

Association of Absolute Times with Events: Current Status and Future Challenges

Date: 21 July 2002

Author: Beth Sundheim

Reviewers: Lisa Ferro, Inderjeet Mani, James Pustejovsky

Introduction

This document is concerned with the current status of TimeML in relation to the handling of TIMEX3 expressions, and where it could go in the future. The document lists TIMEX3-related tasks in an order that approximately reflects increasing amounts of context and increasing degrees of difficulty. The ordering is not intended to indicate a prioritization on the basis of frequency of occurrence of the various types of expressions or context, as the relative frequency in a representative corpus has not been determined.

The list represents a sort of thought piece about developing a broad capability on the basis of TimeML to automatically associate absolute times with TimeML events. This particular document is not concerned with the relative ordering of events, but with the "timestamping" of events. The overall goal is to enable placement of events at a certain point (or pair of points, if the event has obvious duration) on a timeline. This capability will enable effective question answering about event times within and across documents.

Level of annotation in Timebank

The annotation in Timebank reflects the sorts of analysis covered by the current version of the TimeML specification and guidelines. Items 1, 2 and 3 on the list below are included in the current TimeML and are therefore represented in Timebank. The remaining items may be considered in the development of future versions of TimeML and Timebank.

Some of the remaining items will require only elaborations of the current guidelines concerning usage of the current schema, while others will require new ideas for representation in the annotation. A judgment about the difficulty of coming up with new representation solutions is shown for each item on the list. A judgment of the difficulty of implementation of each item is also made. Some items present challenges that are well beyond the state of the art.

Evolution from TIMEX2 to TIMEX3 and beyond

TIMEX2 and TIMEX3 treat essentially the same types of time expressions, but from different perspectives. TIMEX2 views them as the sole focus of attention -- everything that needs to be represented about a time, including its relation to another time, is treated as an attribute of the time expression itself. TIMEX3 views them as part of TimeML,

where the focus of attention is on events. TimeML offers several advantages over TIMEX2:

- TimeML permits more connections to be made between time expressions than was possible with TIMEX2. Some relations between time expressions that were implicit in TIMEX2 (e.g., the relation that causes the normalized value for "8 p.m." in the phrase "8 p.m. on Friday" to reflect the value that comes from the totally separate TIMEX2, "Friday") are now captured in a more explicit way via TLINK tags between an EVENT tag and a TIMEX3 tag (e.g., "They are meeting at 8 p.m. on Friday").
- TimeML defines a SIGNAL tag, which is used in part to identify the temporal prepositions that introduce a time expression. These prepositions, which are outside the extent of the TIMEX2/TIMEX3 tag, are referenced implicitly in TIMEX2 in order to distinguish durations from points, a distinction which is reflected in the form of the ISO notation used. TimeML makes the presence of the signal explicit.
- TimeML covers frequentative expressions such as "twice", which were not markable for TIMEX2, and offers a meaningful way to represent them, using the "cardinality" attribute on the MAKEINSTANCE tag.
- TimeML has "more to say" about the meaning of some time expressions, especially those that are anchored in events. With TIMEX2, an expression such as "20 minutes before the meeting" would be bracketed as a time expression but would have no assigned attributes. With TimeML, the annotation captures the fact that the interval is prior to the meeting and that its duration has the value of 20 minutes.

TIMEX3 is consistent with TIMEX2 not only in the range of expressions that they are concerned with, but also in the commitment to an extended ISO 8601 normalization. However, there are some differences in the way in which the two tags are defined; a few of the most obvious ones are outlined below:

- TIMEX3 does not currently use TIMEX2's SET, PERIODICITY and GRANULARITY attributes. However, some functionally similar mechanism to denote sets of times will be defined prior to the release of Timebank v1 in September.
- TIMEX3 does not use TIMEX2's NON_SPECIFIC attribute. This attribute, which identifies generic, indefinite and other nonreferential types of expressions, is not defined clearly enough in the TIMEX2 guidelines to result in good interannotator agreement, and its meaning and application in TimeML will need to be considered in discussions. For example, evidence for a specific temporal reference could be provided by a generic, as in "...a fit-looking Farkas said on state TV's *morning* talk show". Since the context of this example gave the date of the interview, it would be possible to narrow the "said" event down to the morning of that day.
- TIMEX3 has a defined "type" attribute, whose possible values are DATE, TIME and DURATION. Although the value of this attribute is predictable from the

form of the ISO value (and therefore there is no equivalent attribute defined for TIMEX2), it may be used in the future to permit an expression such as "2 months from today" to have an overall DATE type, although the matrix portion of the expression by itself expresses a DURATION.

- In contrast to TIMEX2, TIMEX3 does not currently permit the nesting of one TIMEX3 inside another. Since TimeML also does not currently permit TLINKs to connect one time expression directly to another, the consequence of disallowing nesting is that the embedded expression is not explicitly represented in the annotation. Thus, in "two weeks from next Tuesday" and "three years ago today", the embedded expressions, "Tuesday" and "today", respectively, are not tagged as distinct time expressions. The guidelines concerning these sorts of expressions will be revised before Timebank v1 is released, and the spec and guidelines will be updated to permit direct links between time expressions.
- With the exception of the types of expressions covered in the previous bullet, the extent of an expression in TIMEX3 excludes right modifiers. Although TIMEX2 and TIMEX3 differ in this regard, the difference is only superficial.
- TIMEX3 requires an anchorTimeID attribute for all indexical expressions. TimeML permits a variety of reference times to be identified in a document; thus, the annotation must specify which of the available reference times is to be used in normalizing any particular expression. The available reference times are annotated using the TIMEX3 functionInDocument attribute. In contrast, in the application of TIMEX2 to the ACE program's Relation Detection and Characterization (RDC) task, the simplifying assumption was made that the reference time for indexicals was to be found ultimately in the only time defined for the document, namely the "story reference time" (via a separate tag, STORY_REF_TIME).

Ordered list of tasks for associating absolute times with events

1. Goal: Generate ISO values for the unanchored durations and complete dates/times.

Analysis tasks: Recognize markable TIMEX3s, determine correct extent, identify TIMEX3s that convey unanchored durations and "complete", non-indexical dates/times.

Covered by TIMEX2 guidelines	Covered by TimeML/TIMEX3 guidelines	Difficulty of TimeML/TIMEX3 extension	Difficulty for algorithm development
Y	Y	N/A	Easy

The analysis tasks entailed by this item can be readily handled by current technology.

Durations can have 'X' placeholder in the numeric field, e.g., "PXM". Interpretation of the placeholder character is left for a later task (see item 5).

*Example expressions*¹: "[July 20, 2002]", "[04/02/1998 22:52:00]" (from a document header), "for [6 months]", "for [two and a half years]", "for [a few hours]"

- 2. Goal: Compute and generate the ISO value of simple indexicals.**
Analysis tasks: Identify "simple" TIMEX3s that are indexicals, determine correct reference time.

Covered by TIMEX2 guidelines	Covered by TimeML/TIMEX3 guidelines	Difficulty of TimeML/TIMEX3 extension	Difficulty for algorithm development
Y	Y	N/A	Easy-Med

"Simple" expressions are those that do not contain a nested time expression. The handling of "complex" TIMEX3s is addressed in other items in this document (see, for example, item 3).

For TIMEX2, these tasks are handled by programs such as TEMPEX, with the simplifying assumption that the reference time is always the "document time".

For TIMEX3, these tasks are handled in part by temporal functions associated with TIMEX3s and in part by temporal functions associated with events, as indicated below. The latter algorithms are only partially implemented at this time.

Example expressions (these are completely handled by TIMEX3 temporal functions): "[yesterday]", "on [Thursday]", "[last week]", "[eight o'clock]".

Example phrases containing more than one simple indexical expression (TIMEX3 temporal functions handle only part of the analysis; incorporation of the value of the more general expression into the value of the more specific expression is to be handled in TimeML by the event-triggered temporal functions): "at [eight o'clock] on [Thursday]", "[last week] on [Wednesday]"

- 3. Goal: Compute and generate the ISO value of complex indexicals.**
Analysis tasks: Identify complex indexicals in which the granularity of the "magnitude" (duration) expression is no finer than the granularity of the anchoring time expression. The anchor may be expressed directly or indirectly.

Covered by TIMEX2 guidelines	Covered by TimeML/TIMEX3 guidelines	Difficulty of TimeML/TIMEX3 extension	Difficulty for algorithm development
Y	Y/N	Med	Easy-Med

¹ In the examples given in the list, the extent of the time expression, as defined by TIMEX3 guidelines, is indicated by the use of square brackets.

Cases where the granularity of the magnitude expression is *coarser* than that of the anchoring expression are addressed later in this document (see item 10).

TIMEX2 adequately handles type (a) expressions, below, but not type (b). Type (b) expressions are tagged as TIMEX2s, but are not given any value.

TimeML covers both types of expression. Type (a) expressions are computed by temporal functions associated with TIMEX3s, while type (b) expressions are (or will be) handled by temporal functions associated with TLINKs. The spec and guidelines for type (a) cases will soon be revised to allow TLINKs between TIMEX3 tags. The treatment of type (a) cases will then be parallel to the treatment of type (b) cases.

Example expressions:

(a) Anchor is expressed directly: "[two days before Labor Day]", "[three years ago today]". These are currently treated as a single TIMEX3.

(b) Anchor is expressed indirectly: "He arrived [20 minutes] before the [8:00] meeting", "He arrived [the day] before the meeting on [Tuesday]". These are treated using TLINKs between each event and its associated time expression.

4. Goal: Compute the ISO value of the missing endpoint of a semi-bounded duration; generate a TLINK that expresses that value (?)

Analysis tasks: Recognize cases where an event instance (or multiple instances TLINKed in an IDENTITY relation) has a TLINK that expresses a BEGUN_BY or ENDED_BY relation *and* a TLINK that expresses a HOLDS relation, e.g., "Jan will speak at 8:30 for an hour"). Determine appropriate granularity for value of the missing endpoint.

Covered by TIMEX2 guidelines	Covered by TimeML/TIMEX3 guidelines	Difficulty of TimeML/TIMEX3 extension	Difficulty for algorithm development
N	N	Easy	Easy-Med

Given one endpoint and a duration, it should be possible to supply the other endpoint in order to anchor an event at both ends on the timeline, but TimeML currently does not permit TLINK to have an ISO value attribute (it has to obtain that value indirectly through a TLINK to a TIMEX3).

5. Goal: Anchor unbounded duration expression from general context, by either (a) computing and generating ISO value for the endpoint(s) or (b) adopting an existing ISO value from another time expression that refers to a time that's within the time range denoted by the duration expression.

Analysis tasks: Find non-local evidence (via network of SLINKs and TLINKs) that identifies (a) one or both endpoints of the duration, (b) an appropriate reference time that's somewhere between the endpoints of the duration, or (c) an appropriate reference time that the duration is "contained in".

Covered by TIMEX2 <u>guidelines</u>	Covered by TimeML/TIMEX3 <u>guidelines</u>	Difficulty of TimeML/TIMEX3 <u>extension</u>	Difficulty for algorithm <u>development</u>
N	N	Easy	Med

This task is focused on durations that have either a known or unknown length. The former type is indicated by a numeric in the value (e.g., "P3M") and the latter by a placeholder in the value (e.g., "PXM").

Given a network of SLINKs and TLINKs that allows one or both of the bounding subevents of a longer event to be identified, any timestamp information associated with the bounding events could be used to help define the endpoints of an event whose time is described only with an unbounded duration expression. Since this task depends on the existence of a nearly complete network, which could require sophisticated automatic event coreference capabilities that may never be achievable, the aim of this task is limited for practical purposes to those cases that would not require such sophisticated capabilities. Therefore, the difficulty of this task is judged to be "medium"; otherwise, it would be "medium-impossible".

Preferably, the analysis would permit the endpoints of the duration to be identified, but if not, it should permit some intermediate point in the duration to be identified in many cases. Presumably, the results of analysis would be represented in the form of new TLINKs. However, TimeML currently does not permit a TLINK to have an ISO value attribute (it has to obtain that value indirectly through a TLINK to a TIMEX3).

Example of identifying endpoints: "Two Russians and a Frenchman left the Mir and endured a rough landing on the snow-covered plains of Central Asia on [Thursday]. ... The two Russians arrived on the Mir [last August] ... Solovyov ... celebrated his 50th birthday during his [six-month] space voyage." (Note that this example may presuppose a more sophisticated event coreference capability than can ever be expected from an automated system.)

Example of identifying a containment time: "Two Russians and a Frenchman left the Mir and endured a rough landing on the snow-covered plains of Central Asia on [Thursday]. ... It took the crew [three hours] to descend from the Mir, ..."

6. Goal: Overlay fuzzy time unit values with precise values

Analysis tasks: Find evidence in general context for intended interpretation of one or both interval endpoints and/or one or more interval midpoints denoted by expressions of fuzzy time units (including elements such as "SU" within an otherwise precise ISO value, and also values such as "PRESENT_REF").

Covered by	Covered by	Difficulty of	Difficulty for
------------	------------	---------------	----------------

TIMEX2 <u>guidelines</u>	TimeML/TIMEX3 <u>guidelines</u>	TimeML/TIMEX3 <u>extension</u>	algorithm <u>development</u>
N	N	Med	Med

This "overlay" could be in the form of a document-level annotation defining the interval of the fuzzy unit.

As with item 5, this item presupposes a relatively complete network of links, which is an unrealistic expectation in the general case, but may be possible in certain kinds of cases.

Example in which one or more points (as in item 5, analysis tasks (a) and (b)) pertaining to a fuzzy unit can be identified: "She went to Germany [last summer]. She arrived in [mid-June]; by [mid-July], she was ready to go home." In this example, "mid-June" could be considered the start point of the interval; "mid-July" could be considered a midpoint.

- 7. Goal: Overlay "mod" attribute values with precise values**
Analysis tasks: Find evidence in general context for a more precise interpretation of modifier expressions that are represented using "mod".

Covered by TIMEX2 <u>guidelines</u>	Covered by TimeML/TIMEX3 <u>guidelines</u>	Difficulty of TimeML/TIMEX3 <u>extension</u>	Difficulty for algorithm <u>development</u>
N	N	Med	Hard

This "overlay" could be in the form of one or more new TIMEX3 attributes that permit expression of the precise meaning of the modifier, in the form of an IDENTITY type of TLINK between TIMEX3s that have the same intended value (assuming precise interpretation of the modifier is possible), and/or in some other form. Since the mod attribute conveys a variety of fuzzy things, the solution to this task could require significant extensions to TIMEX3/TimeML.

Examples of adding precision to modifier terms (modifier terms are in italics): "He left work in the [*early* afternoon]... He sneaked out of the office at [1:00] to go to the ball game." "The meeting ended at [*around* 11:00]... It was adjourned [15 minutes] before the cafeteria opened at [11:30] for lunch."

- 8. Goal: Given a set of event instances for a single event and a set of associated times, select the most precise one for use in applications such as timeline display.**
Analysis tasks: Recognize event instances that are TLINKed via an IDENTITY relation and that have different but "compatible" event-TIMEX3 TLINK types and TIMEX3 values, and determine the nature of the compatibility.

Covered by TIMEX2	Covered by TimeML/TIMEX3	Difficulty of TimeML/TIMEX3	Difficulty for algorithm
----------------------	-----------------------------	--------------------------------	-----------------------------

<u>guidelines</u>	<u>guidelines</u>	<u>extension</u>	<u>development</u>
N	N	N/A?	Easy-Med

This may be an application-dependent task, and may not require any explicit representation in the TimeML output.

Example: "He left work early on [Thursday]... His [1:00] departure went undetected." (If "left" and "departure" are in an IDENTITY type TLINK relationship, and if the value of "Thursday" is 2002-07-18 and the value of "1:00" is 2002-07-18T13:00, choose the latter value as the preferred timestamp for the event.)

9. **Goal: Create TLINKs between events that have coreferential TIMEX3s.**
Analysis tasks: On the basis of anaphoric TIMEX3 expressions such as "at [the time]", find candidate antecedents, such as "1990" and "tomorrow", and select the correct one.

<u>Covered by TIMEX2 guidelines</u>	<u>Covered by TimeML/TIMEX3 guidelines</u>	<u>Difficulty of TimeML/TIMEX3 extension</u>	<u>Difficulty for algorithm development</u>
N	N	N/A	Hard

The presence of an anaphoric time expression can serve as the basis for a search for a coreferring expression. The results can be used to add to the existing network of TLINKs in the document.

Example: "In its most recent affirmative action decision, [a year ago], the court split 5-4 in upholding a FCC policy that ... Thomas's views on the other major issues remain unknown, although the [1990] confirmation hearing for his seat on the Court of Appeals was widely viewed on Capitol Hill and in ideological interest groups as a rehearsal for a Supreme Court nomination. Some Democratic senators said at [the time] and again [Monday] that their votes to confirm him for the appeals court should not be taken as endorsements for the Supreme Court." The anaphor, "at [the time]", could be viewed as having two possible antecedents, "[a year ago]" and "[1990]". The task would be to select one and to perhaps create a second TLINK on "said" to show that its time is INCLUDED_IN the "hearing" event (which already has a TLINK that shows "1990" as its time).

10. **Goal: Generate vague value for complex indexicals with vague anchor.**
Analysis tasks: Recognize granularity of "magnitude" (duration) portion of complex indexicals that are anchored directly or indirectly to a time expression that has coarser (or vague) granularity.

<u>Covered by TIMEX2 guidelines</u>	<u>Covered by TimeML/TIMEX3 guidelines</u>	<u>Difficulty of TimeML/TIMEX3 extension</u>	<u>Difficulty for algorithm development</u>

N

Y/N

Med-Hard

N/A

Under TIMEX2 guidelines, both types of expression are tagged as TIMEX2s, but are not given any value.

If we assume that other tasks listed in this document have failed to identify an exact date value for "the end of April" and "the August meeting", then these complex expressions cannot be resolved to a particular date either. TimeML needs to develop a way to represent the vague endpoint of a duration.

Example expressions:

(a) Anchor is expressed directly, e.g. "[two days before the end of April]". This type is currently treated as a single TIMEX3.

(b) Anchor is expressed indirectly, e.g., "He arrived [two days] before the [August] meeting". These are treated using TLINKs between each event and its associated time expression.