Accepted to appear in the *Proceedings of the 2nd Conference of the Pacific Association for Computational Linguistics (PACLING)*, Brisbane, April 1995.

The Temporal Structure of Narratives: A Semantic Approach

Mona Singh and Munindar P. Singh

MCC
3500 W. Balcones Center Drive
Austin, TX 78759-5398
USA

msingh@cs.utexas.edu, mona@mcc.com
(512)-338-{3269,3431} (voice)
(512)-338-3890 (fax)

October 27, 1994

Abstract

Computing the temporal structure of narratives is an important milestone in understanding stories and discourses. We propose structured representations for events based on their algebraic properties. We use these representations along with the situation type and aspect of the predicates involved to give detailed semantic postulates for various temporal connectives. These postulates—more detailed than previously proposed—are developed in conjunction with an analysis of a story from the New York Times. This story was proposed as a challenge to NL systems.

1 Introduction

Much research in Computational Linguistics has been recently focused on the structure of events and their semantic interpretation [4, 12, 14, 3, 10, 7, 20]. Determining the structure of events is a key prerequisite to solving the problem of constructing appropriate representations for them. Accurate and principled representations are in turn fundamental to understanding and generating natural language.

Recent work has shown that the temporal structure of an event and how it relates to other events in a discourse depends on its *situation type* and *aspect* [4, 12, 20, 19]. For example, the sentences *John walked in the park* and *John found a ball*, which differ in situation type, result in different representations—the former introduces a temporally extended event into the discourse, the latter a point. The significance of this can be seen by comparing the two discourses in which these sentences occur in different orders: if the order is as above, the obvious meaning is that John found the ball *while* walking in the park, but if the order is the opposite, the meaning is that John walked in the park *after* he found the ball. From the first discourse we can also infer that the location of the event of John's finding the ball was in the park.

Our approach takes a strongly semantic stance towards inferring the temporal structure of a narrative. Of course, semantics cannot preclude pragmatics, but complements it. There are some excellent approaches that derive most temporal relations from pragmatics [6, 10, 11]. Lascarides and Asher [10] propose a theory of discourse structure that uses the nonmonotonic logic *commonsense entailment* [2] as a basis for capturing pragmatic inferences. Their approach involves seven discourse relations: *Narration, Elaboration, Explanation, Background, Evidence, Consequence, and Contrast*, of which four are most relevant here. Narration entails that the descriptive order of events matches their temporal order; explanation and elaboration entail it doesn't. Background entails temporal overlap and conveys the pragmatic effects derived from the situation type (states normally provide background information).

In principle, all natural language inferences can be overridden by pragmatics or through accommodation [5]. However, it is important to derive as much information from the semantics as possible. This is because pragmatic reasoning is computationally complex. It is also demanding of the underlying knowledge representation, which must be able to answer queries about, e.g., whether an event potentially caused another. When the domain model is less that perfect, pragmatics may not successfully eliminate quite as many possibilities as one would like. By applying a strong semantics module, we are able to leave fewer options for the pragmatics to choose from. Indeed, pragmatics-based theories expect various "open sites"—where a sentence may

attach in a discourse—to be specified to them [10]. Our semantics-based approach provides precisely such sites.

To this end we are developing a general approach to (a) computing the situation type of a sentence and using it to determine the temporal structure of the situation it describes, (b) analyzing of temporal connectives such as *when*, *while*, *before*, *after* etc. and (c) using all this information to determine the temporal structure of the given discourse, i.e., the temporal locations of the situations with respect to each other. Here we focus on problems (b) and (c). Problem (a) was addressed in [17], which presents a formal approach for computing the situation type of a sentence. That approach applies uniformly to sentences with complex nominal expressions, locative and spatial adjuncts, and prepositional complements. Other related approaches (i) assume that the situation type of the basic verbal predicate is available, i.e., they cannot compute it themselves, and (ii) are almost exclusively procedural and based on a small example set [4, 12, 13].

Our approach is based on Krifka's algebraic semantics [8, 9] with some computational enhancements of our own [17]. Section 2 describes the necessary background concepts. Section 3 presents structures for events of different situation types. These are used in conjunction with aspect to determine the temporal relations among events. Section 4 presents formal semantic postulates for various important temporal connectives. These are developed hand-in-hand with an analysis of the famous *Mr. Hug* story, which given as a challenge to NL systems by John McCarthy.

2 Background

We now discuss situation type and aspect; our underlying algebraic semantics; and its application in classifying situation types.

2.1 Situation Types and Aspect

Each sentence describes a *situation*, i.e., a *state* or an *event* [21]. States, e.g., *be tall*, are homogeneous. Events may be *activities*, e.g., *walk in the park*, any part of which is also walking in the park, *achievements*, e.g., *win a race*, which are instantaneous, and result in a change of state, or *accomplishments*, e.g., *build a house*, which are characterized by a process and its culmination. Often, achievements and accomplishments, which have natural final endpoints, are called *telic* events and activities, which have arbitrary final endpoints only, are called *atelic* events.

Aspect is best defined as the viewpoint of a speaker towards a situation [18]. Two kinds of aspect are traditionally considered. The *perfective* describes a situation as a complete whole, e.g., *Al ate an apple.* The *imperfective* can be the *progressive* or the *habitual*; we consider only the former for reasons of space. A progressive sentence presents an event internally, e.g., *Al was eating an apple.* For languages like Japanese, Hindi, and Chinese, we have argued for the introduction of the *neutral perfective*, but we don't discuss it here for reasons of space [16].

2.2 Algebraic Semantics

In Krifka's algebraic semantics, *events* and *objects* are distinct sorts of entities, each structured as a join semi-lattice without a bottom element. $\sqcup$ is the operation of join, and $\sqsubseteq$ and $\sqsubset$ the corresponding relations of part and proper part, respectively. Thematic relations are defined as mappings of objects to events. The semantics of cumulative and quantized reference can be given in this framework. For example, both *beer* and *apples* are cumulative predicates, since beer combined with more beer is still beer and, adding more apples to apples yields apples. On the other hand, *a glass of beer* and *five apples* are quantized predicates. Suppose there are two distinct entities to which the predicate *a glass of beer* applies. This predicate then cannot apply to their join. No part of *five apples* is *five apples.* The definitions below are due to Krifka [8]. Note that $\mathsf{CUM}(P) \rightarrow \neg\mathsf{QUA}(P)$, if $|\{x : P(x)\}| \geq 2$.

- **Cumulativity** characterizes atelic predicates; e.g., CUM(*Al drank beer*) holds.
  $\forall P[\mathsf{CUM}(P) \leftrightarrow \forall x, y[P(x) \wedge P(y) \rightarrow P(x \sqcup y)]]$
- **Quantization** characterizes telic predicates; e.g., QUA(*Al ate an apple*) holds.
  $\forall P[\mathsf{QUA}(P) \leftrightarrow \forall x, y[P(x) \wedge P(y) \rightarrow y \not\sqsubset x]]$

An *event predicate*, i.e., a predicate computed from a clause that applies to events, is by default taken to apply to events in the perfective aspect. Other aspects are defined as operators on event predicates, which yield predicates in the appropriate aspect. Thematic relations are homomorphisms from objects to events that preserve $\sqsubseteq$. Some properties, proposed by Krifka [9], are useful:

- **Uniqueness of objects**: R relates every event to a different object; e.g., the eating of an apple is related via the PATIENT role to a specific apple.
  $\forall R[\mathsf{UNI\text{-}O}(R) \leftrightarrow \forall e, x, x'[R(e, x) \wedge R(e, x') \rightarrow x = x']]$

- **Mapping to objects**: R maps every subevent to a proper part of the object, e.g., every part of eating an apple involves eating a part of an apple.

  $\forall R[\mathsf{MAP\text{-}O}(R) \leftrightarrow \forall e, e', x[R(e,x) \wedge e' \sqsubseteq e \rightarrow \exists x'[x' \sqsubseteq x \wedge R(e',x')]]]$

- **Mapping to events**: R goes the other way; e.g., it maps every part of an apple to a part of the event of eating it.

  $\forall R[\mathsf{MAP\text{-}E}(R) \leftrightarrow \forall e, x, x'[R(e,x) \wedge x' \sqsubseteq x \rightarrow \exists e'[e' \sqsubseteq e \wedge R(e',x')]]]$

- **Graduality**: R is such that the object is subjected to the event gradually. For example, *writing a letter* or *eating an apple* affect their objects gradually, while *seeing a cat* or *finding a watch* do not.

  $\forall R[\mathsf{GRAD}(R) \leftrightarrow \mathsf{UNI\text{-}O}(R) \wedge \mathsf{MAP\text{-}O}(R) \wedge \mathsf{MAP\text{-}E}(R)]$

2.3 Algebraic Classification of Situation Types

We summarize the exact procedure for determining the situation type of a sentence (for the detailed analysis see [17]). The situation type of an event predicate depends on both the verb and the nominal patient.

Verbs such as *eat, drink*, and *consume* result in activities when combined with cumulative predicates and in accomplishments when combined with quantized predicates. When these verbs have a direct object, their thematic relation of PATIENT is gradual. For example, joining two events of *drinking beer* yields an event of *drinking beer*, i.e., the event predicate is cumulative. Also, the object is consumed gradually. Accomplishments are gradual and quantized; e.g., no part of *drink a glass of beer* is *drink a glass of beer*, but the glass of beer is drunk gradually.

- **Activities:** [−QUA,+GRAD→ ACTIVITY]
- **Accomplishments:** [+QUA,+GRAD→ ACCOMPLISHMENT]
- **Achievements:** [+QUA,−GRAD→ ACHIEVEMENT]

3 Structured Representations for Events

We propose that situation types be understood as complex schematic representations. In each schema, I is the initial endpoint, i.e., the point at which the event begins. F_N is the natural final endpoint of the event, i.e., the point where the event would end

naturally—this is defined only for achievements and accomplishments. E refers to the middle part of the schema—it is the main event for accomplishments and activities, and the preliminary processes for achievements. F_A is an arbitrary final endpoint of E. The events e_i are the *subevents* of E. The perfective aspect introduces the entire schema into the discourse: succeeding sentences or clauses may be placed in any part of it. The progressive aspect introduces only the part of the schema labeled E.

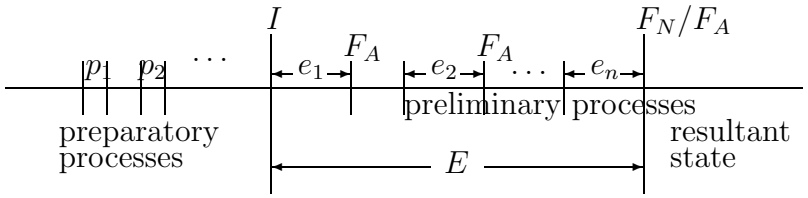

Figure 1: The Schema for Accomplishments and Activities

Schema for Accomplishments and Activities The accomplishment "to build a bridge" can be seen as having subevents of building from 8 am to 5 pm every weekday, or the subevents of setting up the iron bars, adding the concrete, and so on. The preparatory process refers to subevents like getting a loan and drawing up the plans that are prerequisites to the main event. The resultant state refers to the state that holds after the bridge has been constructed. Activities are a lot like accomplishments in their durative nature, but they do not have any natural endpoints and can end at any arbitrary point in time.

Schema for Achievements Achievements differ from accomplishments in that they are instantaneous. Importantly, their preliminary processes are detachable. For example, "spot a cat" does not require any prior activity, whereas "win a race" requires that the agent do at least some running before he can win. But, even in the presence of preliminary processes, an achievement can be predicated only of its final endpoint. For example, a proper subevent of winning a race is not a winning of a race.

The above schemata distinguish between the *preparatory* and the *preliminary* processes of an event. This distinction is crucial for obtaining accurate representations. In (1) below, the initial clause is in the perfective aspect; in (2), the initial clause is in the progressive aspect.

1. When John won the race, he took steroids.

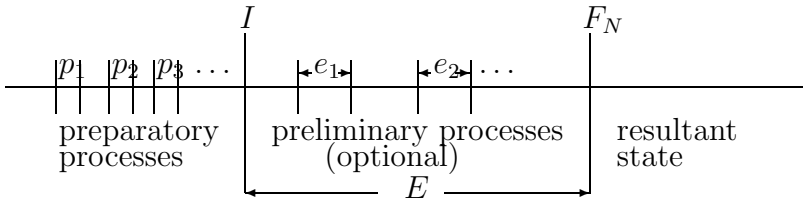

Figure 2: The Schema for Achievements

2. When John was winning the race, he took steroids.

Example (1) suggests that John took steroids before the actual running commenced, i.e., during the preparatory process. This interpretation is most easily obtained if one invokes the pragmatic rule of enablement; i.e., it was taking steroids that made it possible for John to emerge victorious [11]. In (2), however, the taking of the steroids is forced to have the interpretation of having taken place during the running, i.e., during the preliminary process, because the preparatory process is not brought forth on the DRS by the progressive aspect. This is a good example of a semantic restriction on temporal reference. In this case, the pragmatic rule of causation of [10] cannot successfully justify interpreting the steroids-taking event as being before the running event. This is important because the notions of preparatory processes and preliminary processes are crucial in the proposed schemata. The justification for them could be given only in light of the different roles played by the different aspects. These examples also provide a strong motivation for limiting the possible temporal choices in narratives on the basis of semantics.

4 Important Connectives and an Example Narrative

We now show how the above treatment of events when combined with pragmatic information can successfully yield a comprehensive analysis of the temporal structure of narratives. We motivate and introduce the formal definitions for some important temporal connectives in English. We apply these to a narrative taken from the New York Times, which was proposed by John McCarthy as a challenge problem for NL. This helps ground or theory in a realistic example. Throughout, the key terms in the narrative have been highlighted.

H1. A 61-year old furniture salesman was *pushed* down the shaft of a freight elevator *yesterday* in his downtown Brooklyn store by two robbers *while* a third attempted to crush him with the elevator car *because* they were dissatisfied with the $1,200 they *had forced* him to give them.

Sentence (H1) describes a series of actions. The first action described by this sentence is that of "pushing" Mr. Hug. The verb *push* introduces a lexical entry as in (3), which yields a sentence predicate with the properties of [+CUM,+GRAD], i.e., an activity. The thematic role PATH, *down the shaft of a freight elevator*, does not affect the situation type of the sentence. Similarly, the LOCATION theme *in his downtown Brooklyn store* is an adjunct (like all spatial adverbials) and does not affect the situation type.

3. **push:** (S[−QUA,+GRAD]/NP[AGENT])/NP[PATH]
   (S[+QUA,+GRAD]/NP[AGENT])/NP[GOAL]

Yesterday merely establishes the reference time of the event introduced by the clause. The verb *attempt* introduces an activity, in essence the same as the previous case. The connective *while* temporally relates the action of pushing Mr. Hug with that of the attempt to crush him. Temporal connectives like *when* and *while* have been studied by several researchers. However, there is no exhaustive analysis of how these connectives relate various situation types. The semantics of *while* is given in (4). Here the category "Ph" is used only to syntactically prevent nesting of the connective. We assume that sentences of the form "X while Y" are parsed as [X [while Y]]. Thus, in postulate 4, P is bound to Y and Q to X. The proposed semantics states that the resulting phrase is an event predicate, that holds for e composed from events e' and e'', which satisfy P and Q, respectively. Further, e' and e'' must be related according to the applicable table entry.

4. **while:** (Ph/S)/S
   $\lambda P \lambda Q \lambda e[\exists e', e''[P(e') \wedge Q(e'') \wedge (e' \sqcup e'' = e) \wedge \text{table-entry}(e', e'')]]$

We now return to Mr. Hug's travails. Based on the algorithm, we would expect the relationship S1 $\subseteq$ S2 hold between the two schemas introduced by the *while*. This means that the event of pushing Mr. Hug down the shaft temporally overlaps with the event of attempting to crush him. Pragmatic inferences could help in specifying more concrete temporal locations, within the restraints provided by the semantics.

Here table-entry refers to the relationship between events depending on the situation type, aspect and the temporal connective. The table for the connective *while*

is given in table 1. These tables essentially declaratively encode an algorithm for computing the temporal relations that go into a discourse representation; standard techniques apply for the remaining, largely orthogonal, matters. For brevity, we don't mention thematic relations like AGENT and PATIENT.

In the tables below, PRG = Progressive; PRT = Perfect; and PRF = Perfective. Si refers to the entire schema for the corresponding sentence. Ii means the initial endpoint of schema Si. For an accomplishment or activity, Ei refers to the main event for its schema; for an achievement, Ei refers to its preliminary processes. Fi refers to the final endpoint of schema Si, which for accomplishments and activities, is F_N and for activities is F_A. $x \subseteq y$ means x occurs during y. $x < y$ means x precedes y. $x = y$ means x coincides with y. These relations are closely related to those of [1]. An $\times$ indicates that there is no clear semantic relationship for that class of events. In the case of *while* it generally means that the connective conveys a contrastive rather than a temporal relationship between the two clauses.

$S1 \Downarrow$ WHILE $S2 \Rightarrow$		Acc			Ach			Act		
		PRG	**PRT**	**PRF**	**PRG**	**PRT**	**PRF**	**PRG**	**PRT**	**PRF**
Acc	**PRG**	E1⊆E2	×	E1⊆E2	E1⊆E2	×	E1⊆E2	E1⊆E2	×	E1⊆E2
	PRT	F1<E2	×	S1<S2	S1<E2	×	S1⊆S2	S1⊆E2	×	S1<S2
	PRF	S1<E2	×	S1⊆S2	S1 ⊆E2	×	S1⊆S2	S1⊆E2	×	S1⊆S2
Ach	**PRG**	E1⊆E2	×	E1⊆S2	E1⊆E2	×	E1⊆E2	E1⊆E2	×	E1⊆S2
	PRT	F1<E2	×	S1<S2	S1<E2	×	S1<S2	S1<E2	×	S1<S2
	PRF	F1⊆E2	×	F1⊆S2	F1⊆E2	×	S1⊆S2	F1⊆E2	×	S1⊆S2
Act	**PRG**	E1⊆E2	×	E1⊆S2	E1⊆E2	×	E1⊆S2	E1⊆E2	×	E1⊆S2
	PRT	S1<E2	×	S1<S2	S1<E2	×	S1<S2	S1<E2	×	S1<S2
	PRF	S1⊆S2	×	S1⊆S2	S1⊆E2	×	S1⊆S2	S1⊆ E2	×	S1⊆S2

Table 1: Constraints Produced by the WHILE Connective

The last part of sentence (H1) is in the present perfect. The perfect construction has best been explained by Reichenbach using a tripartite temporal relation between speech time (ST), event time (ET) and reference time (RT) [15]. Kamp has extended Reichenbach's analysis by using the *temporal perspective point* (TP). Kamp's theory can explain why time moves backwards in perfect constructions. This means that the "forcing" occured before the events of pushing and crushing.

The connective *because* presents the second clause as the cause of the first. From world knowledge, we know that a cause cannot temporally succeed the effect. The verb *be (dissatisfied)* refers to a state. We can infer from the presence of the *had* that this state must have come to hold after the event of *forcing him to give them $1200.*

5. **because:** (Ph/S)/S
   $\lambda P \lambda Q \lambda e[\exists e', e''[P(e') \wedge Q(e'') \wedge (e' \sqcup e'' = e) \wedge FS(e') = FS(e'')]]$

H2. The buffer springs at the bottom of the shaft prevented the car from crushing the salesman, John J. Hug, *after* he was pushed from the first floor to the basement.

The temporal connective *after* establishes a strict temporal order between two situations. In this case the buffer springs prevented the crushing of Mr. Hug after the action of pushing him down the elevator shaft had ended. From our world knowledge we can also infer that the state holds after Mr. Hug had landed on the basement. The connective *after* can be formalized as in (6). The reference to the pushing event in this sentence is presented as an accomplishment whereas in sentence (H1) the same event was presented as an activity. This is a good example of the choice of the speaker or writer in reporting an event. Here the goal is specified, whereas in (H1) the path was specified. The goal *basement* introduces a definite natural endpoint, i.e., the pushing event is over after Mr. Hug reaches the basement.

6. **after:** (Ph/S)/S
   $\lambda P \lambda Q \lambda e[\exists e', e''[P(e') \wedge Q(e'') \wedge (e' \sqcup e'' = e) \wedge FS(e'') < IS(e')]]$

H3. The car stopped about 12 inches above him *as* he flattened himself at the bottom of the pit.

In sentence (H3) the verb *stopped* is inherently an achievement verb—it results in an achievement irrespective of its complement. In this sentence, the natural endpoint is 12 inches above Mr. Hug. The connective *as* functions quite like the connective *while* and relates two events as overlapping in time, with the additional restriction that they end together. In particular, *as* can apply even when the second event is an achievement. The length of the temporal overlap may vary a great deal from case to case, but we will not go into that here. *As* is formalized in (7). In this instance, it relates the achievement introduced by the verb *stopped* to the accomplishment introduced by the verb *flattened*. The lexical entry of the verb *flatten* specifies that the situation type of the resulting event predicate depends on the complements: if the object is cumulative, it is an activity; if the object is quantized, it is an accomplishment.

AFTER $S1 \Downarrow S2 \Rightarrow$		Acc			Ach			Act		
		PRG	**PRT**	**PRF**	**PRG**	**PRT**	**PRF**	**PRG**	**PRT**	**PRF**
Acc	**PRG**	×	×	×	×	×	×	×	×	×
	PRT	S1<E2	F1<F2	S1<S2	F1<E2	F1<F2	S1<S2	S1<E2	F1<F2	S1<S2
	PRF	S1<E2	F1<F2	S1<S2	S1<E2	F1<F2	S1<S2	S1<E2	F1<F2	S1<S2
Ach	**PRG**	×	×	×	×	×	×	×	×	×
	PRT	S1<E2	F1<F2	S1<S2	S1<E2	F1<F2	S1<S2	S1<E2	F1<F2	S1<S2
	PRF	S1<E2	F1<F2	S1<S2	S1<E2	F1<F2	S1<S2	S1<E2	F1<F2	S1<S2
Act	**PRG**	×	×	×	×	×	×	×	×	×
	PRT	S1<E2	F1<F2	S1<S2	F1<E2	F1<F2	S1<S2	S1<E2	F1<F2	S1<S2
	PRF	S1<E2	F1<F2	S1<S2	S1<E2	F1<F2	S1<S2	S1<E2	F1<F2	S1<S2

Table 2: Constraints Produced by the AFTER Connective

7. **as:** (Ph/S)/S
   $\lambda P \lambda Q \lambda e[\exists e', e''[P(e') \wedge Q(e'') \wedge (e' \sqcup e'' = e) \wedge FS(e') = FS(e'')]]$

H4. Mr. Hug was pinned in the shaft for about half an hour *until* his cries attracted the attention of a porter.

The existential verb *was* yields a state. The adjuncts delimit the interval over which the state holds. The connective *until*, as formalized in (9), requires that the second event end properly inside the first event. No claim is made about the initial endpoint of the second event. In this sentence, it is applied to a state (as generated by the verb *was pinned*) and an achievement (as generated by the verb *attract*). It asserts that Mr. Hug remained pinned until his cries attracted the attention of the porter. Pragmatically, we know that Mr. Hug must have remained for a little while after the porter's attention was attracted, so the porter could get him out. This is allowed, but not required, by the semantics.

8. **for:**$(\mathrm{S}[\alpha\textsc{qua},+\textsc{grad},+\textsc{neut}]/\mathrm{S}[\alpha\textsc{qua},+\textsc{grad}])/\mathrm{NP}[+\textsc{temp}]$
   $\lambda P' \lambda P \lambda e \exists t[\mathsf{DURATION}(e, t) \wedge P'(t)]$

9. **until:** (Ph/S)/S
   $\lambda P \lambda Q \lambda e[\exists e', e''[P(e') \wedge Q(e'') \wedge (e' \sqcup e'' = e) \wedge \text{initial-point}(e') < \text{final-point}(e'') \leq \text{final-point}(e')]]$

H5. The store at 340 Livingston Street is part of the Seaman's Quality Furniture chain.

The sentence (H5) is a state and since its endpoints are not demarcated it will be interpreted as holding throughout the events described.

H6. Mr. Hug was removed by members of the Police Emergency Squad and taken to Long Island College Hospital.

The verb *remove* introduces an accomplishment since its object is quantized. The connective *and* is almost like the default and signals the sequencing of events. In this example too it merely moves the narrative forward. *The Long Island College Hospital* is the GOAL of the verb *remove*.

H7. He was badly shaken, but *after* being treated for scrapes of his left arm and for a spinal injury was released and went home.

The existential verb *was* introduces a state and the verb *treated* introduces a accomplishment. The verb *treat* belongs to the class of verbs that always result in an accomplishment irrespective of the OBJECT in the discourse. In this case the *object* is quantized and the event is an accomplishment. The connective *after* relates the above accomplishment and the achievement of being released, which is itself followed by the accomplishment of going home.

H8. He lives at 62-01 69th Lane, Maspeth, Queens.

The state of living at the given address extends through the whole discourse.

10. **live:** S[STATE]/NP[AGENT]

H9. He has worked for seven years at the store, on the corner of Nevins Street, and this was the fourth time he had been held up in the store.

H10. The last time was about one year ago, when his right arm was slashed by a knife-wielding robber.

Sentences (H9) and (H10) involve anaphora since they refer to the previous occurrences of the event. Their discussion is beyond the scope of this paper.

Each of the temporal connectives occurring in the narrative contributes some constraints on the relations among the various events. These are collected as in the figure below.

5 Conclusions

We have presented a general semantics-based approach to determining the temporal structure of narratives. This approach explicitly takes into account both the situation type and aspect of the component sentences of a narrative. It presents semantic postulates, along with lexical entries, that may be used to systematically compute the temporal structure of a narrative. We have striven to preserve linguistically important distinctions in this analysis to make it applicable to a number of phenomena concerning events. To this end, we have motivated a general and natural schema representation for events, each of whose components is intuitively motivated and plays some role in the semantics. This approach has the advantages of being theoretically well-founded, declarative, and of applicability to a number of natural languages.

Our approach is semantics-based in its philosophy. We do not see it as competing with the pragmatics-based approaches, but rather as complementing and supporting them. Our approach can make a good initial cut at the temporal structure of a narrative, highlighting the gaps that pragmatics alone can fill. We believe that focusing the pragmatics in this manner leads to greater computational efficiency as well enhanced accuracy in inferring the necessary temporal relations.

References

[1] James Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, pages 832–843, 1983.

[2] Nicholas Asher and Michael Morreau. Commonsense entailment. In *IJCAI*, 1991.

[3] Kurt Eberle and Walter Kasper. Tenses as anaphora. In *Proceedings of the Conference of the European Association of Computational Linguistics*, 1989.

[4] Ennio Grasso, Leonardo Lesmo, Vincenzo Lombardo, Pia Maria Maccario, Roberto Salato, and Paolo Terenziani. Semantic interpretation of tense, actionality and aspect. In *Proceedings of the European Conference on Artificial Intelligence*, pages 320–325, Stockholm, 1990.

[5] Irene Heim. *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts, Amherst, MA, 1982.

[6] J. Hobbs. On the coherence and structure of discourse. TR CSLI-85-37, Center for the Study of Language and Information, Stanford, 1985.

[7] Hans Kamp. The perfect and other tenses in French and English. In *Tense and Aspect in English and French, DYANA Deliverable 2.3B*. Center for Cognitive Science, University of Edinburgh, Edinburgh, 1991.

[8] Manfred Krifka. Nominal reference, temporal constitution and quantification in event semantics. In R. Bartsch, J. van Benthem, and P. van Emde Boas, editors, *Semantics and Contextual Expressions*. Foris, Dordrecht, Holland, 1989.

[9] Manfred Krifka. Thematic relations as links between nominal reference and temporal constitution. In Ivan Sag and Anna Szabolcsi, editors, *Lexical Matters*. Chicago University Press, Chicago, 1991.

[10] Alex Lascarides and Nicholas Asher. A semantics and pragmatics of the pluperfect. In *Proceedings of the Conference of the European Association of Computational Linguistics*, 1993.

[11] Alex Lascarides and Nicholas Asher. Temporal interpretation, discourse relations, and commonsense entailment. *Linguistics and Philosophy*, 16:437–493, 1993.

[12] Marc Moens and Mark Steedman. Temporal ontology and temporal reference. *Computational Linguistics*, 14(2):15–28, 1988.

[13] A. Nakhimovsky. Aspect, aspectual class, and the temporal structure of narrative. *Computational Linguistics*, 14(2):29–43, 1988.

[14] Rebecca J. Passonneau. A computational model of the semantics of tense and aspect. *Computational Linguistics*, 14(2):44–60, 1988.

[15] Hans Reichenbach. *Elements of Symbolic Logic*. Macmillan, London, 1947.

[16] Mona Singh. The perfective aspect—an algebraic analysis. Paper read at the Conference on the Mathematics of Language organized by a special interest group of the ACL, 1991.

[17] Mona Singh and Munindar P. Singh. Computing the temporal structure of events in natural language. In *Proceedings of the European Conference on Artificial Intelligence*, pages 528–532, 1992.

[18] Carlota S. Smith. *The Parameter of Aspect*. Reidel - Kluwer, 1992.

[19] Paolo Terenziani. Integrating linguistic and pragmatic temporal information in natural language understanding: The case of "when sentences". In *IJCAI*, pages 1304–1309, 1993.

[20] Paolo Terenziani and Pietro Torasso. The effects of action-type on causation. In *WOCFAI*, 1991.

[21] Zeno Vendler. *Linguistics in Philosophy*. Cornell University Press, Ithaca, 1967.