

TEMPORAL ONTOLOGY IN NATURAL LANGUAGE

Marc Moens* and Mark Steedman†

Centre for Cognitive Science*† and Dept. of AI†, Univ. of Edinburgh,
and Dept. of Computer and Information Science, Univ. of Pennsylvania†

ABSTRACT

A semantics of linguistic categories like tense, aspect, and certain temporal adverbials, and a theory of their use in defining the temporal relations of events, both require a more complex structure on the domain underlying the meaning representations than is commonly assumed. The paper proposes an ontology based on such notions as causation and consequence, rather than on purely temporal primitives. We claim that any manageable logic or other formal system for natural language temporal descriptions will have to embody such an ontology, as will any usable temporal database for knowledge about events which is to be interrogated using natural language.

1. Introduction

It has usually been assumed that the semantics of temporal expressions is directly related to the linear dimensional conception of time familiar from high-school physics -- that is, to a model based on the number-line. However, there are good reasons for suspecting that such a conception is not the one that our linguistic categories are most directly related to.

When-clauses provide an example of the mismatch between linguistic temporal categories and a semantics based on such an assumption. Consider the following examples:

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

To map the temporal relations expressed in these examples onto linear time, and to try to express the semantics of *when* in terms of points or intervals (possibly associated with events), would appear to imply either that *when* is multiply ambiguous, allowing these points or intervals to be temporally related in at least three different ways, or that the relation expressed between main and *when*-clauses is one of "approximate coincidence". However, neither of these tactics explains the peculiarity of utterances like the following:

- (2) #When my car broke down, the sun set.

The oddity of this statement seems to arise because the *when*-clause predicates something more than mere temporal coincidence, that is, some contingent relation such as a *causal* link between the two events. Of course, our knowledge of the world does not easily support such a link. This aspect of the sentence's meaning must stem from the sense-meaning of *when*, because parallel utterances using *just after*, *at approximately the same time as*, and the like, which predicate purely temporal coincidence, are perfectly felicitous.

We shall claim that the different temporal relations conveyed in examples (1) do not arise from any sense-ambiguity of *when*, or from any "fuzziness" in the relation that it expresses between the times referred to in the clauses it conjoins, but from the fact that the meaning of *when* is not primarily temporal at

all. We shall argue that *when* has single sense-meaning reflecting its role of establishing a temporal focus. The apparent diversity of meanings arises from the nature of this referent and the organisation of events and states of affairs in episodic memory under a relation we shall call "contingency", a term related to such notions as causality, rather than temporal sequentiality. This contingent, non-temporal relation also determines the ontology of the elementary propositions denoting events and states of which episodic memory is composed, and it is to these that we turn first.

2. Temporal and Aspectual Categories

Utterances of English sentences can, following Vendler, be classified into temporal/aspectual types on the basis of the tenses, aspects and adverbials with which they can cooccur (cf. Dowty, 1979). This "aspectual type" refers to the relation to other happenings in the discourse that a speaker predicates of the particular happening that their utterance describes. Thus an utterance of *Harry reached the top* is usually typical of what we will call a "culmination" -- informally, an event which the speaker views as accompanied by a transition to a new state of the world.¹ This new state we will refer to as the "consequent state" of the event. *Harry hiccupped* is not usually viewed by speakers as leading to any specific change of state. It typifies what we call "point" expressions, that is punctual events whose consequences are not at issue.

Similarly, *Harry climbed* typifies what we will call for obvious reasons a "process": such utterances describe an event as extended in time but not characterised by any particular conclusion or culmination. In contrast, *Harry climbed to the top* typically describes a state of affairs that also extends in time but that *does* have a particular culmination associated with it at which a change of state takes place. We classify such an utterance as a "culminated process". Finally, *Harry is at the top* typically describes a state.

Thus we can interpret Vendler as saying that a part of the meaning of any utterance of a sentence is one of a small number of temporal/aspectual profiles distinguished on a small number of dimensions. They can be summarized as in Figure 1.

It is important to be clear that this claim concerns sentences *used in a context*: sense-meanings of sentences or verbs in isolation are usually compatible with several (or even all possible) Vendlerian profiles, as Dowty and Verkuyl have pointed out -- hence the frequent use of the word "typically" above. The details of this taxonomy and the criteria according to which utterances can be categorised are therefore less important than

¹ Readers familiar with Vendler's work will realise that we have changed his terminology. We have done so both for notational convenience and to avoid the considerable confusion that has arisen concerning the precise meaning of the old terms.

EVENTS		STATES	
	atomic	extended	
+conseq	Harry left early At six, John arrived	Sue built a sandcastle The ice melted completely	John knows French He was in the kitchen
-conseq	Sandra hiccupped Paul winked	Max worked in the garden Alice played the piano	

Figure 1

the observation that each primitive entity of a given type, such as the culmination-event of *Harry's reaching the top*, carries intimations of other associated events and states, such as the process by which the culmination was achieved, and the consequent state that followed. What linguistic devices like tenses, aspects, and temporal/aspectual adverbials appear to do is to transform entities of one type into these other "contingently" related entities, or to turn them into composites with those related entities.

The temporal/aspectual ontology that underlies these phenomena can be defined in terms of the state-transition network shown in Figure 2. The semantics of tenses, aspectual auxiliaries and temporal adverbials is defined in terms of functions which map categories onto other categories, and having the important characteristic of "coercing" their argument to be of the appropriate type. Both the possibilities for coercing an input proposition, and the possibilities for the output category, are defined by the transition net. In addition, the felicity of a particular transition is conditional on support from knowledge and context.

Consider, for example, the combination of a culminated process expression with a *for*-adverbial, as in

(3) Sue played the sonata for a few minutes.

A *for*-adverbial coerces its input to be of the process variety. According to the network in Figure 2, such a transition is felicitous if the culmination point associated with the event of *playing the sonata* is "stripped off". As a result, there is no implication in (3) that Sue *finished* playing the sonata.

Another route through the network is possible in order to account for examples with *for*-adverbials: the culminated process, like any other event, can be viewed as an unstructured "point". A transition to turn it into a process then results in an iteration of occurrences at which Sue plays the sonata. This route through the network seems to be ruled out for (3) because it finds no support in our knowledge about sonatas and about how long they typically last. It *does* result, however, in a likely interpretation for a sentence like

(4) Sue played the sonata for about eight hours.

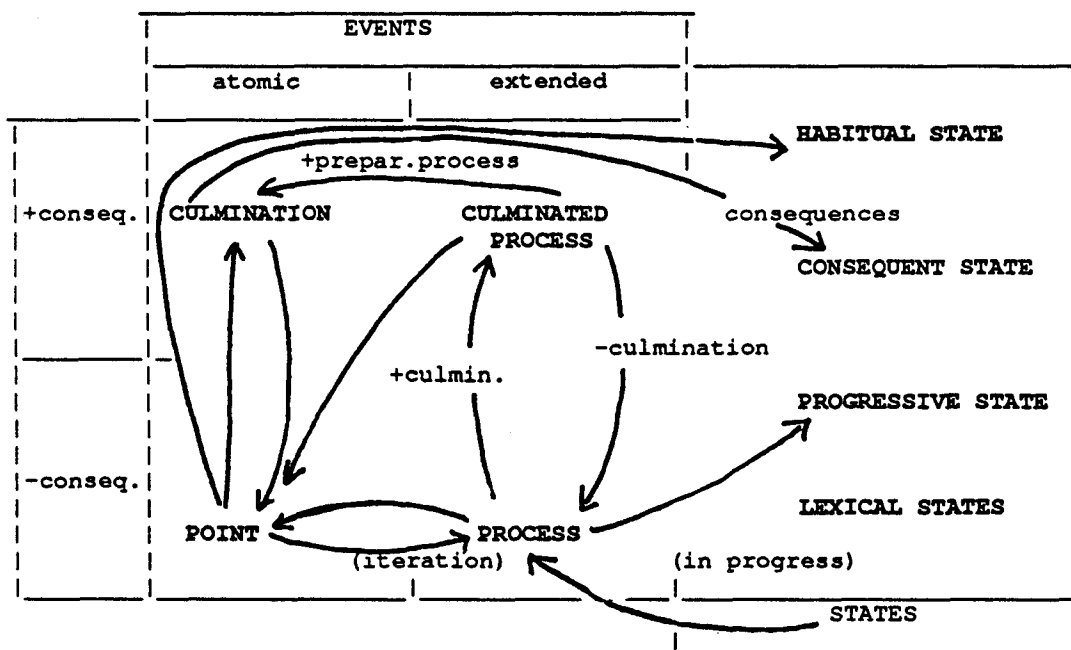


Figure 2

Not all aspectual/temporal adverbials expressing a time span have the same functional type. *In*-adverbials, for example, coerce their input to be a *culminated* process expression. This means that combination with a culmination expression requires the transition to be made to the culminated process node. According to the aspectual network in Figure 2 this transition is a felicitous one if the context allows a preparatory process to be associated with the culmination, as in (5):

(5) Laura reached the top in two hours.

The *in*-adverbial then defines the length of this preparatory period.

Since the arcs describe what the world has to be like for transitions to be made felicitously, it is obvious that there are expressions that will resist certain changes. For example, it will be hard to find a context in which an *in*-adverbial can be combined with a culmination expression like *Harry accidentally spilled his coffee*, since it is hard to imagine a context in which a preparatory process can be associated with an involuntary act.

A similar problem arises in connection with the following example:

(6) John ran in a few minutes

The process expression *John ran* has to be changed into a culminated process expression before combination with the *in*-adverbial is possible. One way in which the network in Figure 2 will permit the change from a process to a culminated process is if the context allows a culmination point to be associated with the process itself. General world knowledge makes this rather hard for a sentence like *John ran*, except in the case where John habitually runs a particular distance, such as a measured mile. If the *in*-adverbial had conveyed a specific duration, such as *in four minutes*, then the analysis would make sense, as Dowty has pointed out. However, the unspecified *in a few minutes* continues to resist this interpretation.

However, another route is also possible for (6): the process of *John running* can be made into an atomic point, and thence into a culmination in its own right. This culmination can then acquire a preparatory process of its own -- which we can think of as *preparing to run* -- to become the culminated process which the adverbial requires. This time, there is no conflict with the content of the adverbial, so this reading is the most accessible of the two.

Progressive auxiliaries coerce their input to be a process expression. The result of the application is a progressive state, which describes the process as being in progress. This means that, when a culmination expression like *reach the top* is used with a progressive, a transition path has to be found from the culmination node to the process node. According to the transition network, this involves first adding a preparatory process to the culmination, and then stripping off the culmination point. As a result, the progressive sentence only describes the preparation as ongoing and no longer asserts that the original culmination even occurred. There would be no contradiction in continuing

(7) Harry was reaching the top

as in

(8) Harry was reaching the top but he slipped and fell before he got there.

As Moens & Steedman (1986) point out, the fact that according to the present theory, progressives coerce their input to be a process, so that any associated culmination is stripped away and no longer contributes to truth conditions, provides a resolution of the "imperfective paradox" (Dowty 1979), without appealing to theory-external constructs like "inertia worlds".

A perfect, as in

(9) Harry has reached the top

refers to the consequent state of the culmination. It requires its input category to be either a culmination or a culminated process, and maps this expression into its consequent state. Informal evidence that it does so can be obtained by noticing that perfects are infelicitous if the salient consequences are *not* in force. The most obvious of these consequences for (9) is that Harry still *be* at the top, although as usual there are other possibilities.

Since the transition network includes loops, it will allow us to define indefinitely complex temporal/aspectual categories, like the one evoked by the following sentence:

(10) It took me two days to play the "Minute Waltz" in less than sixty seconds for more than an hour.

The culminated process expression *play the Minute Waltz* can combine straightforwardly with the *in*-adverbial, indicating how long it takes to reach the culmination point of *finishing playing the Minute Waltz*. Combination with the *for*-adverbial requires this expression to be turned into a process - the most obvious route through the network being that through the point node. The resulting culminated process expression describes the iterated process of *playing the Minute Waltz in less than sixty seconds* as lasting for more than an hour. The expression *it took me...*, finally, is like an *in*-adverbial in that it is looking for a culminated process expression to combine with. It finds one in the expression *to play the Minute Waltz in less than sixty seconds for more than an hour* but combination is hampered by the fact that there is a conflict in the length of time the adverbials describe. In the case of (10), the whole culminated process is instead viewed as a culmination in its own right (via the path through the point node). Knowledge concerning such musical feats then supplies an appropriate preparatory process which we can think of as *practising*. The adverbial *it took me two days* then defines the temporal extent of this preparatory process needed to reach the point at which repeatedly playing that piece of music so fast for such a considerable length of time became a newly acquired skill.

This basic framework thus allows for a unified semantics of a wide variety of aspectual adverbials, the progressive, the perfect, and iterative expressions in English. It is also used to explain the effect of bare plurals and certain varieties of negation on the overall temporal structure of discourse (Moens forthcoming).

All of the permissible transitions between aspectual categories illustrated in Figure 2 appear to be related to a single elementary contingency-based event structure which we call a "nucleus". A nucleus is defined as a structure comprising a culmination, an associated preparatory process, and a conse-

quent state. It can be represented pictorially as in Figure 3.²




Figure 3

Any or all of these elements may be compound: the preparation may consist of a number of discrete steps, for example the stages of climbing, having lunch or whatever, leading to the culmination of *reaching the top*. The consequent state may also be compound. Most importantly, we shall see that it includes the further events, if any, that are in the same sequence of contingently related events as the culmination. Similarly, the culmination may itself be a complex event — such as the entire culminated process of *climbing a mountain*. (In this case, the associated preparatory process and consequent state will be quite different ones to those internal to the culminated process itself.) The device is intended to embody the proposal that when we hear about an event like *climbing a mountain* in conjunction with some coercive aspectual category which forces it to undergo a transition, then the alternatives that are available are:

- a) to *decompose* the core event into a nucleus and to make a transition to one of the components, such as the preparatory activity of climbing or to the consequent state of having climbed the mountain; or
- b) to treat the entire event as a culmination, to *compose* it into a nucleus with whatever preparation and consequences the context provides for the activity of climbing a mountain, and to make the transition to either one of those.

We further claim that those are the *only* alternatives.

The concept of a nucleus not only explains the transitions of Figure 2, but also provides an answer to the question raised in the introduction concerning the apparent vagaries in the meaning of *when*-clauses.

3. *When*-clauses

The aspects and temporal/aspectual adverbials considered above all act to modify or change the aspectual class of the core proposition, subject to the limits imposed by the network in Figure 2, which we claim is in turn determined by the organisation of episodic memory. However, tenses and certain other varieties of adverbial adjuncts have a rather different character. Tense is widely regarded as an anaphoric category, requiring a previously established temporal referent. The referent for a present tense is usually the time of speech, but the referent for a past tense must be explicitly established. Such a referent is usually established using a second type of "temporal" adverbial, such as *once upon a time*, *at five o'clock last Saturday*, *while I was cleaning my teeth*, or *when I woke up this morning*.

Most accounts of the anaphoric nature of tense have invoked Reichenbach's (1947) trinity of underlying times, and his concept of the "positional" use of the reference time which he

called "R". Under these accounts (reviewed in Steedman, 1982), the adjuncts establish a reference time to which the reference time of a main clause and subsequent same-tensed clauses may attach or refer, in much the same way that various species of full noun phrases establish referents for pronouns.

However, in one respect, the past tense does not behave like a pronoun. Use of a pronoun such as "she" does not change the referent to which a subsequent use of the same pronoun may refer, whereas using a past tense may. In the following example, the temporal reference point for sentence (b) seems to have moved on from the time established by the adjunct in (a):

- (11) a. At exactly five o'clock, Harry walked in.
- b. He sat down.

This fact has caused theorists such as Dowty (1986), Hinrichs (1984) and Partee (1984) to stipulate that the reference time autonomously advances during a narrative. However, such a stipulation (besides creating problems for the theory *vis-à-vis* those narratives where reference time seems *not* to advance) seems to be unnecessary, since the *amount* by which it advances still has to be determined by context. The concept of a nucleus that was invoked above to explain the varieties of aspectual categories offers us exactly what we need to explain both the fact of the advance and its extent. We simply need to assume that a main clause event such as *Harry walked in* is interpreted as an entire nucleus, complete with consequent state, for by definition the consequent state includes whatever other events were contingent upon Harry walking in, including whatever he did next. Provided that the context (or the hearer's assumptions about the world) supports the idea that a subsequent main clause identifies this next contingent event, then it will provide the temporal referent for that main clause. In its ability to refer to entities that have not been explicitly mentioned, but whose existence has merely been *implied* by the presence of an entity that *has* been mentioned, tense appears more like a definite NP like *the music* in the following example than like a pronoun, as Webber (1987) points out.

- (12) I went to a party last night. The music was wonderful.

A similar move is all that is required to explain the puzzle concerning the apparent ambiguity of *when*-clauses with which the paper began. A *when*-clause behaves rather like one of those phrases that are used to explicitly change topic, like *and your father* in the following example from Isard, (1975):

- (13) And your father, how is he?

A *when*-clause introduces a novel temporal referent into focus whose unique identifiability in the hearer's memory is similarly presupposed. However, again the focussed temporal referent is an entire nucleus, and again an event main clause can attach itself anywhere within this structure that world knowledge will allow. For example, consider the example (1) with which we began (repeated here):

- (14) When they built the 39th Street bridge...
 - (a) ...a local architect drew up the plans.
 - (b) ...they used the best materials.
 - (c) ...they solved most of their traffic problems.

Once the core event of the *when*-clause has been identified in memory, the hearer has the same two alternatives described before: either it is decomposed into a preparatory process, a culmination and a consequent state, or the entire event is treated as itself the culmination of another nucleus. Either

² A similar event structure is proposed by Passonneau (1987).

way, once the nucleus is established, the reference time of the main clause has to be situated somewhere within it - the exact location being determined by knowledge of the entities involved and the episode in question. So in example (a) the entire culminated process of building the bridge becomes a culmination (via a path in Figure 2 which passes through the "point" node) which is associated in a nucleus with preparations for, and consequences of, the entire business, as in Figure 4:




Figure 4

The drawing up of the plans is then, for reasons to do with knowledge of the world, situated in the preparatory phase.

In example (b), in contrast, the building of the bridge is decomposed into a quite different preparatory process of building, a quite different culmination of completing the bridge and some consequences which we take to be also subtly distinct from those in the previous case, as in Figure 5. The use of the best materials is then, as in (a), situated in the preparatory process -- but it is a different one this time.




Figure 5

Example (c) is like (a) in giving rise to the nucleus in Figure 4, but pragmatics demands that the main clause be situated somewhere in the consequent state of building the bridge.

Thus, a main clause event can potentially be situated anywhere along this nucleus, subject to support from knowledge about the precise events involved. But example (2) is still strange, because it is so hard to think of *any* relation that is supported in this way:

(15) #When my car broke down, the sun set.

The *when*-clause defines a nucleus, consisting of whatever process we can think of as leading up to the car's break-down, the break-down itself and its possible or actual consequences. It is not clear where along this nucleus the culmination of *the sun set* could be situated: it is not easy to imagine that it is a functional part of the preparatory process typically associated with a break-down, and it is similarly hard to imagine that it can be a part of the consequent state, so under most imaginable circumstances, the utterance remains bizarre.

The constraints *when* places on possible interpretations of the relation between subordinate and main clause are therefore quite strong. First, general and specific knowledge about the event described in the *when*-clause has to support the association of a complete nucleus with it. Secondly, world knowledge also has to support the contingency relation between the events in subordinate and main clause. As a result, many constructed

examples sound strange or are considered to be infelicitous, because too much context has to be imported to make sense of them.

In all of the cases discussed so far, the main clause has been an event of some variety. When the main clause is stative, as in the following examples, the effect is much the same. That is to say, the *when*-clause establishes a nucleus, and the stative is asserted or attached wherever world knowledge permits within the nucleus. The only difference is that statives are by definition unbounded with respect to the reference time that they are predicated of, and outlast it. It follows that they can usually fit in almost anywhere, and therefore tend not to coerce the *when*-clause, or to induce the causal/contingent interpretations that we claim characterise the corresponding sentences with events as main clauses:

- (16) When they built that bridge
 ...I was still a young lad.
 ...my grandfather had been dead for several years.
 ...my aunt was having an affair with the milkman.
 ...my father used to play squash.

However, world knowledge may on occasion constrain the relation of a stative main clause, and force it to attach to or describe a situation holding over either the preparatory process or the consequent state of the subordinate clause, as in the following examples (cf. Smith 1983):

- (17) When Harry broke Sue's vase,
 ...she was in a good mood.
 ...she was in a bad mood.

4. Towards a Formal Representation

We have argued in this paper that a principled and unified semantics of natural language categories like tense, aspect and aspectual/temporal adverbials requires an ontology based on contingency rather than temporality. The notion of "nucleus" plays a crucial role in this ontology. The process of temporal reference involves reference to the appropriate part of a nucleus, where appropriateness is a function of the inherent meaning of the core expression, of the coercive nature of co-occurring linguistic expressions, and of particular and general knowledge about the area of discourse.

The identification of the correct ontology is also a vital preliminary to the construction and management of temporal databases. Effective exchange of information between people and machines is easier if the data structures that are used to organise the information in the machine correspond in a natural way to the conceptual structures people use to organize the same information. In fact, the penalties for a bad fit between data-structures and human concepts are usually crippling for any attempt to provide natural language interfaces for data base systems. Information extracted from natural language text can only be stored to the extent that it fits the preconceived formats, usually resulting in loss of information. Conversely, such data structures cannot easily be queried using natural language if there is a bad fit between the conceptual structure implicit in the query and the conceptual structure of the database.

The "contingency-based" ontology that we are advocating here has a number of implications for the construction and management of such temporal databases. Rather than a homogeneous

database of dated points or intervals, we should partition it into distinct sequences of causally or otherwise contingently related sequences of events which we will call "episodes", each leading to the satisfaction of a particular goal or intention. This partition will quite incidentally define a partial temporal ordering on the events, but the primary purpose of such sequences is more related to the notion of a plan of action or an explanation of an event's occurrence than to anything to do with time itself. It follows that only events that are contingently related necessarily have well defined temporal relations in memory.

A first attempt to investigate this kind of system was reported by Steedman (1982), using a program that verified queries against a database structured according to some of the principles outlined above. These principles can be described using Kowalski's event-calculus (Kowalski & Sergot 1986). In this framework, there are primitives called events, the occurrence of which usually implies the existence of periods of time over which states hold. In the terms of the present paper, these "events" are either "points" or "culminations" (depending on whether they are in fact associated with *consequent* states -- see section 2). For example, in the world of academic promotions which Kowalski and Sergot take as an example, an event description like *John was promoted from the rank of lecturer to the rank of professor* is a culmination which implies that there was a period of time, ended by this event, during which John had the rank of lecturer, and there is a period of time, started by that same event, during which John had the rank of professor.

The events in the event calculus are given unique identifiers, but are not necessarily associated with absolute time. Moreover, they can be partially ordered with respect to each other, or occur simultaneously. Events themselves may also be described only partially; later information can be added when it becomes available. These features, which they share with the corresponding primitives in a number of other formalisms, such as those of McDermott (1982), Allen (1984) and Lansky (1986), all constitute an advance over temporal representation formalisms based on the situation calculus (McCarthy & Hayes 1969).

Although Kowalski's events are undecomposable points or culminations, they can be used to describe extended events such as our processes, in terms of a *pair* identifying their starting point and to the point at which they stop (in the case of processes) or their culmination (in the case of culminated processes). This means that a process expression like *John ran* will introduce two events, one indicating the start of the process and one indicating the endpoint. Just like the point events considered by Kowalski and Sergot, these events have certain properties or states associated with them. The starting-point of the process referred to by uttering *John ran* marks the beginning of a progressive state that we refer to when we use a progressive like *John is running*, a state which is terminated by the corresponding endpoint event.

This duality between events and states (which was also exploited in Steedman, 1982), is very useful for representing the kind of ontology that we have argued natural language categories reflect. But one shortcoming of Kowalski's event calculus is the absence of other than temporal relations between the events. The best worked out event-based model that takes into account causal as well as temporal relations is Lansky's (1986). The representation she presents is based on

GEM (Lansky & Owicki 1983), a tool for the specification and verification of concurrent programs. GEM reifies events and explicitly represents both their causal and temporal relations. It also provides mechanisms for structuring events into so-called "locations of activity", the boundaries on these locations being boundaries of causal access - as in our episodes. In Lansky (1986), the GEM tool is used to build event-based knowledge representations for use in planners. She suggests the use of three accessibility relations: temporal precedence (<), causality or contingency (@), and simultaneity (\$).

These relations have the following properties:

- < : irreflexive, antisymmetric, transitive
- @ : irreflexive, antisymmetric, intransitive
- \$: reflexive, symmetric, transitive

Because we follow Lansky in making the causality/contingency relation @ intransitive, we avoid certain notorious problems in the treatment of *when*-clauses and perfects, which arise because the search for possible consequences of an event has to be restricted to the *first* event on the chain of contingencies. Thus, when (18a) and (b) are asserted, it would be wrong to infer (c):

- (18) (a) When John left, Sue cried
- (b) When Sue cried, her mother got upset
- (c) When John left, Sue's mother got upset

The reason is exactly the same as the reason that it would be wrong to infer that Sue's mother got upset *because* John left, and has nothing to do with the purely *temporal* relations of these events. It should also be noted that the notion of causality or contingency used here (in line with Lansky's proposals) is weaker than that used in other representation schemes (for example that of McDermott 1982) in that causality is here decoupled from eventuality: if an event A stands in a causal relation to event B, then an occurrence of A will *not* automatically lead to an occurrence of B: John laying the foundations of the house is a prerequisite for or enables him to build the walls and roof but does not "cause" it in the more traditional sense of the word and does not automatically or inevitably lead to him building the walls.³

5. Conclusion

Many of the apparent anomalies and ambiguities that plague current semantic accounts of temporal expressions in natural language stem from the assumption that a linear model of time is the one that our linguistic categories are most directly related to. A more principled semantics is possible on the assumption that the temporal categories of tense, aspect, aspectual adverbials and of propositions themselves refer to a mental representation of events that is structured on other than purely temporal principles, and to which the notion of a nucleus or contingently related sequence of preparatory process, goal event and consequent state is central.

We see this claim as a logical preliminary to the choice of any particular formal representation. However, certain properties of the event-based calculi of Kowalski and Sergot, and of Lansky, seem to offer an appropriate representation for a

³ A Prolog program incorporating the above extension to the event calculus is under construction and will be presented in Moens (forthcoming).

semantics of this kind.

ACKNOWLEDGEMENTS

We thank Ethel Schuster and Bonnie Lynn Webber for reading and commenting upon drafts. Parts of the research were supported by: an Edinburgh University Graduate Studentship; an ESPRIT grant (project 393) to CCS, Univ. Edinburgh; a Sloan Foundation grant to the Cognitive Science Program, Univ. Pennsylvania; and NSF grant IRI-10413 A02, ARO grant DAA6-29- 84K-0061 and DARPA grant N0014-85-K0018 to CIS, Univ. Pennsylvania.

REFERENCES

- Allen, J.. (1984). Towards a general theory of action and time, *Artificial Intelligence*, 23, pp. 123-154
- Dowty, D. (1979). *Word Meaning and Montague Grammar*. Dordrecht: Reidel.
- Dowty, D. (1986). The effects of aspectual class on the temporal structure of discourse. *Linguistics and Philosophy* 9, 37-62
- Hinrichs, E. (1986) Temporal anaphora in discourses of English. *Linguistics and Philosophy* 9, 63-82.
- Kowalski, R. & M. Sergot (1986). A logic-based calculus of events. In *New Generation Computing* 4, 67-95.
- Isard, S.D. (1975). Changing the context. In E. Keenan, (ed.), *Formal Semantics of Natural Language*, London, Cambridge University Press.
- Lansky, A. (1986). A representation of parallel activity based on events, structure and causality. In *Workshop on Planning and Reasoning about Action*, Timberline Lodge, Mount Hood, Oregon, 1986, 50-86.
- Lansky, A. & S. Owicki (1983). GEM: a tool for concurrency specification and verification. In *Proceedings of the Second Annual ACM Symposium on Principles of Distributed Computing*, August 1983, 198-212.
- McCarthy, J. & P.J. Hayes (1969). Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer & D. Michie (eds.) *Machine Intelligence*, Volume 4, 463-502. Edinburgh, Edinburgh University Press.
- McDermott, D. (1982). A temporal logic for reasoning about processes and plans. In *Cognitive Science*, 6, 101-155.
- Moens, M. & M. Steedman (1986). *Temporal Information and Natural Language Processing*. Edinburgh Research Papers in Cognitive Science, II, Centre for Cognitive Science, University of Edinburgh.
- Moens, M. (forthcoming). *Tense, Aspect and Temporal Reference*. PhD dissertation, University of Edinburgh.
- Partee, Barbara H. (1984). Nominal and temporal anaphora, *Linguistics and Philosophy*, 7, pp. 243-286.
- Passonneau, Rebecca J. (1987). Situations and intervals, paper to the 25th Annual Conference of the ACL, Stanford, July 1987.
- Reichenbach, H. (1947). *Elements of Symbolic Logic*. New York, Free Press, 1966.
- Schuster, E. (1986). Towards a computational model of anaphora in discourse: reference to events and actions. Tech. Report CIS-MS-86-34, CIS, Univ. of Pennsylvania.
- Smith, C. (1983). A theory of aspectual choice, *Language*, 59, pp.479-501.
- Steedman, M. (1982). Reference to past time. In *Speech, Place and Action*, R. Jarvella & W. Klein (eds.), 125-157. New York: Wiley.
- Vendler, Z. (1967). Verbs and times. In *Linguistics in Philosophy*, Z. Vendler, 97-121. Ithaca, N.Y.: Cornell University Press.
- Webber, B. (1987). The interpretation of tense in discourse. Paper to the 25th Annual Conference of the ACL, Stanford, July 1987.