

REPRESENTING IMPLICIT AND EXPLICIT TIME RELATIONS IN NARRATIVE*

Lynette Hirschman and Guy Story

Linguistic String Project
New York University
New York, New York 10012

ABSTRACT

This paper describes a representation for time relations in narrative. The time relations are based on both explicit sources of time information (e.g., adverbial expressions or tense) and implicit sources, such as multiple reference to a single event, narrative time progression and earlier events implied by change of state words. Natural language processing is used to analyse the input text into a set of subject-verb-object units, connected by binary connectives; these units correspond to the events of the narrative. With each event is associated a time in relation to another event, adjusted by an optional time quantity. These time relations have a natural representation as a directed graph whose nodes are time points and whose edges are time intervals. The algorithm for extracting the time relations from a text is illustrated for an excerpt from a hospital discharge summary.

I Introduction

An understanding of time relations is essential to understanding narrative. For a computer program to "understand" narrative, it must process these time relations; this requires both a representation of the types of time information that interconnect narrated events, and an algorithm to extract this information from the natural language text.

The representation of time relations in narrative raises a number of complex problems. On the one hand, there are a multitude of time words: prepositions (e.g., during, before), adjectives (current, prior), adverbs (then, now), and nouns (duration, week). These various words combine into time expressions (generally adverbial expressions) which must be analysed and reduced to a systematic representation.

On the other hand, most events in a narrative do not have an explicit time stated in the form of an adverbial expression. Nonetheless the reader of

a narrative can determine the sequence of events even in the absence of explicit time information for each event. To do this, the reader makes use of certain characteristics of narrative discourse, in order to assign time relations to events whose time is not explicitly given. One characteristic of connected discourse is reference: by identifying various mentions of an event as referring to the same event, the reader associates these mentions with the same time. The conventions of narrative discourse also provide certain default time relations between consecutive clauses and/or sentences. For example, the convention of "narrative time progression" provides that in narrative, time does not move backward unless an explicit time marker is provided. An algorithm for capturing time relations in narrative must make use both of the explicit time information and of the time information conveyed implicitly by reference and by the default time relations. This paper will present such an algorithm and outline its application to a sample text.

II Background

There have been several previous investigations of the representation of time relations, all within the context of question-answering systems. [1] focused on the specification of the temporal information provided by tense and described a program CHRONOS that accepted a restricted set of English-like input, stored this information and answered simple time-related questions based on the input. Since the input was provided on a sentence-by-sentence basis, the issue of implicit time relations in narrative discourse did not arise in this work. [2] described a representation of time relations between sets of events. The representation was used in querying a data base consisting of events whose times were specified (sometimes incompletely) in terms of their duration and in terms of beginning and ending times in relation to another event. From this information, the program could determine whether an event preceded, overlapped, or followed another event.

More recently, [5] and [6] described a "time specialist" program which accepted time information about events in a LISP-like form; the specialist stored this information, checked it for consistency, and answered queries about the events described in the input. Although the time specialist handled only explicit time information provided in the input, there is a discussion in [5] about its ap-

* This research was supported in part by the National Library of Medicine grant number LM-02616, awarded by the National Institutes of Health; in part by the National Science Foundation grant number IST-7920788; and in part by the National Science Foundation grant MCS-8002453.

plication to a narrative-like situation, namely the specification of the time course of a disease and the matching of input data to the specified time course.

In contrast to the question-answering context of the work described above, this paper is based on research at the Linguistic String Project on the creation of a data base from natural language input. In our initial experiments using medical narrative, it quickly became clear that time relations between various medical events (e.g., symptoms, therapies) were of central importance in capturing the significant information. The time algorithm described here is a modification of an earlier algorithm [3] developed for the analysis of medical records, and the examples in the paper are drawn from actual hospital discharge summaries. The algorithm should, however, be equally applicable to time relations in other narrative domains.

In the automatic analysis of the medical narrative material, various stages of natural language processing preceded the analysis of time information. These stages included 1) a parse, identifying (surface) syntactic relations; 2) transformational regularization, mapping more complex sentence types into a "canonical set" of simple subject-verb-object sentence types; 3) a "formatting" component, mapping the analysed sentence into a tabular representation of the medical information. (See [7] for a more detailed description of the stages of natural language processing).

These earlier stages of processing are relevant to the time algorithm because the time algorithm assumes that each sentence is broken down into syntactic units consisting of subject-verb-object combinations corresponding to events; these units are connected by binary connectives, with the main clause preceding the clause subordinated or conjoined to it. Binary connectives include, among others, subordinate conjunction (e.g., while in She developed dyspnea while she was in Maryland); co-ordinate conjunction (e.g., and in Her elbow became infected and she was treated....); and relative clauses (e.g., She was treated with erythromycin which she took for 5 days.). As part of the syntactic analysis, each tense and each explicit time expression is identified and associated with the unit that it modifies. The output of the syntactic processing is thus a sequence of events with associated time information, connected by binary connectives. The time algorithm operates sequentially on the set of connected events obtained for a document; it starts at the beginning of the narrative and assigns to each event a time in terms of an already processed event, before going on to the next event.

The emphasis in developing this algorithm has been to use the available syntactic information to its fullest. We have not explored the question of whether the incorporation of domain-specific extra-linguistic knowledge would improve the processing of time (it almost certainly would), and whether the kinds of extra-linguistic knowledge needed to do this can be clearly delimited (a difficult task). It is our hope that the full utili-

zation of linguistic techniques will bring into sharper focus the kinds of information needed to supplement a linguistically based analysis.

III The Representation of Time

It is convenient to analyse the time of an event in terms of its relation to the time of another event (the "reference point") adjusted by an optional time quantity (see [6] and [3]). The reference point may be a date, which anchors the time to a fixed external scale, or it may be another event within the narrative. This analysis has a natural representation as a directed graph, where the nodes of the graph are the time points associated with the events and the edges represent the intervals between events. The time graph for a short three sentence textlet is shown in Fig. 1.

An event generally has a single time point associated with it, although events with duration have both a beginning and an ending point. If two events occur at the same time, they will both be associated with the same time point. Thus in the example of Fig. 1, if admission to the hospital occurs on 1/27 and the patient develops a fever on 1/27, then these two events will both be associated with the same time point (point A in the graph of Fig. 1).

The time graph is useful in determining the relative order of the events displayed in the graph: an event I precedes an event J iff there is a path from the time point associated with event I to the time point associated with event J. The sum of the edges along the path gives the time interval between I and J. For example, in the graph of Fig. 1, B precedes A. There are, in fact, two paths from B to A. In this case, the sum of the edges of the two paths must be equal: $BA = BC + CA$. (If the length of the two paths were not equal, an inconsistency would be detected). Since BA has a




FIG. 1 TIME RELATIONS REPRESENTED AS A GRAPH

length of 3 weeks, BC + CA = 3 weeks; therefore BC and CA are both less than or equal to 3 weeks. In this way it is possible to deduce further time information from the time graph, such as an upper limit for an edge length.

IV Explicit Time Information

We can distinguish three sources of time information in narrative: 1) explicit information, generally in the form of adverbial time expressions and tense; 2) time information derived from default values based on a small set of assumptions about how events within a narrative are connected temporally; and 3) time derived by the identification of an event with a previous mention of the same event (and its associated time).

The most obvious source of explicit time information is time adverbials expressions. These include dates (2/8/74) and time expressions that give the time of an event explicitly in terms of its relation to another event, e.g., during the first day of hospitalization. Figure 2 shows some sample adverbial phrases and tense. Note that in certain cases, the reference point may not be explicit; for example, in the adverbial phrase then (example 4, Fig. 2) the implicit reference point is taken to be the time of the last narrated topic (namely the preceding main clause). For tense, as for some adverbial expressions, the implied reference point must be made explicit; for example, for the present tense, it is the time of speaking.

The same analysis is used for the time information conveyed by a subordinate conjunction. For example, in the sentence Throat cultures were obtained only after IV ampicillin was started, the subordinate clause event (IV ampicillin was started) precedes the time of the main clause event (throat

cultures were obtained by an unspecified amount:

$$\begin{aligned} &\text{time (IV ampicillin)} = \\ &\text{time (throat cultures)} - ? \end{aligned}$$

Verbs (such as the verbs precede and follow) can also specify temporal connections between events. One important class of verbs is the class of aspectual verbs: verbs of beginning, ending, changing, etc. These verbs carry time information about a state in relation to an earlier implicit state: fever subsided implies that at an earlier point in time, a fever existed. Once the earlier implied event is located, the time relation between it and the event modified by the aspectual can also be recorded within the framework sketched above.

V Default Values for Time Connectives

There are several basic assumptions about how time progresses in narrative that enable the time algorithm to capture time relations implicit in the narrative. The most important of these assumptions is that time does not move backwards in narrative unless an explicit time marker is provided. The time of an event relative to the time of a previously mentioned event is assumed to be equal or later than the time of the previous event. The following sequence of sentences from Fig. 1 illustrates the usefulness of this hypothesis:

1. 3 weeks prior to the current admission, a bright red patch appeared under the patient's eye.
2. The patient developed a maculopapular rash over hands and knees.
3. On 1/27 she began having a fever to 104.

Sentence 1 describes an event occurring three weeks prior to admission; sentence 2 would have a default time greater than or equal to the time of sentence 1, using the hypothesis of narrative time

	TIME OR TENSE EXPRESSION	RELATION	REF-PT	ADJUSTMENT		
				DIR	Q	UNIT
1.	on admission	AT	ADM*	0		
2.	would (= future past)	AT	LAST NARRATED TOPIC	+		
3.	during the first day of hospitalization	FROM TO	ADM ADM	0 +	1	DAY
4.	then	AT	LAST NARRATED TOPIC	+		
5.	for the remainder of her hospitalization	FROM TO	LAST NARRATED TOPIC DISCH†	0 0		

* ADM = current hospital admission

† DISCH = discharge from current hospital admission

RELATION: AT indicates time-point
FROM } indicates beginning and ending
TO } points of a duration

REF-PT = REFERENCE POINT

DIR = DIRECTION Q = QUANTITY

FIG. 2 ANALYSIS OF EXPLICIT TIME EXPRESSIONS

progression. Sentence 3 has an explicit time (1/27), but narrative time progression also applies; this provides the information that the time of sentence 3 is later than or equal to the time of sentence 2. As a result, a lower and an upper bound on the time of sentence 2 are obtained from the discourse context, although sentence 2 has no explicit time. These relations are represented by the graph in Fig. 1.

A second hypothesis states that the time of a subordinate clause is equal to the time of its main clause, provided that no explicit time information occurs in the subordinate clause or in the subordinate conjunction. For example, in the sentence A spinal tap was performed which revealed 806 red blood cells, the subordinate clause (which revealed 806 red blood cells) is assigned a time equal to that of the main clause (A spinal tap was performed).

Co-ordinate conjunction poses some special and difficult problems for interpreting time information. The algorithm distinguishes two main cases. When the subject or object contains a conjoined noun or adjective phrase, the conjoined phrase is expanded during the transformational stage to yield two events; however, the algorithm assigns them the same time (specifically the second conjunct is assigned the time of the first conjunct). Thus in the sentence The maculopapular rash and the patch under her right eye gradually disappeared, the two events rash disappeared and patch disappeared are assigned the same time.

When there are two distinct predicates connected by a co-ordinate conjunction, then the two conjoined events behave like two independent sentences, and the algorithm assigns to the second a time later than or equal to the time of the first, as in narrative time progression. For example, in the sentence, The maculopapular rash and the patch under her right eye gradually disappeared and were gone after the first twelve hours on therapy, the default time information is used (together with the time information from the explicit adverbial after the first twelve hours on therapy) to determine that the time of the second conjunct (rash and patch were gone) is later than the time of the first conjunct (rash and patch gradually disappeared); in fact, the second conjunct marks the end point of the time period of the first conjunct.

VI Time Information from Resolution of Reference

If a single event is mentioned several times within a discourse, it is important to identify these mentions of the event and to equate the times associated with each mention (ignoring, for a moment, descriptions of states which can have a duration). For example, the events current hospitalization and current hospital stay have a special significance in a hospital discharge summary. They are associated with specific dates (provided in the header information) and appear as reference points in many time expressions, for example, she again spiked fevers on the 4th day of hospitalization, or patient was afebrile on discharge. If a given mention of admission,

discharge, or hospitalization can be associated with the current hospitalization, then its time can be set equal to the date given in the header. For example, if the hospital admission occurred on 1/27, then the 4th day of hospitalization will be 1/31, provided that this is a reference to the current hospitalization. Despite the fact that anaphoric reference is more difficult to detect in the note-taking style of the discharge summaries (articles, both definite and indefinite are commonly omitted), these references to the current hospitalization are generally easy to resolve because they are the default reference; there are occasional references to previous hospitalizations, but these are flagged by various modifiers, as in the previous hospitalization.

The original program implemented to process time relations [3] did anaphora resolution only when an event appeared as the reference point in a time expression. However, the current algorithm extracts more complete time information by comparing each new mention of an event (regardless of whether it occurs as a reference point) to previously mentioned events of the same type. When it identifies two mentions as referring to the same event, then it equates their times.

A simple technique is used for determining sameness of reference: a backward search is performed for events of the appropriate type, namely, events in the same synonym class. When a candidate event is found and its factuality is checked (negated events cannot be antecedents), then its time is compared to the time of the current event. If the times are compatible (that is, no contradiction arises from assuming that these events are the "same"), then they are taken as referring to the same event; otherwise, the backwards search is continued until all events of the appropriate type have been tested.

This procedure has been adequate to identify sameness of reference in the hospital discharge summaries, which have a relatively simple referential structure. For more complex texts, it is clear that a more sophisticated technique for resolution of anaphora would have to be used.

VII Capturing Time Relations in a Sample Text

The two sentences below are the third and fourth sentences (events 9-15) from a paragraph of a hospital discharge summary.

1. The patient became afebrile during the first day of hospitalization.
2. She again spiked fevers on the 3rd and 4th day of hospitalization, but then became afebrile and remained afebrile for the remainder of her hospitalization.

We will use these sentences to illustrate how the algorithm extracts and represents time information from a narrative. In the first sentence she became afebrile, there is an implicit reference (carried by the change-of-state verb became) to a state when the patient was not afebrile. This earlier implied state is entered as event E9; the state of becom-

ing afebrile is event E10. The second sentence states that she again spiked fevers on the 3rd and 4th days of hospitalization (events E11 and E12). Here, the word again is another reference to the earlier mention of fever (E9). Finally the word become in [she] became afebrile (event E13) is an implicit reference to the preceding fever events (she spiked fevers, E11 and E12). Each of these change-of-state words triggers an antecedent search as described above. Figure 3 shows the two sentences as a sequence of connected events, with their associated times.

In detail, the algorithm proceeds as follows:

E9: became afebrile in the first sentence is an implicit reference to an event be not afebrile not found by an antecedent search in the paragraph; therefore this implied event is entered as E9, with time preceding E10.

E10: There are two sources of time information for this event. The explicit adverbial expression during the first day of hospitalization gives rise to a beginning (FROM) time point = ADMISSION and an end (TO) point = ADMISSION + 1 DAY. (In this case, this represents the lower and upper bounds of a time-point, rather than the beginning and ending points of a continuous state; the current algorithm does not distinguish these two uses of durational expressions). Narrative time progression relates E10 to the main clause of

the previous sentence in the paragraph, namely E8 (not shown).

E11: The word again triggers a search for a preceding occurrence of fever, found in E9; this generates the time information AT E9 +. The explicit time expression on the 3rd day of hospitalization provides the time information AT ADMISSION + 3 DAY. The fact that E11 is derived from a separate predicate conjoined to E10 provides the third relation, time (E11) >= time (E10).

E12: There is an explicit time expression (on the 4th day of hospitalization); E12 also refers to the earlier fever (E9) by the use of again. Time progression from conjoined predicates is ignored here because it provides redundant information; an explicit time relation between E11 and E12 can be established by subtraction: time (E12) = time (E11) + 1 DAY.

E13: There are two sources of time information here: the adverb then, which provides the information that E13 occurred after the last narrated topic, namely E12 (time (E13) = time (E12) + ?). It also has time progression from conjunction with E12, but this is omitted because it is redundant; then provides more explicit information.

E14: There is an explicit duration expression here: for the remainder of her hospitalization. This provides a starting (FROM) time = time of last

TEXT (third and fourth sentences of a paragraph from a hospital discharge summary)

Date of admission 1/27/xx

Date of discharge 2/8/xx

... The patient became afebrile during the first day of hospitalization. She again spiked fevers on the 3rd and 4th day of hospitalization, but then became afebrile and remained afebrile and asymptomatic for the remainder of her hospitalization....

EVENT	CONNECTIVE	EVENT (TEXT)	TIME	RELATION	REF-PT	ADJUSTMENT
E9	CHANGE-OF-STATE = become	[patient be not afebrile]		AT	E10	-
E10		patient become past afebrile	during the first day of hosp.	[FROM TO AT	ADM ADM E8*	0 + 1 DAY 0/+
E11	CO-ORDINATE-CONJ = and	she spike past fever again	on the 3rd day of hosp.	AT AT AT	ADM E9 E10	+ 3 DAY + 0/+
E12		[she spike past fever again]	on the 4th day of hosp.	AT AT	ADM E9	+ 4 DAY +
E13	CO-ORDINATE-CONJ = and	she become past afebrile	then.	AT	E12	+
E14	CO-ORDINATE CONJ = and	[she] remain past afebrile	for the remainder of her hosp.	[FROM TO	E13 DISCH	0 0
E15	CO-ORDINATE CONJ = and	[she remain past] asymptomatic		[FROM TO	E14 E14	0 0

scope of conjunction indicated by braces {

[] indicate words filled in by transformational regularization

* Earlier sentences in paragraph contain events E1-E8; E8 = last narrated topic from preceding sentence (not shown).

FIG. 3 CONNECTED EVENTS AND THEIR TIME RELATIONS

narrated topic, namely E13; and it provides an ending (TO) time = date of discharge. The time derived from conjoined predicate time progression (time (E14) >= time (E13)) is therefore redundant and is omitted.

E15: This event is derived from E14 by expansion of a conjoined adjective in the object: afebrile and asymptomatic; it is therefore taken as having the same time as E14.


The time graph for these sentences is shown in Fig. 4.

In the course of computing the time relations, all events of a given type (for example, fever events) are collected into a set, in order to facilitate determination of sameness of reference. This grouping has an additional advantage in that

it makes it possible to graph a set of fever events against time. Figure 5 shows a graph of fever against time for the sentences analysed in Figs. 3 and 4. Although it is not possible to determine the exact shape of the curve (or its height, since a numerical temperature is not always given), the graph nonetheless provides an approximate picture of the patient's periods of fever during the hospital stay.

VIII Conclusion

In this brief discussion, only the broad outline of an algorithm to extract time information from narrative can be sketched. The algorithm described in this paper is a modification of a previously developed algorithm [3] that was implemented



F indicates the FROM time of an event with a duration;
T indicates the TO time of an event with duration.

FIG. 4 TIME GRAPH FOR AN EXCERPT FROM A HOSPITAL DISCHARGE SUMMARY

and successfully tested on a set of eight hospital discharge summaries, each containing approximately 50 sentences of text [4]. The original program contained an extensive treatment of explicit time expressions, covering some 50 time words. The modifications to the original program have been mainly in the handling of implicit time relations.

We are now in the process of implementing the new version of the algorithm and are still investigating several problems in this regard. These are: 1) maintaining a compact representation that permits detection of inconsistencies when a new time relation is added to the set of known time relations; 2) representing continuous states as opposed to time-point events: a state may change over time, but can still be identified as a single state (e.g., a fever may rise and fall but it is still referred to as a single continuing fever); 3) implementing a more sophisticated procedure for anaphora resolution, including the handling of references to continuous states; and 4) distinguishing between the two uses of durational expressions: 1) to define the lower and upper limits of a single time-point, as in she had a transfusion during her hospitalization; or 2) to define beginning and end points of a continuous state: she was afebrile during her hospitalization. In addition, there are a number of issues that we have not yet touched on, for example, the representation of imprecision in time specifications or the use of extra-linguistic knowledge to determine time relations. Nonetheless, the general principles outlined in this paper appear to constitute an important first step in extracting time information from narrative in a way that captures many of the important relations between the narrated events.

References

1. Bruce, B.C., A Model for Temporal References and its Application in a Question Answering Program, Artificial Intelligence 3, 1-25 (1972).
2. Findler, N.V. and D. Chen, Retrieval of Temporal Relations, Causality, and Coexistence, Intl. J. Comp. Inf. Sci. 2, 161-185 (1973).
3. Hirschman, L., Retrieving Time Information from Natural Language Texts, to appear in Information Retrieval Research, R.N. Oddy, C.J. Van Rijsbergen, S.E. Robertson and P. Williams, eds., Butterworths, London (1981).
4. Hirschman, L., Story, G., Marsh, E., Lyman, M., and N. Sager, An Experiment in Automated Health Care Evaluation from Narrative Medical Records, to appear in Computers and Biomedical Research.
5. Kahn, K., Mechanization of Temporal Knowledge, M.I.T. Project MAC Technical Report No. 155 (1975).
6. Kahn, K. and G.A. Gorry, Mechanizing Temporal Knowledge, Artificial Intelligence 9, 87-108 (1977).
7. Sager, N., Natural Language Information Formatting: the Automatic Conversion of Texts to a Structured Data Base, Advances in Computers 17, 89-162, M.C. Yovits and M. Rubinoff, eds., Academic Press, New York (1978).

Vertical lines indicate unknown amount of fever.

Horizontal lines indicate inexact time
(E13 is represented as a horizontal line rather than as a point because its exact time is not known).

Dotted line suggests possible curve for fever vs. time.

*The word spike implies fever goes up and down.




FIG. 5 GRAPH OF FEVER VS. TIME