36-44
- [12] Ahn, I., Ariav, G., Batory, D., Clifford, J., Dyreson, C., Elmasri, R., Grandi, F., Jensen, C. S., Kafer, W., Kline, N., Kulkarni, K., Leung, C., Loretzos, N., Roddick, J., Segev, A., Soo, M., Sripada, S.M.: The Temporal Query Language TSQL2 in R. T. Snodgrass (ed.), Kluwer Academic, (1995)

- [13] Snodgrass, R. T., Böhlen, M. H., Jensen, C. S., Steiner, A.: Transitioning temporal support in TSQL2 to SQL3, in O. Etzion, S. Jajodia, and S. Sripada (eds.): *Temporal Databases: Research and Practice*. LNCS 1399, Springer Verlag, 150-194, 1998.
- [14] O'Connor, M. J., Tu, S. W., Musen, M. A.: The Chronus II Temporal Database Mediator. *JAMIA, Symposium Supplement, AMIA Fall Symposium*, Isaak S. Kohane ed. (2002) 567-571
- [15] Gadia, S. K.: A homogeneous relational model and query languages for temporal databases. *ACM Transactions on Database Systems* 13(4) (1988) 418-448
- [16] Jensen, C. S., Snodgrass, R. T.: Semantics of Time-varying information. *Information Systems* 21(4) (1996) 311-352
- [17] Toman, D.: Point-based temporal extensions of SQL and their efficient implementation. in O. Etzion, S. Jajodia, and S. Sripada (eds.): *Temporal Databases: Research and Practice*, LNCS 1399, Springer Verlag (1998) 211-237
- [18] Böhlen, M. H., Busatto, R., Jensen C. S.: Point- Versus Interval-Based Data Models. *Proc. of the IEEE Int'l Conference on Data Engineering* (1998) 192-200
- [19] Terenziani, P., Snodgrass, R. T.: Reconciling Point-based and Interval-based Semantics in Temporal Relational Databases: A Treatment of the Telic/Atelic Distinction. *IEEE Transactions on Knowledge and Data Engineering* 16(5) (2004) 540-551
- [20] Shoham, Y.: *Temporal Logics in AI: Semantical and Ontological Considerations*. *Artificial Intelligence* 33 (1987) 89-104
- [21] Aristotle: *The Categories, on Interpretation. Prior Analytics*. Cambridge, MA, Harvard University Press.
- [22] Dowty, D.: The effects of the aspectual class on the temporal structure of discourse. *Tense and Aspect in Discourse. Linguistics and Philosophy* 9(1) (1986) 37-61
- [23] Böhlen, M. H., Snodgrass, R. T., Soo, M.: Coalescing in Temporal Databases. *Proc. of the International Conference on Very Large Databases* (1996) 180-191
- [24] Bettini, C., Wang, X.S., Jajodia, S.: Temporal Semantic Assumptions And Their Use In Databases. *IEEE Transactions on Knowledge and Data Engineering* 10(2) (1998)
- [25] Vendler, Z.: *Verbs and times* *Linguistics in Philosophy*, Cornell University Press, New York NY, 97-121, 1967.
- [26] Webber, B.L. (guest ed.): Special Issue on Tense and Aspect. *Computational Linguistics* 14(2), 1988
- [27] Bennet, M., Partee, B.: *Tense and Discourse Location In Situation Semantics*. Indiana University Linguistics Club, Bloomington (1972)
- [28] Dowty, D. R.: *World Meaning and Montague Grammar*, D. Riedel, Dordrecht, 1979.
- [29] Moens, M., Steedman, M.: Temporal Ontology and Temporal Reference *Computational Linguistics* 14(2), 15-28, 1998.
- [30] Steedman, M.: Verbs, Time and Modality *Cognitive Science* 2, 216-234, 1977.
- [31] Mourelatos, A.: Events, Processes and States *Linguistics and Philosophy* 2, 415-434, 1978.
- [32] Allen, J.F.: Time and Time Again: The Many Ways to Represent Time *International Journal of Intelligent Systems* 6(4), 341-355, 1991.
- [33] Allen, J.F.: Towards a General Theory of Action and Time *Artificial Intelligence* 23:123-154, 1984.
- [34] Galton, A.: A Critical Examination of Allen's Theory of Action and Time *Artificial Intelligence* 42, 159-188, 1991.

- [35] Bacchus, F., Tenenber, J., Koomen, J.: A Non-reified Temporal Logic *Artificial Intelligence* 52, 87-108, 1991.
- [36] Schubert, L.K., Hwang, C.H.: An Episodic Knowledge Representation For Narrative Texts *Proceedings First International Conference on Knowledge Representation and Reasoning (KR'89)*, 444-458, Toronto, Canada, 1989.
- [37] Hobbs, J.: Sketch of an Ontology Underlying the Way We Talk About The World in *[HCS, 95]*, 819-830, 1995.
- [38] Emerson, A.E.: Temporal and Modal Logic in: Van Leeuwen (ed) *Handbook of Theoretical Computer Science*, Vol. B, pp. 997-1072, Elsevier Science Publishers, 1990.
- [39] Bach, E.: The algebra of events *Linguistics and Philosophy* 9, 5-16, 1986.
- [40] McCarthy, J.: Programs With Common Sense *Semantic Information Processing*, L.M. Minsky (ed.), 403-418, MIT Press, Cambridge, MA, 1968.
- [41] Kowalski, R., Sergot, M.: A Logic-Based Calculus Of Events *New Generation Computing* 4, pp. 67-95, 1986.
- [42] Franconi, E., Giorgi, A., Pianesi, F.: Tense and Aspect: A Mereological Approach *Proc. 13th International Joint Conference on Artificial Intelligence*, Chambéry, France, 1222-1228, 1993.
- [43] Terenziani, P., Torasso, P.: Time, Action-Types and Causation: An Integrated Analysis *Computational Intelligence* 11(3):529-552, 1995.
- [44] Jixin, M, Hayes, P.: Addressing time intervals *Proc. IJCAI'05 Workshop on Spatial and Temporal Reasoning*: 29-38, 2005.
- [45] Allen, J.F.: Maintaining Knowledge about Temporal Intervals *Communications of the ACM* 26(11):832-843, 1983.
- [46] Vila, L.: A Survey of Temporal Reasoning in Artificial Intelligence *AI Communications* 7(1), 4-128, 1994.
- [47] Combi, C., Pozzi, G.: HMAP - A temporal data model managing intervals with different granularities and indeterminacy from natural language sentences. *VLDB J.* 9(4): 294-311 (2001)
- [48] Shahar, Y.: A Framework for Knowledge-Based Temporal Abstraction. *Artificial Intelligence* 90(1-2): 79-133 (1997)
- [49] Shahar, Y., Musen, M.A.: RÉSUMÉ: A temporal-abstraction system for patient monitoring. *Computers and Biomedical Research* 26(3) (1993) 255-273.
- [50] Guarino, N.: Formal Ontology, Conceptual Analysis and Knowledge Representation in *[HCS, 95]*, 625-640, 1995.
- [51] Lenat, D.B., Guha, R.V.: *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc project*. Addison Wesley Publishing Company, 1994.
- [52] Terenziani, P.: Is point-based semantics always adequate for temporal databases?. *Proc. TIME 2000*, Cape Breton, Canada IEEE Press (2000) 191-199
- [53] Verkuyl, H.J.: *A Theory of Aspectuality*. Cambridge University Press. 1993.
- [54] Chen, C. X., Zaniolo, C.: Universal Temporal Extensions for Database Languages *Proceedings of the IEEE International Conference on Data Engineering*, pp. 428-437, 1999.
- [55] Terenziani, P., Molino, G., Torchio, M.: A Modular Approach for Representing and Executing Clinical Guidelines. *Artificial Intelligence in Medicine* 23 (2001) 249-276
- [56] Terenziani, P., Montani, S., Bottrighi, A., Torchio, M., Molino, G., Correndo, G.: A context-adaptable approach to clinical guidelines. *Proc. MEDINFO'04*, M. Fieschi et al. (eds), Amsterdam, IOS Press (2004) 169-173