

# TEMPORAL ONTOLOGY AND TEMPORAL REFERENCE

Marc Moens

Centre for Cognitive Science  
University of Edinburgh  
Edinburgh, Scotland EH8 9LW

Mark Steedman

Department of AI  
Centre for Cognitive Science  
University of Edinburgh  
and  
Department of Computer and Information Science  
University of Pennsylvania  
Philadelphia, PA 19104-6389

“Two weeks later, Bonadea had already been his lover  
for a fortnight.”

—Robert Musil, *Der Mann ohne Eigenschaften*.

A semantics of temporal categories in language and a theory of their use in defining the temporal relations **between events** both require a more complex structure on the domain underlying the meaning representations than is commonly assumed. **This paper proposes an ontology based on such notions as causation and consequence, rather than on purely temporal primitives.** A central notion in the ontology is that of an elementary event-complex called a “nucleus.” A **nucleus** can be thought of as an association of a goal event, or “culmination,” with a “preparatory process” by which it is accomplished, and a “consequent state,” which ensues. Natural-language categories like **aspects, futurates, adverbials, and when-clauses** are argued to change the temporal/aspectual category of propositions under the control of such a **nucleic knowledge representation structure.** The same concept of a nucleus plays a central role in a theory of temporal reference, and of the semantics of tense, which we follow McCawley, Partee, and Isard in regarding as an anaphoric category. We claim that any manageable formalism 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 is often assumed that the semantics of temporal expressions is directly related to the linear time concept 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, suggested by Ritchie 1979:

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

Copyright 1988 by the Association for Computational Linguistics. Permission to copy without fee all or part of this material is granted provided that the copies are not made for direct commercial advantage and the *CL* reference and this copyright notice are included on the first page. To copy otherwise, or to republish, requires a fee and/or specific permission.

0362-613X/88/010015-28\$03.00

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 unusual character 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 or an enablement relation between the two events. Our knowledge of the world does not easily support such a link for (2), at least if we don't indulge in the fiction that the natural universe is conspiring against the speaker. Nor is the relation predicated between the two events by *when* the one that we normally think of as scientifically causal, for *when* seems to predicate an intransitive relation. Consider:

- 3. a. When John left, Sue cried.
- b. When Sue cried, her mother got upset.
- c. When John left, Sue's mother got upset.

From (3a) and (b) it would be unwarranted to conclude the state of affairs that is described in (c). And this causal aspect of the sentence's meaning must stem from the sense-meaning of *when*, because parallel utterances using *while*, *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) and (2) 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. Nor is it simply causal, as Example 3 shows. We will argue instead that *when* has a single sense-meaning, reflecting its role of establishing a temporal focus, which we follow Isard and Longuet-Higgins (1973) in relating to Reichenbach's reference time (cf. introduction to this collection). 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, but not identical to a notion like causality, rather than mere temporal sequentiality. This contingent, nontemporal relation on the representation of events in episodic memory also determines the ontology of propositions associated with linguistic expressions denoting events and states. It is to these that we turn first.

## 2 TEMPORAL AND ASPECTUAL CATEGORIES

Propositions conveyed by English sentences uttered in context can, following Vendler, be classified into temporal or aspectual types, partly on the basis of the tenses, aspects, and adverbials with which they can co-occur (cf. Dowty 1979, and the introduction to the present collection). The term *aspectual type* refers to the

relation that a speaker predicates of the particular happening that their utterance describes, relative to other happenings in the domain of the discourse. What the speaker *says* about those relations is of course quite distinct from what those relations objectively *are*. In particular, the speaker's predications about events will typically be coloured by the fact that those events are involved in sequences that are planned, predicted, intended, or otherwise governed by *agencies* of one kind or another. For want of some established term to cover this very general class of dependencies between events, we will use the term *contingency*. Thus an utterance of

4. Harry reached the top

is usually typical of what we will call a *culmination*—informally, an event which the speaker views as punctual or *instantaneous*, and as accompanied by a transition to a new state of the world.<sup>1</sup> This new state we will refer to as the *consequent state* of the event. It does not necessarily include *all* events that are objectively and in fact consequences. It rather includes only those consequences that the speaker *views* as contingently related to other events that are under discussion, say by causing them or by *permitting them to occur*. For reasons that are discussed in Section 3.2 below, expressions like these readily combine with the perfect, as in

5. Harry has reached the top.

The point may perhaps best be made by noting that there is another class of punctual expressions that is *not* normally associated with a consequent state. For example,

6. John hiccupped

is not usually viewed as leading to any relevant change in the state of the world. It typifies what we call a *point* expression. A point is an event (not necessarily an instantaneous one) that is viewed as an indivisible whole and whose consequences are not at issue in the discourse—which of course does not mean that *de facto* consequences do not exist. Such expressions are evidently not the same as culminations, for they are rather odd in combination with the perfect, as in

7. #Harry has hiccupped.

The reasons for this will also be discussed below. Sentences like

8. Harry climbed

typify a *third aspectual category*, which we will call for obvious reasons a *process*. Most utterances of such sentences *describe an event as extended in time* but not characterised by any particular conclusion or culmination. As was pointed out by Vendler, expressions like these can be combined with a *for*-adverbial but not with an *in*-adverbial:

- 9. Harry climbed for several hours.  
#Harry climbed in several hours.

In contrast,

- 10. 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 most utterances of such sentences as a fourth aspectual type, called a **culminated process**. Culminated processes, in contrast to ordinary processes, combine readily with an *in*-adverbial but not with a *for*-adverbial.

- 11. Harry climbed all the way to the top in less than 45 minutes.  
#Harry climbed all the way to the top for less than 45 minutes.

All of the above categories describe what common sense suggests we call **events**—that is, happenings with defined beginnings and ends. We distinguish these “hard-edged” categories from a class of **indefinitely extending states of affairs**, which, equally commonsensically, we call **states**. Example 12 typically describes one kind of state:

- 12. Harry is at the top.

Part of the appeal of Vendler’s account, and such descendants as the present proposal, is that it suggests that 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. In **present terms**, the **event-types can be distinguished on just two dimensions**, one concerned with the contrast between **punctuality and temporal extension**, the other with the association **with a consequent state**. This subcategorisation can be summarized as in Figure 1.

|         | EVENTS   |   | STATES                                 |
|---------|--|---|--|
|         | atomic   | extended  |  |
| +conseq | <b>CULMINATION</b><br>recognize, spot,<br>win the race | <b>CULMINATED PROCESS</b><br>build a house,<br>eat a sandwich | understand,<br>love, know,<br>resemble |
| -conseq | <b>POINT</b><br>hiccup,<br>tap, wink                   | <b>PROCESS</b><br>run, swim, walk,<br>play the piano          |  |

Figure 1.

We have included in Figure 1 examples of verbs which *typically* yield propositions of the relevant types, and we shall assume that such verbs (or, strictly speaking,

the associated uninstantiated propositions) are **lexically specified as bearing that type**. However, it cannot be stressed too often that these aspectual profiles are properties of 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 words like “typically” and “readily” above. The details of this taxonomy and the criteria according to which utterances can be categorised are less important than 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.

For example, we shall argue below that the progressive auxiliary demands that **its argument be a process, which it predicates as ongoing**. If it is combined with an **event type that isn’t a process**, say with a punctual event as in *Harry was hiccupping*, then it will cause that original event to be reinterpreted as a process, in this case the process of *iteration* or repetition of the basic event. Similarly, we shall argue that **a perfect auxiliary demands a culmination**, predicating of the time referred to that the associated consequent state holds. The notion of “time referred to” is related to Reichenbach’s **reference time** in Section 4.1 below. If the perfect is combined with an event description for which world knowledge provides no obvious culmination, then the ensemble will tend to be anomalous. So, for example, *Harry has reached the top* is fine, but *The clock has ticked*, and *Harry has hummed*, to the extent that they are acceptable at all, seem to demand rather special scenarios in which the tick of the clock and the mere act of humming have a momentousness that they usually lack.

The phenomenon of change in the aspectual type of a proposition under the influence of modifiers like tenses, temporal adverbials, and aspectual auxiliaries is of central importance to the present account. We shall talk of such modifiers as functions which “coerce” their inputs to the appropriate type, by a loose analogy with type-coercion in programming languages (cf. Ait-Kaci 1984). Thus the effect on meaning of the combination of the progressive with an expression denoting an atomic punctual event as in *Sandra was hiccupping* occurs in two stages: first the point proposition is coerced into a process of iteration of that point. Only then can this process be defined as ongoing, and hence as a progressive state. These two stages might be represented as in the following diagram:

13. (point (*Sandra hiccup*))  
 ↓  
 (process (iteration (point (*Sandra hiccup*))))  
 ↓  
 (progressive (process (iteration (point (*Sandra hiccup*))))))

The temporal/aspectual ontology that underlies the phenomenon of aspectual type coercion can be defined in terms of the transition network shown in Figure 2, in which each transition is associated with a change in the content and where, in addition, the felicity of any particular transition for a given proposition is conditional on support from knowledge and context.




Figure 2.

Rather than attempting to explain this diagram from first principles, we present below a number of examples of each transition. However, it is worth noting first that many 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 consequent state.<sup>2</sup> It can be represented pictorially as in Figure 3:




Figure 3.

Any or all of these elements may be compound: for example, the preparation leading to the culmination of *reaching the top of Mt. Everest* may consist of a number

of discrete steps of climbing, resting, having lunch, or whatever. The consequent state may also be compound; most importantly, it includes the further events, if any, that are in the same sequence of contingently related events as the culmination. Similarly, the culmination itself may be a complex event. For example, we shall see below that the entire culminated process of *climbing Mt. Everest* can be treated as a culmination in its own right. In this case, the associated preparatory process and consequent state will be different ones to those internal to the culminated process itself.

### 3 ASPECT

#### 3.1 THE PROGRESSIVE

According to the present theory, progressive auxiliaries are functions that require their input to denote a process. Their result is a type of state that we shall call a progressive state, which describes the process as ongoing at the reference time. Thus the following sentence, among other meanings that we shall get to in a moment, can simply predicate of a present reference time that the process in question began at some earlier time and has not yet stopped:

14. The president is speaking.

If the input to a progressive is atomic then by definition it cannot be described as ongoing. However, as was noted in the introduction, it may be coerced into a process by being iterated, as in

15. Harry is hiccupping.

There is another route through the network in Figure 2, where the point is coerced into a culmination, i.e., as constituting an atomic event that *does* have consequences associated with it. In this case, the interpretation for (15) parallels the one given for *Harry was reaching the top*, below. However, this particular example is deliberately chosen in order to make that interpretation unlikely.

If a progressive combines with a culminated process, as in:

16. Roger was running a mile

—then the latter must also first be coerced to become a process. The most obvious way to do this is to strip off the culmination and leave the preparatory process behind. It is this process that is stated to be ongoing at the past reference time. Another possible coercion is to treat the entire culminated process as a point, and to iterate it. This interpretation appears to be the one that is forced by continuing (16) as in:

17. Roger was running a mile last week. This week he is up to three.

When a culmination expression like *reach the top* is used with a progressive, it must be coerced to become a process in a slightly more complicated way. The most

obvious path through the network in Figure 2 from the culmination node to the process node involves first adding a preparatory process to the culmination to make it a culminated process, then stripping off the culmination point as before. Thus sentences like the following describe this preparatory process as ongoing at the past reference time:

18. Harry was reaching the top.

Again, an iterated reading is possible in principle, but pragmatically unlikely here.

As a result of the coercions implicit in the last two examples, it is no longer asserted that the culminations in question ever in fact occurred, but only that the associated preparatory processes did. Thus there is no contradiction in continuations that explicitly *deny* the culmination, like:

- 19. a. Harry was running a mile, but he gave up after two laps.
- b. Harry was reaching the top when he slipped and fell to the bottom.

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.

### 3.2 THE PERFECT

A perfect, as in

20. Harry has reached the top

is a function that requires its input category to be a culmination. Its result is the corresponding consequent state. The most obvious of these consequences for (20) is that Harry still *be* at the top, although as usual there are other possibilities. Informal evidence that this indeed is the function of the perfect can be obtained by noticing that perfects are infelicitous if the salient consequences are *not* in force. Thus, when I'm on my way to get a cloth to clean up the coffee I accidentally spilled, I can say

21. I have spilled my coffee.

After cleaning up the mess, however, all the obvious consequences associated with this event seem to be over. In that context, it would be infelicitous to utter (21).

If the input to a perfect is not a culmination, then the perfect will do its best to coerce it to be one, subject to the limitations imposed by contextual knowledge. If the hearer cannot identify *any* relevant consequences, as seems likely for the following example, then coercion may simply fail, in which case a perfect will be infelicitous, as was noted earlier:

22. #The star has twinkled.

To be able to use a culminated process expression like *climbing Mount Everest* with a perfect auxiliary, it first has to be coerced into a culmination. Requiring such a transition might seem unnecessary since a culminated process already implies the existence of a culmination with consequences to which the perfect could refer. But consider Figure 4 as a possible rendering of the nucleus associated with *climbing Mt. Everest*:




Figure 4.

If a perfect could be used to single out the consequences of a nucleus associated with a culminated process expression, then *having climbed Mt. Everest* could be used to refer to the state of having reached the/ summit or being at the top. However, this does not seem to be the case. A reporter who has managed to establish radio contact with a mountaineer who has just reached the top of Mt. Everest is unlikely to ask

23. Have you climbed Mt. Everest yet?

The question rather seems to concern consequences of the culminated process as a whole. We capture this fact by making the perfect coerce the culminated process to become a culmination. The transition network allows this to happen if the entire event of climbing Mt. Everest is treated as a single unit by making it a point, so that it can become a culmination in its own right. The perfect then delivers a rather different kind of consequent state.

A process like *work in the garden* can be coerced by a perfect auxiliary in essentially the same way: the process of working, possibly associated with a culmination point, is treated as a single unit. This pointlike entity can then be used as the starting point for the construction of a new nucleus, by treating it as a culmination in its own right, provided that there are associated consequences. As a result, a question like 24 can only be used felicitously if John's working in the garden was (for example) part of a prearranged plan, or a particular task John had to finish before something else could happen:

24. #Has John worked in the garden?

This account also explains the infelicity of a sentence like (25):

25. #They have married yesterday.

The sentence could only refer to the consequences of *getting married yesterday* as opposed to getting married

some other time. But most of what we think of as consequences of events are *independent* of the specific time at which the event occurred. (In this respect they are different from the preparatory processes, which are argued below to be implicated in certain futurates.) If a certain situation *is* a consequence of an event taking place at a particular time, then a perfect auxiliary may be used to describe that event. Thus a superstitious person believing that disastrous consequences are likely to result from actions performed on an unpropitious date can say:

26. They have married on Friday the 13th!

But even on Saturday the 14th, such a person still cannot use (25), for it would not provide the essential information about the date, thus flouting Grice's maxim of quantity.

The account given here also explains the well-known contrast between the infelicitous (27a) and its felicitous counterpart, (b):

- 27. a. #Einstein has visited Princeton.
- b. Princeton has been visited by Einstein.

Whatever causal sequence of events and their consequences associated with the individual (Einstein) we take to be the one we are currently talking about, (a) cannot be used felicitously to refer to a part of that sequence since all such causal sequences seem to be to do with his enduring consciousness and are therefore by definition over. However, (b) *can* be uttered felicitously to refer to that same event because the relevant causal sequence must be one whose event and consequences apply to the institution of Princeton University (whose corporate consciousness endures) and many such consequences are still in train.

The hypothesis we advance that the perfect has only one temporal meaning has a precedent in the work of Inoue 1979. Moens 1987 has extended the present analysis to show that the distinctions McCawley 1971, 1981 and Comrie 1976 draw between different kinds of perfects (such as "perfect of current relevance," "hot news," "result," etc.) are nothing but different consequent states, depending on the nature of the verbal expression and the particular core event it expresses, and the specific kind of episodes in which our general knowledge tells us such core events typically occur.

### 3.3 ADVERBIALS

*For*-adverbials can only be used felicitously with process expressions:

28. John worked in the garden for five hours.

The resulting combination is a culminated-process expression. Evidence for this can be found in the ease with which an expression like (28) can be combined with a perfect, unlike its process counterpart:

- 29. #John has worked in the garden.
- John has worked in the garden for five hours.

An expression like *playing the sonata* can readily occur with a *for*-adverbial, suggesting that its basic category—by which we mean the type assigned in the lexicon and inherited by the proposition in the absence of any coercion—is that of a process. As a result, (30) carries no implication that Sue *finished* playing the sonata:

30. Sue played the sonata for a few minutes.

Another route through the network is possible in order to account for examples like (30): Sue's playing the sonata, 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 (30) 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

31. Sue played the sonata for about eight hours.

A similar transition path is needed to make sense of examples like the following, in which a culmination is coerced to become a point, and then in turn coerced to become a process by being iterated:

32. John arrived late at work for several days.

The aspectual network would wrongly predict the existence of a *for*-adverbial paradox, parallel to the imperfective paradox, if *for*-adverbials were permitted to freely coerce culminated processes (and hence culminations) to be (not necessarily completed) processes. The theory might seem to wrongly predict that (a) below would mean roughly the same as (b):

- 33. a. #Red Rum won the race for the first few minutes.
- b. Red Rum was winning the race.

However, it is hard to find a context in which (a) means anything at all. The reason for this lies in the way English syntax and morphology control coercion in the aspectual transition network. The transition from culmination to consequent state, for example, demands the presence of a perfect. Similarly, the arc from process to progressive state may be traversed only if a progressive auxiliary is present in the sentence. For other transitions, such as the one resulting in an iterated process or an habitual state, English has no explicit markers and they can be made freely.

The transition from culminated process to process is not one that can be made freely in English, but seems to require the presence of a progressive *-ing*-form. As a result, turning the culmination in (33a) into a process by first adding a preparatory process and then stripping off the culmination point is not allowed. It *is* allowed in (b), but only because the example contains the required progressive *-ing*-form. The only other transition path in the aspectual network that can account for the combination of a culmination with a *for*-adverbial is the one

that turns the culmination into a point, and then iterates it to be a process. This interpretation is not felicitous for (33a), either, given our knowledge about what constitutes winning a race. However, as with (32), it is acceptable for

34. Nikki Lauda won the Monaco Grand Prix for several years.

Sometimes, a *for*-adverbial in combination with a culmination seems to describe a time period *following* the culmination rather than an iterated process:

35. John left the room for a few minutes.

This adverbial is of a different kind, however, expressing intention rather than duration. It is merely by accident that English uses the same device to convey these different meanings. In French or German, for example, the two constructions are clearly distinct, as shown in the following translations of (35) and (32):

36. Jean a quitté la chambre *pour* quelques minutes.  
Johann verliess *für* einige Minuten das Zimmer.

37. *Pendant des années* Jean est arrivé en retard au travail.  
*Jahrelang* erschien Johann zu spät zur Arbeit.

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, as do related phrases like “it took me two days to . . . .” This means that combination with a culmination expression requires a transition to the culminated process node. According to the aspectual network in Figure 2 this transition is felicitous if the context allows a preparatory process to be associated with the culmination, as in (38):

38. Laura reached the top in two hours.

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

Since the arcs describe how one must be able to view the world 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. Indeed, sentences like the following only seem to be made tolerable to the extent that it is possible to conjure up contexts in which the event only *appears* to be accidental:

39. In fifteen minutes, Harry accidentally spilled his coffee.

A similar problem arises in connection with the following example:

40. 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 unspecific *in a few minutes* continues to resist this interpretation.

However, another route is also possible for (40): 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.

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:

41. It took me two days to play the “Minute Waltz” in less than sixty seconds for more than an hour.

The process expression *play the Minute Waltz* is coerced by the *in*-adverbial into a culminated process, including a culmination of *finishing playing the Minute Waltz*. Combination with the *for*-adverbial requires this expression to be turned into a process—the only possible route through the network being that through the point node and iterating. 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 would find 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 (41), 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 that we can think of as *practicing*. The phrase *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. We assume that the ordering of these successive coercions, like others

induced by the perfect and the progressive, are (not necessarily unambiguously) under the control of syntax.

#### 4 TENSE AND TEMPORAL FOCUS

##### 4.1 TENSE

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, and by contextual knowledge. 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. This is done by using a second type of “temporal” adjunct, 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. Under these accounts, temporal adjuncts establish a referent 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 and definite anaphors (see foreword).

Reichenbach's account is somewhat inexplicit as far as extended, noninstantaneous events go. In particular, he makes it look as though the reference time is always an instant. However, we believe that the following account is the obvious generalisation of his and probably what he intended anyway.

In Reichenbach's system a simple past tense of an atomic event is such that **reference time (R)** and **event time (E)** are identical, while progressives and perfects are such that R and E are *not* identical.<sup>3</sup> The only coherent generalisation of his scheme to durative events is to maintain this pattern and assume that R and E are coextensive for an utterance like:

42. Harry ran a mile.

It follows that R may be an extended period (cf. Steedman 1982). R may also be an extended period for a state such as a progressive, although in this case the corresponding event time is still quite separate, of course.

What is the nature of this referent, and how is it established? The anaphoric quality of tense has often been specifically compared to *pronominal* anaphora (cf. McCawley 1971; Partee 1973; Isard 1974). 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 the successive conjoined main clauses seems to move on from the time originally established by the adjunct:

43. At exactly five o'clock, Harry walked in, sat down, and took off his boots.

Nor is this just a matter of pragmatic inference; other orders of the clauses are not allowed:

44. #At exactly five o'clock, Harry took off his boots, sat down and walked in.

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 the reference time 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 that the reference time advances and by how much. 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 comprises 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. If the context does not support this interpretation, then the temporal referent will be unchanged, as in:

45. At five o'clock, my car started and the rain stopped.

In its ability to refer to temporal 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 (e.g., *the music* in the following example) than like a pronoun, as Webber 1987 points out.

46. I went to a party last night. The music was wonderful.

##### 4.2 WHEN-CLAUSES

The definite nature of tense together with the notion of the nucleus as the knowledge structure that tensed expressions conjure up explain the apparent ambiguity of *when*-clauses with which this paper began. A *when*-clause behaves rather like one of those phrases that are used to explicitly change topic, such as *and your father* in the following example (cf. Isard 1975):

47. And your father, how is he?




A *when*-clause does not require a previously established temporal focus, but rather brings into focus a novel temporal referent whose unique identifiability in the hearer's memory is presupposed. Again, the focused temporal referent is associated with an entire nucleus, and again an event main clause can refer to any part of this structure conditional on support from general or discourse specific knowledge. For example, consider again Example 1 with which we began (repeated here):

- 48. 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 two alternative routes to construct a complete nucleus:

- 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 building or to the consequent state of having built the bridge; or
- b) to treat the entire event as a single culmination and *compose* it into a nucleus with whatever preparation and consequences the context provides for the activity of building a bridge, and to make the transition to either one of those.


Either 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 48a, the entire culminated process of building the bridge tends to become a culmination (via a path in Figure 2 that passes through the point node), which is associated in a nucleus with preparations for, and consequences of, the entire business, as in Figure 5:



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, people tend to see the building of the bridge as decomposed into a quite different preparatory process of building, a quite different culmination of completing the bridge and some consequences that we take to be also subtly distinct

from those in the previous case as was argued in Section 3.2. The resulting nucleus is given in Figure 6. The use of the best materials is then, as in (a), situated in the preparatory process—but it is a different one this time.



Example c is like (a) in giving rise to the nucleus in Figure 5, but pragmatics seems to demand 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, repeated here, is still strange, because it is so hard to think of *any* relation that is supported in this way:

- 49. #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 breakdown, the breakdown 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 breakdown, 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 clauses. 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. With stative main clauses, as in the following examples, the interpretation strategy is somewhat different. Statives show no sign of being related under what we are calling contingency, presumably because contingency is by definition a relation over events. In particular, they do not enter in a causal or contingent relation with a *when*-clause the way corresponding sentences with events as main

clauses do. They therefore merely predicate that the state in question holds at the time of the culmination:

50. 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, a stative main clause can be turned into an event expression; in that case, a contingency relation *is* predicated to exist between the two events. Thus the following example seems to involve an inceptive event, which *begins* the state of knowing:

51. When Pete came in, I knew that something was wrong.

Such changes of type are similar to others discussed above but are not treated in the present paper.

### 5 REFERRING TO FUTURE EVENTS

Bennett and Partee 1972, speaking of the difference between the present perfect and the simple past, remark that one might expect a similar distinction among future tenses. One could conceive of a construction parallel to the perfect, whose event time would be in the future and whose reference time would be the time of speech, conveying a notion of current relevance; and there could be a construction parallel to the simple past, with both reference and event times in the future. Bennett and Partee suggest that English is *not* as one would expect and follow Reichenbach in saying that these two functions are conflated in a single device, the modal future using *will*. Although it is true that the modal future shares features of both perfect and simple past, it is nevertheless also the case that there are two classes of futurate expressions, with properties parallel to each of the two past expressions.

The candidate for the role parallel to the perfect is the so-called futurate progressive (Smith 1983):

52. Robert was working on the speech project until he got a job offer from Sussex.

As Dowty 1979, 1986 argues, examples like (52) can be both a past imperfective progressive (answering a question about Robert's past activities) and a past futurate progressive (answering a question about Robert's plans at some past time and meaning something like *Robert was going to work on the speech project, but he didn't*). However, the difference between the two interpretations seems to be a matter of pragmatic world knowledge rather than sense-semantics, corresponding to the two different ways of constructing a nucleus (cf. Section 4). The imperfective progressive decomposes the core event into a nucleus and makes a transition to the preparatory process, indicating that it is in progress at

the time of reference. The futurate progressive, through the use of an adverbial signaling an event time posterior to the reference, forces the whole event to be treated as a single unit, which is then composed into a new nucleus. The progressive then indicates that the preparation leading up to the event *as a whole* was in progress at the time of reference (as usual, without asserting that that event or even its onset was ever reached). The futurate progressive thus resembles the perfect in saying something about a (past or present) reference time that is entirely separate from the event time.

The candidate for the role parallel to the simple past among the futurates is to be found in the simple, or non-modal future, sometimes (confusingly) called the tenseless future:

53. He leaves on Tuesday.

While the futurate progressive shares with the perfect the property of needing no nonpresent adverbial, the nonmodal future cannot be used in this way. For example, in response to a question about the current state of affairs as specific as *Why are you being so rude to your boss these days?* or as general as *What's new?*, one may respond with an unanchored progressive (54a), much as with a perfect (54b). But one may not reply with an unanchored nonmodal future (54c), although an anchored one (54d) is quite all right.

54. a. I am leaving.  
 b. I have handed in my notice.  
 c. \*I leave.  
 d. I leave next month.

In its requirement for an established non-present reference time, the nonmodal future resembles the past tense. The resemblance (which was noted in Leech 1971) is supported by the following further observations. A *when* question concerning the past progressive is ambiguous, reflecting the separation of reference time and event time. By contrast, the nonmodal future does not really seem to occur in the past at all, except of course in reported or indirect speech; it just becomes indistinguishable from the simple past. It follows that (55) can be answered with (a) or (b). But (56) can only be answered with (a), not with (b).

55. When were you leaving?  
 a. Last week (ambiguous).  
 b. Next week.
56. When did you leave?  
 a. Last week (unambiguous).  
 b. \*Next week.

These similarities suggest the symmetry depicted informally in Figure 7 between the perfect, the simple past, the futurate progressive, and the nonmodal future. The hatching again informally indicates the extent of the consequent state and the preparatory process associated with the perfect and the futurate progressive,



state of the world at speech time. These senses are exhibited below.

62. a. You will be my long-lost brother Willy.  
 a'. You will marry a tall dark stranger.  
 b. You must be my long-lost brother Willy.  
 b'. You must marry a tall dark stranger.  
 c. You may (or may not) be my long-lost brother, Willy.  
 c'. You may (or may not) marry a tall dark stranger.

But, as has often been suggested before, the future epistemic modals have nothing to do with future tense in the strict sense of the word.<sup>4</sup>

## 6 TOWARD 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 organise 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 database 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 might 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 in Steedman 1982, using a program that verified queries against a database structured according to some of the principles outlined above; a more recent extension of this work was reported in Moens 1987. Events are stored as primitives in the database, possibly but not necessarily associated with a time point. Extended events are represented in terms of a pair of punctual events, identifying their starting point as well as the point at which they end (in the case of processes) or culminate (in the case of culminated processes).

Apart from the obvious accessibility relations of temporal precedence and simultaneity, events can also enter into the relation of contingency introduced above. It is significant that the relation used in the implementation is identical to the notion of causality used by Lansky 1986 in an entirely different problem area. She developed a knowledge representation scheme for use in planners in which events are reified and modeled with an explicit representation of their temporal as well as causal relations. In this scheme, a mechanism is provided for structuring events into so-called "locations of activity", the boundaries of which are boundaries of "causal" access. As a result, two events with no causal relation between them cannot belong to the same location of activity—as in the episodes introduced above.

Because we follow Lansky in making the 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 (3) is asserted, repeated here as (63a) and (b), it would be wrong to infer (c):

63. 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 contingency used here (in line with Lansky's proposals) is weaker than the notion of causality used in other representation schemes (for example, that of McDermott 1982 or Allen 1984): if Event A stands in a contingent 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.

The transitions in the network are implemented as inference procedures in the database. Answering a query involving the aspectual auxiliaries and adverbials

discussed before consists of finding a matching event description in the database and checking its aspectual type; if the event description is found not to have the required aspectual type, it can be changed by means of the inference procedures, provided such a change is supported by knowledge in the database about the event in question.

## 7 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.

## ACKNOWLEDGMENTS

We thank Jon Oberlander, 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. An earlier version of some parts of this paper was presented as Moens and Steedman 1987.

## REFERENCES

- Ait-Kaci, H. 1984 A Lattice Theoretic Approach to Computation Based on a Calculus of Partially Ordered Type Structures. Ph.D. thesis, University of Pennsylvania.
- Allen, J. F. 1984 Towards a General Theory of Action and Time. *Artificial Intelligence* 23:123–154.
- Bennett, M. and Partee, B. 1972 *Toward the Logic of Tense and Aspect in English*. Indiana University Linguistics Club, Bloomington, IN.
- Boyd, J. and Thorne, J. 1969 The Semantics of Modal Verbs. *Journal of Linguistics* 5:57–74.
- Comrie, B. 1976 *Aspect: An Introduction to the Study of Verbal Aspect and Related Problems*. Cambridge University Press, Cambridge, England.
- Dowty, D. 1979 *Word Meaning and Montague Grammar*. D. Reidel, Dordrecht, W. Germany.
- Dowty, D. 1986 The Effects of Aspectual Class on the Temporal Structure of Discourse: Semantics or Pragmatics?. *Linguistics and Philosophy* 9:37–61.
- Hinrichs, E. 1986 Temporal Anaphora in Discourses of English. *Linguistics and Philosophy* 9:63–82.
- Inoue, K. 1979 An Analysis of the English Present Perfect. *Linguistics* 17:221–222, 561–590.
- Isard, S. D. 1974 What would you have done if . . . *Theoretical Linguistics* 1:233–255.
- Isard, S. D. 1975 Changing the Context. In Keenan, E. (ed.) *Formal Semantics of Natural Language*. Cambridge University Press, Cambridge, England.
- Isard, S. and Longuet-Higgins, C. 1973 Modal Tic-tac-toe. In Bogdan, R. J. and Niiniluoto, I. (eds.) *Logic, Language and Probability*. Reidel, Dordrecht, W. Germany:189–195.
- Lansky, A. L. 1986 A Representation of Parallel Activity Based on Events, Structure, and Causality. In *Workshop on Planning and Reasoning about Action*. Timberline Lodge, Mount Hood, OR: 50–86.
- Leech, G. N. 1971 *Meaning and the English Verb*. Longman, London.
- McCarthy, M. 1974 Sons of the Morning. In *The Seventeenth Degree*. Harcourt Brace Jovanovich, New York, NY.
- McCawley, J. D. 1971 Tense and Time Reference in English. In Fillmore, C. and Langendoen, T. (eds.) *Studies in Linguistic Semantics*, Holt, Rinehart and Winston, New York, NY: 96–113.
- McCawley, J. D. 1981 Notes on the English Present Perfect. *Australian Journal of Linguistics* 1:81–90.
- McDermott, D. 1982 A Temporal Logic for Reasoning about Processes and Plans. *Cognitive Science* 6:101–155.
- Moens, M. 1987 Tense, Aspect and Temporal Reference. Ph.D. thesis. Centre for Cognitive Science, University of Edinburgh, Edinburgh, Scotland.
- Moens, M. and Steedman, M. 1986 Temporal Information and Natural Language Processing. Research Paper/RP-2. Center for Cognitive Science, University of Edinburgh, Edinburgh, Scotland.
- Moens, M. and Steedman, M. 1987 Temporal Ontology in Natural Language. In *Proceedings of the 25th Annual Meeting of the Association for Computational Linguistics*. Stanford University, Stanford, CA: 1–7.
- Palmer, F. 1974 *The English Verb*. Longman, London.
- Partee, B. 1973 Some Structural Analogies Between Tenses and Pronouns in English *Journal of Philosophy* 70:601–609.
- Partee, B. 1984 Nominal and Temporal Anaphora. *Linguistics and Philosophy* 7:243–286.
- Passonneau, R. J. 1987 Situations and Intervals. In *Proceedings of the 25th Annual Meeting of the Association for Computational Linguistics*, Stanford University, Stanford, CA: 16–24.
- Reichenbach, H. 1947 *Elements of Symbolic Logic*. Macmillan, London.
- Ritchie, G. D. 1979 Temporal Clauses in English. *Theoretical Linguistics* 6:87–115.
- Sag, I. 1973 On the State of Progress on Progressives and Statives. In Bailey, C. and Shuy, R. (eds.) *New Ways of Analyzing Variation in English*. Georgetown University Press, Washington, DC: 83–95.
- Smith, C. S. 1983 States and Dynamics. *Language* 59:479–502.
- Steedman, M. J. 1977 Verbs, Time and Modality. *Cognitive Science* 1: 216–234.
- Steedman, M. J. 1982 Reference to Past Time. In Jarvella, R. and Klein, W. (eds.) *Speech, Place and Action*. John Wiley and Sons, New York, NY: 125–157.
- Vendler, Z. 1967 Verbs and Times. *Linguistics in Philosophy*. Cornell University Press, Ithaca, NY: chap. 4, 97–121.
- Verkuyl, H. J. 1972 *On the Compositional Nature of the Aspects*. D. Reidel, Dordrecht.
- Vetter, D. 1973 Someone Solves This Problem Tomorrow. *Linguistic Inquiry* 4:104–108.
- Webber, B. L. 1987 The Interpretation of Tense in Discourse. In *Proceedings of the 25th Annual Meeting of the Association for Computational Linguistics*. Stanford University, Palo Alto, CA: 147–154.
- Wekker, H. 1976 *The Expression of Future Time in Contemporary British English*. North Holland, Amsterdam.

## NOTES

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. The new nomenclature is also intended to reflect the fact, also noted by Dowty (1979), that Vendler's "accomplishments," which we will refer to as "culminated processes," are composite events, consisting of a process which is associated with a particular culmination point.
2. A similar tripartite event structure is proposed in Passonneau (1987, cf. this volume).
3. In attributing this view to Reichenbach, we are assuming that there is an oversight or a misprint in his diagram for the past progressive, p. 290: the diagram seems to suggest that R and E are coextensive, whereas what is intended is that the punctual reference time is included in an extended event time, as in his diagram for the present progressive. We also ignore here one of his analyses of the modal future, which we regard as incorrect (cf. Section 5).
4. It is an implication of such an analysis that there should be no truly past version of epistemically modal propositions. Where past tenses of the epistemic modals do occur, they must, like the past nonmodal future, always be either counterfactual or indirect or reported speech. This seems to be the case. Mary McCarthy (1974), speaking of David Halberstam's use in *The Best and the Brightest* of "what she could only describe as the Future Past," as in
  - i. At a dinner party after the Bay of Pigs Bundy would tell friends . . .
 and
  - ii. The power and prestige that the McNamara years would bring . . .
 called it "that awful tense, seeming to endow the author with prophetic powers," signifying "a future already plangent when it has not yet happened." The source of that awful power (which also accrues to the past tenses of the nonmodal future and, as McCarthy also remarks, the modal-like "be to" construction), is of course the shifting of the speech or consciousness time into the past, rather than the reference time.