

On the Temporal Structure of Events

Mona Singh*
Dragon Systems Inc.
320 Nevada Street
Newton, MA 02160, USA

mona.singh@ncsu.edu

Munindar P. Singh†
Department of Computer Science
North Carolina State University
Raleigh, NC 27695-8206, USA

singh@ncsu.edu

Abstract

The temporal interpretation of events is crucial to many applications of natural language understanding. Interpreting events as described in individual utterances and discourse requires both *semantic* and *pragmatic* processing. We develop representations for events that reflect the internal structure of different situation types. We propose semantic constraints through which the representations of different events can be related based on their tense, aspect, and situation type.

Our approach uses semantic information to a greater extent than previous approaches, and thereby generates constraints that restrict the ultimate pragmatic interpretations. We develop our approach on the notorious *when-clauses*, and show how it may be used for understanding events in discourse.

Introduction

Temporal reference is the problem of (a) inferring the internal structure of events as described in natural language, and (b) determining how events relate to one another. Determining the temporal reference of events is essential to understanding natural language (NL). Major applications of events in NL include understanding stories, processing texts, and following instructions for tasks. Accordingly, temporal reference has drawn continued attention from the research community (Kamp & Reyle 1993; Lascarides & Oberlander 1993; Singh & Singh 1992; Moens & Steedman 1988; Nakhimovsky 1988; Allen 1983).

The effective interpretation of events presupposes generic representations for them, and techniques to build and reason about these representations. It is widely recognized that these techniques must involve a

*Mona Singh was supported by the National Institute of Standards and Technology under grant NIST-70NANB5H1181. She is also affiliated with the Department of Computer Science, North Carolina State University.

†Munindar Singh was partially supported by the NCSU College of Engineering, by the National Science Foundation under grants IRI-9529179 and IRI-9624425, and IBM corporation.

pragmatic component to handle the complexity of NL, especially as engendered by the importance of world knowledge in understanding events. Pragmatic reasoning is required for understanding sentences such as (1) below. Here the first clause introduces an event of building a bridge; the continuations (a), (b), and (c) refer to before, during, or after that event. The necessary distinctions among the readings must be pragmatic in nature, and use the knowledge that plans are drawn up before construction, materials are used during construction, and traffic problems may be solved after a bridge is completed and put into operation.

1. When they built the 39th street bridge, ...
 - (a) a local architect drew up the plans.
 - (b) they used the best materials.
 - (b') they had used the best materials.
 - (c) they solved most of the traffic problems.

Although pragmatic reasoning is important, it can be computationally complex. We believe that many traditional approaches, e.g., (Lascarides & Asher 1993; Song & Cohen 1991), give up on grammar and semantics too soon. Pragmatics by itself cannot explain why the *during* reading is invalidated if clause (1b) is replaced by (1b'). Similarly, the *after* reading is invalid in (2a), also on semantic grounds. Pragmatics would have a tough time overcoming these semantic restrictions. Other approaches that consider semantics (Terenziani 1993) are still unable to handle the temporal distinctions between (1b) and (1b') and (2a) and (2a'). In other words, the semantics accounts are not complete enough to handle subtle distinctions, and the pragmatic accounts don't lend themselves to semantic phenomena easily.

2. When they were building the 39th street bridge, ...
 - (a) they solved most of the traffic problems.
 - (a') they were solving most of the traffic problems.

Basic Concepts. Before we can study the above problems in any depth, we must review the basic theoretical concepts. The interpretation of events involves a combination of three main components: situation type, tense, and aspect. Situation type refers to the

inherent internal structure of situations. Situations may characterize a state or an event (Vendler 1967). Events may be (a) *activities*, e.g., “walk in the park,” any part of which is also walking in the park and the activity may end at any arbitrary point, (b) *achievements*, e.g., “win a race,” which is instantaneous, or (c) *accomplishments*, e.g., “build a house,” which is the action of building, terminating in the completion of the house. Achievements and accomplishments are *telic*, because they have natural final endpoints, activities are *atelic*, because they do not have natural final endpoints. A natural endpoint is one which is somehow entailed by the definition of the verbal predicate. For example, for the verbal predicate *walk in the park*, any subevent (of some reasonable granularity) is also *walk in the park*, whereas proper subevents of *win a race* are not *win a race*.

Aspect is best defined as the viewpoint of the speaker towards a situation (Smith 1992). The perfective aspect, as in (3), describes a situation as a whole without any reference to its internal dynamics. The imperfective or progressive aspect, as in (4), is an internal view of the situation. (We will not discuss the habitual here.) Thus the situation type gives the objective structure of an event, whereas aspect gives the speaker’s view of it.

3. John ate an apple.
4. John was eating an apple.
5. John has eaten an apple.
6. John will have eaten an apple.

Tense is the most basic mechanism used by languages to anchor events in time. Much work has been done on tenses, so we shall not go into detail here. However, we do consider the traditional grammatical categories of perfect, e.g., (5), and futurate, e.g., (6), which combine properties of tense and aspect.

Temporal Reference. Several authors have worked on issues related to temporal reference (ter Meulen 1995; Terenziani 1993; Song & Cohen 1991; Grasso *et al.* 1990; Moens & Steedman 1988; Nakhimovsky 1988). Moens & Steedman can explain how (1a), (1b), and (1c) may be allowed, but they cannot explain the temporal distinctions between (1b) and (1b’) or (2a) and (2a’). Kamp & Reyle show how the tenses interact with events and give some general rules for how successive events (non-statives) may relate to each other. However, they use a very simple internal structure of different situation types. They note that the complexities of when-clauses cannot be captured by their apparatus. Other existing theories cannot handle the temporal distinctions between (7) and (8). Briefly, as will become clear, the above limitations arise because existing approaches do not distinguish (a) situation type from aspect, (b) the perfect from the perfective, and (c) the preparatory from the preliminary processes of an event introduced below.

Approach and Contributions. Briefly, our idea is to maintain as much of the semantic information about events as possible. This enables us to maximally constrain temporal reference across different sentences. We borrow some intuitions from (Kamp & Reyle 1993) and (Smith 1992), but develop a richer representation and more sophisticated rules for reasoning than these approaches. We develop a set of semantic constraints for each case of situation type, tense, and aspect. Although our approach applies to successive sentences in an utterance, in this paper we consider the case of when-clauses, which have a similar non-generic interpretation. We also consider the temporal connective *after*.

Our contributions are thus in the design of (a) structural schemas (b) inference rules for event location and closure, which apply in a wider variety of situations than previously considered possible. We motivate structural schemas for different situation types. We show how the aspects help locate events into the existing schemas, and describe how the tenses effectively “open up” or “close” the schemas for later utterances.

Organization This paper is organized as follows. Section motivates and proposes the structural schemas for the three main event types. A key idea here is the distinction between preparatory and preliminary processes. Section shows how the schemas can be used to construct discourse representations. A highlight is the application of our approach to the traditional grammatical categories of perfect and futurate. Section sums up the key rules of temporal reference in the context of the connectives *when* and *after*.

Structural Schemas

We now introduce a set of structures related to (Smith 1992). In these structures, I is the initial point, i.e., the point at which the given event begins. F_N is the natural final endpoint of the event. F_A is an arbitrary endpoint, i.e., a point at where the event ends arbitrarily. F_N is defined only for telic events. E is the main event for accomplishments and activities. To help understand our schemas, it is instructive to compare them to traditional representations. Figure 1 shows the structure of culminations or telics proposed by (Moens & Steedman 1988). This structure is also used by (Kamp & Reyle 1993) and others.



Figure 1: Culminated events (Moens & Steedman 1988)




Figure 2: Schema for accomplishments

Schema for Accomplishments

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. These constitute the main event *E*. 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. Figure 2 illustrates this schema.

Schema for Achievements

Achievements differ from accomplishments in that they are instantaneous. Thus they lack a main event *per se*. This means that *I* can be identified with *FN*. However, in many cases, achievements involve some events that are somehow salient to the natural ending. For example, “recognize 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, no proper subevent of winning a race constitutes the winning of a race. Figure 3 illustrates this schema.




Figure 3: Schema for achievements

Preparatory versus Preliminary Processes The above schema distinguishes between the *preparatory* and the *preliminary* processes of an event. This distinction is crucial. Consider the case where the first sentence, i.e., the one which initiates the DRS, is an achievement. In (7) below, the initial clause is in the perfective aspect; in (8), the initial clause is in the progressive aspect.

7. When John won the race, he took steroids.
8. When John was winning the race, he took steroids.

Example (7) suggests that John took the 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 (Lascarides & Asher 1991)—effectively claiming that it was taking steroids that made it possible for John to emerge victorious. In (8), 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. The pragmatic rule cannot push the taking of the steroid event before the running event because of restrictions on temporal reference imposed by the DRS. The distinction between preparatory processes and preliminary processes is crucial in determining the effects of different aspects. We return to this point in section .

Schema for Activities

Activities do not have any natural endpoints and can end at any time. But, they can include smaller events and preparatory processes. Consider sentence (9).

9. When John collected sea shells on the beach, ...
 - (a) he collected many kinds of sea shells.
 - (b) he had to wear his glasses.
 - (c) he was relaxed.




Figure 4: Schema for activities

Clause (9b) refers to the time before John actually collected any sea shells, (9a) refers to the subevents of the activity, and (9c) refers to the resultant state that obtained after the end of the activity. This shows that reference can be made to intervals before or after the actual activity. Hence, activities have preparatory processes and resultant states. Figure 4 illustrates this schema.

Constructing Discourse Representations

The above schemas are maximal representations of events. However, often only parts of the schemas are introduced into the discourse, as DRs. These parts depend on the aspect of the given clauses. Typically, the DRS of a sentence is merged with the DRS of the discourse so far. However, a new DRS is introduced as a separate entity for when-clauses, temporal adverbial clauses, and if the schema of a sentence does not fit in with the existing DRS. In either case, the components of a new DRS become available for merging with future discourse. We now explain how a DRS is constructed for the different aspects.

The Perfective

Clauses in the perfective introduce the entire schema into the discourse. If the sentences that follow are also

in the perfective, they add their schema to the existing DRS. For example, all three extensions of (1) are acceptable. Hence, the second clause must insert a DRS with the entire range of possibilities *drawing the plans* occurs during preliminary processes; *using the best materials* during the preparatory processes; and *solving the traffic problems* during the resultant state.

R1. The perfective aspect introduces the entire schema into the discourse.

The Progressive

Kamp & Reyle summarize the effect of the progressive as the schema minus the culmination point (Kamp & Reyle 1993, p. 566). This definition explains why sentence (10a) may be interpreted as occurring during the building of the bridge and the only possible interpretation of (10b) is that the traffic problems were solved *while* the bridge was being built. However, their theory cannot explain the oddity of (10c). It can only have a reading similar to that of (10a), that the plans were drawn up *while* the bridge was being built.

10. When they were building the bridge, ...

- (a) they used the best materials.
- (b) they solved the traffic problems.
- (c) a local architect drew up the plans.

In (10), only the actual bridge-building (from I to F_N or F_A) is available on the DRS. Thus the progressive leaves out not only the resultant state as noted by Kamp & Reyle, but also the preparatory processes. Thus (10c), which refers to a preparatory process, does not have the interpretation of preceding the actual building event.

The progressive aspect introduces only the core, i.e., from I to F_N . Thus, the progressive leaves out the preparatory processes and the resultant state.

R2. For achievements without preliminary processes, the progressive may

- yield an iterative interpretation, if the object can be subjected to the same action repeatedly;
- present a stretched out interpretation of the action, otherwise. An example of the latter is visualizing *John was breaking the glass* in slow motion.

The predicate *win a race* has a distinct preliminary process (i.e., running), whereas *knock the door* does not. Thus (11) suggests that John took the steroids while running. This is because preparatory processes are left out by the progressive aspect. However, (12) forces an interpretation where John heard the dog bark while repeatedly knocking. The knocking events act as a cluster, which behaves like an event beginning with the initial point of the first knocking and ends with the final endpoint of the last knocking. The progressive aspect introduces the event cluster into the discourse. The subsequent event of the dog barking is thus interpreted as occurring during the knocking events. None

of the individual knocking events may coincide temporally with any of the barking subevents.

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

12. When John was knocking at the door, he heard the dog bark.

The Role of the Tenses

The simple tenses (past, present, and future) trivially locate the event before, during, and after the speech time, respectively. However, tenses interact with aspect to produce some interesting (traditional) grammatical categories such as the perfect and futurate. Both the perfect and the futurate combine features of tense and aspect: the perfect combines the perfective aspect with the present or past tense; the futurate combines the perfective aspect and the future tense.

Reichenbach first noted that the perfect alludes to three points, namely, the *reference time* (RT), or the standpoint from which the event is described; *speech time* (ST), or the time of utterance; and the *event time* (ET), or the time of occurrence (Reichenbach 1947). With various combinations of these times, we can arrive at the general schema for the present perfect, past perfect, and the futurate. In the present perfect, e.g., (13), the ST and the RT are the same and the time of the event precedes the time of speech. Hence, $ET < ST = RT$ holds.

13. John has won the race.

However, in addition to their temporal interpretation, the perfect and the futurate also have an effect of attributing a property. Consider (14) and (15). The past perfect suggests that there is a reference time, i.e., yesterday, when Henry had the property of having danced or having eaten an apple, respectively. In particular, it requires that Henry existed as a separate entity, i.e., was alive, then. Hence, $ET < RT < ST$ holds for the past perfect.

14. Henry had danced yesterday.

15. Henry had eaten an apple yesterday.

The futurate is exemplified in (16) and (17) below, which mean that there will be a time such that Henry will have the attributes of having danced or having eaten an apple, respectively. Hence, $ST < ET < RT$ holds for the futurate.

16. Henry will have danced tomorrow.

17. Henry will have eaten an apple tomorrow.

In the perfect, the attribution essentially holds till the reference time, but may not hold afterward. For example, (18) means that at the reference time, which equals the speech time, the property of having won the world championship applies to John. But, if John died after he won the world championship, then the simple perfect will no longer hold, though the perfective will continue to hold. The schema for the present perfect is given in Figure 5.

18. John has won the world championship.
 19. John won the world championship.




Figure 5: Schema for the present perfect

It is known that the temporal reference for the present perfect is $ET < ST = RT$; the past perfect is $ET < RT < ST$; and the futurate is $ST < ET < RT$ (Reichenbach 1947). Therefore, we propose the the following rule.

- R3. The perfect tends to seal off the schema of the event it introduces, making it difficult to refer to events or times inside it.

Thus, whereas all the combinations in (1) are grammatical, the corresponding cases with the perfect replacing the perfective in the initial clause are not. It seems that, whereas (20a) and (20b) do not have the intended reading, the perfect aspect in the clauses that follow the when-clause, as in (21) are have the same temporal interpretations as those in (1). Interestingly, in (21), the event time of the when-clause becomes the reference time for the succeeding sentences.

20. When they had built the 39th street bridge, ...
 (a) a local architect drew up the plans.
 (b) they used the best materials.
 (c) they solved most of the traffic problems.
21. When they had built the 39th street bridge, ...
 (a) a local architect had drawn up the plans.
 (b) they had used the best materials.
 (c) they had solved most of the traffic problems.

We should emphasize that since the perfect and futurate are a combination of the perfective and a tense (past, present, or future), there is a *tendency* for them to occur with culminated events. Typically a perfect construction depicts a complete event as having taken place with reference at a particular time. The event, therefore, takes on a point-like characteristic, since reference to its preparatory processes and resultant states cannot be made. The perfect essentially depicts the event as an unanalyzed whole, thereby presenting it as momentous. In general, the perfect is available for *all* event types:

22. Mary has swum in the pond. (activity)
 23. John has built a house. (accomplishment)
 24. John has won the race. (achievement)

By contrast, (Moens & Steedman 1988) claim that perfect can only apply with culminated events. This appears to hold for their specific examples, e.g., *John has hummed*, whose verbal predicates do not yield easily attributable properties.

Rules

We present some additional constraints introduced by some common temporal connectives. We don't claim to have the entire set of rules.

When The following rules help relate the event of the when-clause with the event of the subsequent clause. For lack of space we present the rules for accomplishments only.

- R4. If an **accomplishment** is introduced using the **past perfect** and the subsequent clause is in the progressive or perfective, then the *F* of the when-clause must precede the *F* of the event in the subsequent clause. Thus the *before*, and *during* readings of (20a) and (20b) cannot be achieved.
- R5. If an **accomplishment** is introduced using the **past perfect** and the subsequent clause is in the past perfect, then the event of the subsequent clauses may be interpreted as having occurred *before*, *during*, or *after* the event of the main clause.
- R6. If an **accomplishment** is introduced using the **perfective** by a when-clause, the events of the main clause and the subsequent clauses may overlap. The subsequent clause may occur during the preparatory, preliminary, or resultant stages (as in (1a), (1b), and (1c)). An exception is that if the subsequent clause is in the past perfect, the subsequent clause is interpreted as strictly preceding the when-clause event (as in (1b')).
- R7. If the **progressive** is used to introduce an event (accomplishment, achievement, or activity) and the following utterance is in the progressive or the perfective, the second event must be temporally located during the *E* of the first event.
- R8. If the **progressive** is used to introduce an event and the following utterance is in the **present perfect**, the result is likely to be semantically unacceptable, e.g., **When they were building the 39th street bridge, a local architect has drawn up the plans*. This is because the present perfect requires that $ST = RT$. However, the when-clause introduces a reference time before speech time. Thus there is a temporal conflict.
- R9. If the **progressive** is used to introduce an event and the following utterance is in the **past perfect**, the event of the subsequent clause is located before the event of the main clause.

After Here are some rules for *after*. Interestingly, these do not distinguish among the situation type of the events introduced.

- R10. After cannot introduce a clause in the **progressive** (e.g., **After John was eating an apple he took*

a walk). A possible explanation is that the progressive draws the focus on the preliminary processes, whereas the after draws attention toward the resultant state of the event.

- R11. If an event is introduced in the **perfect** or the **perfective**, and the subsequent clause is in the **progressive**, then the F_A or F_N of the first event must precede the I of the second event (e.g., *After they built the 39th street bridge, the city was going bankrupt*). The resultant states of the first event may overlap with the preparatory processes of the second.
- R12. If an event is introduced in the **perfect** or the **perfective**, and the subsequent clause is in the **perfect**, the F of the main clause must precede the F of the subsequent clause (e.g., *After John won the race, he had got a medal*). The event of the subsequent clause can begin before the event in the after clause ended. For example, in *After they built the 39th street bridge, they had solved the traffic problems*, the solving of the traffic problems started when they started building the bridge, although the traffic problems were solved only after the bridge was complete.
- R13. If an event is introduced in the **perfect** or the **perfective**, and the subsequent clause is in the **perfective**, then the entire event introduced by the after-clause must precede the event in the subsequent clause (e.g., *After John won the race, he went home*).

Conclusions

We proposed structured representations for events that accommodate both the objective and the subjective facets of natural language. We first defined schemas for different situation types that characterize the intrinsic or objective structure of events. Next, we showed how the viewpoint towards an event, as expressed in the speaker's choice of aspect, can be used to determine the contribution of a sentence to the representation of an ongoing discourse. The *DRS* of a discourse as it has progressed contains the "visible" parts of the structures of the various events. The *DRS* of a new sentence is inserted into it by merging or appending as necessary. Consequently, certain readings or interpretations are eliminated purely on grammatical or semantic grounds without invocation of any powerful notions of world knowledge and pragmatics. This is important, because an approach that invokes such inferences more than necessary would not only be computationally intractable, but also less effective in identifying the best interpretations. We emphasize that we do not seek to obviate the pragmatics component, just minimize the semantic interpretations it has to select among.

We also argued that certain intuitive distinctions between preparatory and preliminary processes can be made in a principled manner and used to motivate the

key difference between the progressive and the perfective. They can also be used to explain the distinction between the aspects and the common categories from traditional grammar.

References

- Allen, J. 1983. Maintaining knowledge about temporal intervals. *CACM* 832-843.
- Grasso, E.; Lesmo, L.; Lombardo, V.; Maccario, P. M.; Salato, R.; and Terenziani, P. 1990. Semantic interpretation of tense, actionality and aspect. In *Proceedings of the European Conference on Artificial Intelligence*, 320-325.
- Kamp, H., and Reyle, U. 1993. *From Discourse to Logic*. Dordrecht: Kluwer.
- Lascarides, A., and Asher, N. 1991. Discourse relations and common sense entailment. In *Proceedings of the Conference of the Association of Computational Linguistics*.
- Lascarides, A., and Asher, N. 1993. A semantics and pragmatics of the pluperfect. In *Proceedings of the Conference of the European Association of Computational Linguistics*.
- Lascarides, A., and Oberlander, J. 1993. Temporal connectives in a discourse context. In *Proceedings of the Conference of the European Association of Computational Linguistics*.
- Moens, M., and Steedman, M. 1988. Temporal ontology and temporal reference. *Computational Linguistics* 14(2):15-28.
- Nakhimovsky, A. 1988. Aspect, aspectual class, and the temporal structure of narrative. *Computational Linguistics* 14(2):29-43.
- Reichenbach, H. 1947. *Elements of Symbolic Logic*. London, UK: Macmillan.
- Singh, M., and Singh, M. P. 1992. Computing the temporal structure of events in natural language. In *Proceedings of the European Conference on Artificial Intelligence*.
- Smith, C. S. 1992. *The Parameter of Aspect*. Reidel - Kluwer.
- Song, F., and Cohen, R. 1991. Tense interpretation in the context of narrative. In *AAAI*, 131-136.
- ter Meulen, A. 1995. *Representing Time in Natural Language*. Cambridge, MA: MIT Press.
- Terenziani, P. 1993. Integrating linguistic and pragmatic temporal information in natural language understanding: The case of "when sentences". In *IJCAI*, 1304-1309.
- Vendler, Z. 1967. *Linguistics in Philosophy*. Ithaca, NY: Cornell University Press.