

Chronoscopes: A theory of underspecified temporal representations

Inderjeet Mani
Georgetown University

Abstract

Representation and reasoning about time and events is a fundamental aspect of our cognitive abilities and intrinsic to our construal of the structure of our personal and historical lives and recall of past experiences. This talk describes an abstract device called a Chronoscope, that allows a temporal representation (a set of events and their temporal relations) to be viewed based on temporal abstractions. The temporal representation is augmented with abstract events called episodes that stand for discourse segments. The temporal abstractions allow one to collapse temporal relations, or view the representation at different time granularities (hour, day, month, year, etc.), with corresponding changes in event characterization and temporal relations at those granularities. A temporal representation can also be filtered to specify temporal trajectories of particular participants. Trajectories, in turn, can be intersected at various levels of granularity. Chronoscopes can be used to compare temporal representations (e.g., for aggregation, summarization, or evaluation purposes), as well as help in the visualization of temporal narratives.

1. Introduction

Representing and reasoning about time and events is a fundamental aspect of our cognitive abilities. Regularities in our experience, in the form of periodic/cyclic events, provide a basis for our systems for time reckoning, and are crucial in the scheduling of activities for a culture as well as a species. Representation and reasoning about time and events are, of course, intrinsic to our construal of the structure of our personal and historical lives and recall of past experiences. In all these different areas of existence, experiences in the world give rise to various inferences about events and their temporal organization.

Psychological research has shed some light on this inferential process. For example, in the case of visual perception, early work (Allport 1968) suggested that in inferring whether two event stimuli are simultaneous or successive, the inter-stimulus interval is the determining factor. Later, experiments by (DiLollo 1980) revealed that this inference is dependent on stimulus duration and persistence of the image in iconic short term memory. In research on auditory perception, (Jones 1989) found that estimates of temporal durations depend on prior expectations and particular modes of attending. Work by (Boltz 1995) showed that the structure of auditory events influences retrospective judgments of their duration. Research on consciousness by (Libet et al. 1979) indicated that ordering of sensations is influenced by latencies in stimulus signal propagation and the brain's tendency to adjust for this delay. This can result in sensations being reported as occurring in a different "subjective" order from their "objective" temporal order.

Unfortunately, there has been hardly any psychological research investigating temporal inference in language understanding. Natural language has a variety of devices to communicate information about events and time, including the use of finite verbs and event nouns, tense and aspect markers, and temporal adverbials, with these devices playing different roles in different languages. While such linguistic information is often present in natural language texts, they are often vague as to the precise temporal relationships. In (1), from (Dowty 1986), the ordering of the leaving with respect to the work and the hurrying, along with their temporal durations, is left unspecified.

(1) John hurried to Mary's house after work. But Mary had already left for dinner.

Further, crucial information about temporal relationships is often not overtly expressed in the text. In (2), temporal adverbials and aspectual information cue the inference that the twisting occurred during the running, while commonsense knowledge suggests that the pushing occurred before the twisting.

(2) Yesterday Holly was running a marathon when she twisted her ankle. David pushed her.

Finally, natural language may express information at different levels of temporal granularity. Mentions of events which have long durations may be interspersed with mentions of more punctual ones, as in (3), from (Ghosh 2005) (my italics):

(3) In the course of *the two decades* he spent in Kabul, Babur led four expeditions into India. His fifth and final *campaign* was launched in *October 1525*: it had a characteristically light-hearted *beginning*: "We mostly *drank* and had morning draughts on drinking days". Between *marches* Babur and his nobles *wrote* poetry, *collected* obscene jokes, and *gave chase* to the occasional rhinoceros. Despite internal *dissensions* the Lodis *managed to field* an army of 100,000 men and 1,000 elephants against Babur's paltry force of 12,000. The armies *met* on *April 20, 1526*, at the historic battlefield of Panipat a few miles north of Delhi.

In this paper, we describe an approach that allows one to represent and reason about time and events in the face of vagueness, missing information, and different granularities of temporal representation. Such a representation, called a Chronoscope, aims at capturing some of the subjectivity and looseness, as well as the flexibility inherent in our construction of temporal representations from natural language. The representation is simple and precise, and is an extension of representations that have been constructed automatically by recent information extraction systems (Verhagen et al. 2005). It is closely tied to recent frameworks for formal reasoning about time (Hobbs and Pan 2004) (Verhagen 2004). Unlike some other theories of granularity a/o abstraction, it also allows for a straightforward implementation and embedding in various temporal reasoning and visualization tools.

2 Representational Distinctions for Natural Languages

2.1 Underspecified Temporal Relations

Humans often find it difficult to infer fine-grained temporal relations from natural language texts. A pilot experiment (Mani and Schiffman 2005) with 8 subjects providing event-ordering judgments on 280 clause pairs revealed that people have difficulty distinguishing whether there are gaps between events. The 8 subjects were asked to distinguish whether an event is (a) strictly before the other, (b) before and extending into the other, or (c) simultaneous with it. These distinctions can be hard to make, as in the example of ordering *try on* with respect to *realize* in (4):

- (4) a. Shapiro said he tried on the gloves
b. and realized they would never fit Simpson's larger hands.

Agreement between 3 subjects was .5 Kappa, but .61 Kappa (acceptable) if the distinction between (a) and (b) was ignored. We can represent the temporal relations involved in (a) or (b) as a disjunction:

- (5) $ea \{ < m o \} eb$

Now, a disjunction set can be reified as a new relation. In temporal calculi, this notion has been explored in the work of (Freksa 1992). In particular, $\{ < m o \}$ corresponds to Freksa's relation *ob* (older and survived by), which has a semantics of the start (end) of the first interval preceding the start (end) of the second. Freksa defines 16 coarse-grained relations based on reified disjunctions involving the 13 fine-grained relations specified by Allen. There will be cases where the disjunction set involves a polarity change, i.e., an "inconsistency"; such disjunctions are not in Freksa's 16 coarse-grained relations.

More research is needed to determine which of the coarse-grained ones are more frequent in an annotated corpus. (Schilder 1997) and (Verhagen 2004) have described how relations can be arranged in a hierarchy based on relations between start and end points. In Section 3, we will show how the representation of different levels of coarse- and fine-grained relations can be captured by the chronoscope concept of Zooming.

2.2 Abstract Events and Temporal Discourse Structure

In addition to reification of temporal relations, it is possible to reify abstract events corresponding to entire discourses. Consider discourse (6), from (Webber 1988):

- (6) a. John went into the florist shop.
b. He had promised Mary some flowers.
c. She said she wouldn't forgive him if he forgot.
d. So he picked out three red roses.

It is clear that b and c comprise a sub-discourse in the main discourse. If we represent these discourses by abstract events (i.e., e_0 for the root discourse, e_1 for the sub-discourse), then individual events can be viewed as being temporally included (\subseteq) in their

immediate discourse, giving rise to a tree-structured representation of discourse (Mani and Pustejovsky 2004). In this tree representation, the dominance relation between nodes corresponds to temporal inclusion. The tree is unordered in terms of precedence relation, though by convention the nodes are ordered in order of mention. Abstract events representing sub-discourses will be called **episodes**. So, the temporal relations in (6) (we are of course ignoring the full logical form) is given by (7), where there is one episode e_1 (here variables are implicitly existentially quantified):

$$(7) e_a \subseteq e_0 \ \& \ e_1 \subseteq e_0 \ \& \ e_d \subseteq e_0 \ \& \ e_b \subseteq e_1 \ \& \ e_c \subseteq e_1 \ \& \ e_c < e_a \ \& \ e_a < e_d$$

States are represented as minimally included (\subseteq_{\min}) in the embedding event, without committing to whether the state extends before or after the event. Note that the ordering of e_b and e_c is left unspecified rather than reified, since the disjunction set $\{< >\}$ is inconsistent. For more details, and a comparison to Discourse Representation Theory (DRT), see (Mani and Pustejovsky 2004).

In some cases, the attachment of sub-discourses will be ambiguous, i.e., there will be more than one tree possible for a given discourse. As an example, consider the ambiguity which arises when the discourse (8), from (Hwang 1992), is extended with (9):

- (8) a Yesterday, Jack and Sue went to a hardware store
 b as someone had stolen their lawnmower.
 c She had seen a man take it and
 d had chased him down the street, but
 e he had driven away in a truck.

- (9) f. Later, they went to the police station.
 g. The police were not interested in such a minor crime.

Here e_f could be attached under e_0 or e_1 . We can therefore represent the temporal relations for the discourse (8-9) as follows:

$$(10) e_a \subseteq e_0 \ \& \ e_b \subseteq e_0 \ \& \ e_1 \subseteq e_0 \ \& \ e_c \subseteq e_1 \ \& \ e_d \subseteq e_1 \ \& \ e_e \subseteq e_1 \ \& \ e_b < e_a \ \& \ e_b < e_c \ \& \ e_c < e_d \ \& \ e_d < e_e \ \& \ e_g \subseteq e_f \ \& \ e_f \subseteq \{e_0 \ e_1\}$$

Such an underspecified representation of ambiguity based on factoring out common elements is similar in some respects to underspecified approaches used quite widely in natural language semantics.

3 Chronoscopes

An abstraction allows you to view information at different levels of granularity (w/o loss of information), and is based on research by (Hobbs 1984) (Giunchiglia & Walsh 1992) (Euzenat 1995) (Pianezi & Varzi 1996) (Mani 1998). **Temporal Abstractions** allow one to collapse temporal relations, or to zoom the representation to different time granularities.

Let E be a set of events and R be a set of binary temporal relations on E . A *Temporal Representation* is a relation $T \subseteq E \times E \times R$. A Temporal Representation spanning several years could be abstracted at different grain sizes, e.g., time units such as year, month, week, or day.

Let $U = \langle U_1, \dots, U_n \rangle$ be a sort of Time Units such that U_i *during* U_{i+1} , e.g., $\langle \dots, \text{YYYY-MM-DD}_1, \text{YYYY-MM}_2, \text{YYYY}_3, \text{YY-DE}_4, \text{YY}_5, \dots \rangle$.

Let us define the **time-granularity** for an event related by a temporal relation r to a time of granularity g :

$$(11) \forall e \forall r \forall g \text{ time-granularity}(e, r, g) \equiv \exists t \text{ calendarTime}(e, r, t) \ \& \ \text{coerceToUnit}(t, U_g)$$

Here, $\text{calendarTime}(e, r, t)$ means event e is in the temporal relation r to time t .

Let's now consider **equi-granular** (\sim_{rg}) events:

$$(12) e1 \sim_{rg} e2 \equiv \text{time-granularity}(e1, r, g) \ \& \ \text{time-granularity}(e2, r, g)$$

For example, in (3), *launched* $\sim_{\text{during-4}}$ *met*, i.e., *launched* $\sim_{\text{during-152X}}$ *met*. Thus, events whose temporal locations are in the same calendar decade of the same century will be in a common equivalence class corresponding to the decade position of that century.

Some example axioms used in a Chronoscope are shown below:

(13) *Older-and-survived-by at g:*

$$\forall x \forall y \text{ ob}_g(x, y) \equiv \text{before}_g(x, y) \vee \text{meets}_g(x, y) \vee \text{overlaps}_g(x, y)$$

(14) *Upward entailment of simultaneity:*

$$\forall x \forall y \text{ simultaneous}_g(x, y) \supset \text{simultaneous}_{g+1}(x, y)$$

(15) *Zoomed-simultaneity at g:*

$$\forall x \forall y x \sim_{\text{during-g}} y \supset \text{simultaneous}_g(x, y)$$

(16) *Granularity-invariance of relation r:*

$$\forall x \forall y \text{ invariant}(r) \equiv r_g(x, y) \equiv r_{g+1}(x, y)$$

These axioms index temporal relations to a particular level of granularity. This seems acceptable; after all, to say that things are simultaneous at certain perceptual thresholds doesn't mean they are physically simultaneous. Axiom (15) allows *launched* to be simultaneous with *met* at decade or higher granularity, i.e., even though *launched* is before *met* at lower levels of granularity. Axiom (16) allows one to drop the granularity subscript for a given relation r .

For any set of events X , let **characterization**(e, X) be true if e can represent X . In general, a characterization is an abstract event corresponding to the individual correlate of some proper subset of events in X . For example, an abstract event that is the individual correlate of the set of events e unique to a month of a particular year could serve as a characterization of the events X of that month. Or X might be characterized in terms of an abstract event from a background ontology. When a set of events X is characterized by an abstract event e , every temporal relation (or link) from an event y not in X to any event in X has to be rewired from y to e , and every link from any event in X to an event y not in X has to be rewired from e to y . In a characterization of X , temporal links among the events in X disappear from view.

With equi-granularity \sim_{rg} as the equivalence relation on a temporal representation s , we can create a hierarchy of partitions $\pi_{\sim_{rg}}(s)$. Let $Z_r(s) = \langle \pi_{\sim_{rg1}}(s), \dots, \pi_{\sim_{rgn}}(s) \rangle$ be a sequence such that $\pi_{\sim_{rgi}}(s) \leq \pi_{\sim_{rgi+1}}(s)$ for $1 \leq i < n$, where \leq is a refinement relation. We call $Z_r(s)$ a **Temporal Zooming**, as it permits zooming to any temporal grain size. Given an ordering of time units, a Zooming Z_r allows us to drill down to views of temporal representations based on fine-grained units as well as drill up to views based on coarse-grained units. We call Z_{rj} a *Zooming to grain j* . Thus, a Zooming to grain *Year*, i.e., Z_{r3} , will not look inside the months (instead, it will use characterizations for them).

Let us assume that $r = \{= o \subseteq mi\}$, where mi is ‘met by’. The resulting Zooming Z_R for the text (3) is as follows (here the curly parentheses are used to indicate the set of events (in the partition cell) that is characterized, with the calendar time associated with the events in the partition cell indicated as a subscript):

(17)

$Z_{r2}(s)$ (month grain):

1526-04: {₁₅₂₆₋₀₄₋₂₀ meet}

$Z_{r3}(s)$ (year grain):

1525: {₁₅₂₅₋₁₀ launch campaign drink}, march, write,
collect, give chase, dissent, manage, field

1526: {₁₅₂₆₋₀₄ {₁₅₂₆₋₀₄₋₂₀ meet}}

$Z_{r4}(s)$ (decade grain):

152X: {₁₅₂₅ {₁₅₂₅₋₁₀ launch campaign drink} march,
write collect give-chase dissent manage
field}, {₁₅₂₆ {₁₅₂₆₋₀₄ {₁₅₂₆₋₀₄₋₂₀ meet}}}

$Z_{r5}(s)$ (century grain):

15XX:
{_{150/1X} spend lead expedition},
{_{152X} {₁₅₂₅ {₁₅₂₅₋₁₀ launch campaign drink}
march write collect give-chase
dissent manage field}}

{ 1526 { 1526-04 { 1526-04-20 meet } } } }

It can be seen from (17) that $\pi_{\sim_{rgi}}(s) \leq \pi_{\sim_{rgi+1}}(s)$ for $1 \leq i < n$. Also, since the characterization of a singleton set allows for picking that element as its representative, this simplifies to:

(18)

$Z_{r2}(s)$ (month grain):

1526-04: 1526-04-20meet

$Z_{r3}(s)$ (year grain):

1525: { 1525-10 launch campaign drink }, march, write,
collect, give chase, dissent, manage, field

1526: 1526-04meet

$Z_{r4}(s)$ (decade grain):

152X: { 1525 { 1525-10 launch campaign drink } march
write collect give-chase dissent manage
field }, 1526meet

$Z_{r5}(s)$ (century grain):

15XX:
{ 150/1X spend lead expedition },
{ 152X { 1525 { 1525-10 launch campaign drink }
march write collect give-chase dissent
manage field } 1526meet }

As we will see later, there are a number of tasks where we may need to compare temporal representations. In order to do so, we can consider their intersection \cap , i.e., the events E and relations R in common. Once we have granularity in the picture, however, we can make the intersection sensitive to the granularity. The following theorem shows that it does not matter whether we zoom first and then intersect, or vice versa.

(19) *Distributivity of zooming over intersection:*

$$\forall s1 \forall s2 Z_{rg}(s1 \cap s2) \equiv Z_{rg}(s1) \cap Z_{rg}(s2)$$

Given a temporal representation, it can be filtered in various ways depending on the needs of a particular task. Filtering constraints can include (i) currency, e.g., current(e), or pre(e);(ii) focus, e.g., use focus(e) instead of curr(e) in above constraints; (iii) episode, e.g., the temporal representation of the ‘Mary’ episode in (9); (iv) participant, e.g., the temporal representation of events which involve Alexander the Great; (v) time constraint, e.g., the temporal representation of events on 9/11; (vi) temporal inclusion, e.g., the temporal representation of events which happened during the WTC attack, etc.

A temporal representation constrained based on event participants, as in (iv), is called a **trajectory**. A trajectory is the temporal path taken by an event participant (or sets of

them), e.g., the trajectory of Babur during his lifetime or during his last invasion of India. Trajectories can be abstracted at different levels of granularity, and can be intersected at particular levels of granularity. The intersections of trajectories are particularly interesting in the construction of biographical and historical narratives.

In a Chronoscope, temporal abstraction, filtering, and zooming can be composed together, allowing for considerable flexibility in dealing with the complexity of the cognitive information space.

4 Embodiment

It is often useful to be able to compare two temporal representations, in order to assess reliability of human annotation of temporal representations, to score a machine temporal representation against a human one, or to merge two different temporal representations of similar information for information extraction or summarization purposes.

A simple scorer for temporal links has been developed, based on the TimeML annotation scheme for events and their temporal anchors in text (Pustejovsky et al. 2005). Temporal Representations are directly represented in TimeML as TLINKs (temporal links) between events and between events and times, corresponding to the calendarTime relation described above. A scorer has been developed to compare TLINKs in a pair of documents (which are identical except possibly for TLINKs), before or after axioms for transitive closure have been applied. The scorer can easily be extended to allow for user-defined temporal relation equality predicates, e.g., Freksa's $pr=\{<, m\}$, $ob=\{<, m, o\}$. It could also be extended to allow for granularity parameter, so that at e.g., year grain, we can successfully match TLINKs with unequal relations, e.g., $(launch_{1525-10} < drink_{1525-10})$ and $(launch_{1525-10} > drink_{1525-10})$.

Chronoscopes can also be embedded in an annotation and visualization environment called TANGO (Pustejovsky et al. 2004), which is tied to TimeML. TANGO allows for editing and viewing of TimeML annotations by viewing events and their links in a grid-like display, where times are laid out horizontally and events aligned vertically.

TANGO allows for the selection of sets of elements in the display by boxing a region. Built in temporal filters (e.g., for (i) to (vii) above) can be expressed as queries on a temporal database underlying the TimeML representation, and can then be integrated with TANGO by applying those constraints to a boxed region. A list of named entities in a text, attached as arguments to TimeML event predicates, can be used very easily as a filtering mechanism to detect trajectories and their crossings along the lines of Figure 1.

Temporal Abstractions can be used in displaying more coarse- or fine-grained temporal relations as different types of links, and applying transitivity axioms based on the particular set of temporal relations. This is already carried out to some extent by the SputLink tool of (Verhagen 2005) that is being integrated with TANGO.

Temporal frames require partitioning the events related to a given event into a trichotomy; this can be displayed in TANGO, based on an early implementation of SputLink. Episodes require a layer of annotation on top of TimeML, as described in (Mani and Pustejovsky 2004). Once annotated, TANGO can be extended to display episodes by drawing a blob around a set of events and their links that constitute an episode. Zooming can be implemented by sorting the times into bins based on time units, and then constructing different granular representations based on their characterizations.

5 Related Work

This work derives from earlier work by (Hobbs 1984), who suggested that in the course of reasoning we conceptualize the world at different levels of granularity, and that in a particular reasoning process we distinguish only those things that are relevant to that process, making other things indistinguishable for all practical purposes. A theory in which temperatures are distinguished to the nearest degree, can, via an indistinguishability relation, be mapped to one in which temperatures are distinguished to the nearest 10 degrees (e.g., *in the fifties*). In any given situation, a granularity is determined, allowing the local theory to be selected. When the grain size shifts, certain “articulation axioms” are used to map to another local theory. In general, a mapping which induces a change in granularity can be considered a special case of an *abstraction* (Giunchiglia and Walsh 1992) who lay out a formal theory of abstraction.

Others who have discussed granularity shifts include (Euzenat 1995), who exploits complex relations which are disjunctions of other relations. Euzenat's approach postulates a number of fundamental algebraic properties of granularity operators. (Pianesi and Varzi 1996) have discussed degrees of temporal granularity in event structure. The analogue of zooming in their account is in terms of a “minimal divisor” on structures corresponding to sets of events, where temporal differences within the divisor are neglected. (Mani 1998) has discussed granularity shifts in terms of a theory of abstraction implemented as operators on logical forms.

6 Conclusion

The Chronoscope is a simple but flexible device that allows one to represent and manipulate temporal representations inferred from natural language texts. It is tied to existing annotation schemes and ontologies, and can also be integrated with annotation and visualization tools.

There are many challenges remaining to be addressed in this work, including the representation of sets of times like the events between marches in (3), generic events like the drinking in (3), and the zooming applied to events which are in the scope of modal operators or subordinated to other events (including the quotation in (3)). While TimeML and OWL-Time has a representation for some of these items, there is much remaining to be done here. Concise characterization in zooming remains a fundamental challenge, however, though work on event summarization and ‘event scripts’ is clearly relevant. Work on event durations, and the use of metric as well as qualitative temporal

constraints, as (Hobbs 2005) has been exploring, are also important to the further development of Chronoscopes.

While the representation of time described here strives to be human-centric and cognitively plausible, its psychological validity cannot be assessed until more is known about the psychological processes in temporal inference from natural language.

References

- Allen, J. (1984). Towards a General Theory of Action and Time. *Artificial Intelligence* 23, 123-154.
- Allport, D. A. (1968). Phenomenal Simultaneity and the Perceptual Moment Hypothesis. *British Journal of Psychology*, 59(4), 395-406.
- Boltz, M. G. (1995). Effects of Event Structure on Retrospective Duration Judgments. *Perception and Psychophysics* 57, 1080-1096.
- DiLollo, V. (1980). Temporal integration in visual memory. *Journal of Experimental Psychology*, 109 (1), 75-97.
- DiLollo, V., Hogben, J., & Dixon, P. (1994). Temporal integration and segregation of brief visual stimuli: Patterns of correlation in time. *Perception & Psychophysics*, 55, 373-386.
- Dowty, D. (1986). The effects of aspectual class on the temporal structure of discourse: semantics or pragmatics? *Linguistics and Philosophy* 9, 36-61.
- Euzenat, J. (1995). An Algebraic Approach to Granularity in Qualitative Time and Space Representation, *Proceedings of IJCAI-95*.
- Freksa, C. (1992). Temporal Reasoning Based on Semi-Intervals. *Artificial Intelligence*, 54(1):199–227.
- Ghosh, A. (2005). The Man Behind The Mosque. <http://www.amitavghosh.com/>
- Giunchiglia, F. and Toby Walsh, T. A (1992). Theory of Abstraction. *Artificial Intelligence*, 57, 2-3, 323-390.
- Hobbs, J. (1984). Granularity. *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-84)*.
- Hobbs, J, and Pan, F. (2004). An Ontology of Time for the Semantic Web. *ACM Transactions on Asian Language Processing (TALIP): Special issue on Temporal Information Processing*, Vol. 3, No. 1, March 2004, pp. 66-85.

- Hobbs, J. (2005). personal communication.
- Jones, M. R. (1989). Dynamic Attending and Responses to Time. *Psychological Review*, 96(3), 459 - 491.
- Kamp, H. and Reyle, U. (1993). Tense and Aspect. Part 2, Chapter 5 of *From Discourse to Logic*, 483-546.
- Libet B., Wright, E., Jr, Feinstein, B., & Pearl, D. K. (1979). Subjective referral of the timing for a conscious sensor experience: A functional role for the somatosensory specific projection system in man. *Brain*, 194, 191–222.
- Loftus, E.F. & Palmer, J.C. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behavior*, 13, 585-589.
- Mani, I., and Pustejovsky, J. (2004). Temporal Discourse Models for Narrative Structure. *ACL Workshop on Discourse Annotation*, Barcelona, Spain, 2004.
- Mani, I., and Schiffman, B. (2005). Temporally Anchoring and Ordering Events in News. In J. Pustejovsky and R. Gaizauskas, *Time and Event Recognition in Natural Language*. John Benjamins.
- Pianesi, F. and Achille C. Varzi, A. C. (1996). Refining Temporal Reference in Event Structures, *Notre Dame Journal of Formal Logic*, 37, 1.
- J. Pustejovsky, B. Ingria, R. Sauri, J. Castano, J. Littman, R. Gaizauskas, A. Setzer, G. Katz, and I. Mani. (2005). The Specification Language TimeML. In Mani, I., Pustejovsky, J., and Gaizauskas, R (eds.). (2005) *The Language of Time: A Reader*. Oxford University Press.
- Schilder, F. (1997). A Hierarchy for Convex Relations. *TIME* 1997, 86-93.
- Verhagen, M. 2004. *Times Between the Lines*. Ph.D. thesis. Department of Computer Science. Brandeis University.
- Verhagen, M., Mani, I., Sauri, R., Knippen, R., Littman, J., and Pustejovsky, J.. (2005). Automating Temporal Annotation with TARSQI. Demo Session. *ACL 2005*.
- Webber, B. (1988). Tense as Discourse Anaphor. *Computational Linguistics* 14(2): 61-73.