The Productions of Time

The Productions of Time: Temporality and Causality in Linguistic Semantics

Mark Steedman

Draft 5.0, May 31, 2005

Division of Informatics
University of Edinburgh
2 Buccleuch Place
Edinburgh EH8 9LW
Scotland UK
(steedman@inf.ed.ac.uk)

Eternity is in love with the productions of time

— William Blake, *Proverbs of Hell*.

Contents

Preface

This book is not primarily about time. Rather, it is about some more primitive conceptual representations of causal and teleological dependency among events that our understanding of time and the semantics of natural language categories like tense are built upon. These draft lecture notes represent an attempt to identify those parts of logic and computer science that are most useful for representing natural human reasoning about times, events, and causality, with particular attention to the problem of constructing a causal and temporal semantics for English and other languages, of the kind that will support effective inference by automatic theorem provers. In particular, they seek to adapt the insights of computer science and artificial intelligence concerning programing language semantics and robot planning to this purpose, using the formal devices of Dynamic and Linear Logic. It is important to begin by reflecting on how the apparently quite different approaches from computation and logic are related.

Logicians and to a lesser extent computer scientists have many concerns that are orthogonal to problems of knowledge representation. Some of these concerns make the most familiar logics rather unsuitable for the purposes of knowledge representation. My friend Kit Fine has compared the position of the linguist or artificial intelligencer who turns to logic for this purpose to that of a man in need of trousers who goes to a tailor, only to be told that tailors only make jackets, and that in fact only jackets are necessary, for it is easy to show that jackets are topologically equivalent to trousers. Such is the authority of logicians that many otherwise decorous persons have found themselves in the position of trying to use jackets as trousers. When they have complained that jackets don't seem to work very well for the purpose—for example, that the pockets seem to be the wrong way up—the response has often been impatient.

Sometimes the users have been led to give up on logic entirely and to go off and invent their own knowledge representations. However serviceable these have been, they have often been derided by logicians as outlandish and even indecorous (perhaps the kilt is the metaphor I need). This is a shame, because in the end one's trousers are best made by tailors, and logicians are or ought to be the right people to make knowledge representations.

These notes represent an attempt to bring together the concerns of those who really do need trousers with those who are best fitted to provide them, and to explain why the *sans culottes* on occasion voice needs (or offer alternatives of their own devising, such as "nonmonotonic" logic or "promiscuous" ontologies) that strike logicians as bizarre and uncouth. They are also an attempt to show what parts of logic proper can be most easily adapted to those needs.

The use of logic for practical knowledge representation can be quite distant from the usual concerns of logicians, whose more mathematical concerns encourage attention to model theory, and to logics with minimal ontologies or type systems and minimal numbers of inference rules. In the case of logics of time, the ontology has typically been restricted to ordinals or points on the real number line. Sometimes, in response to some very persuasive (but ultimately misleading) intuitions people have about the way they think about the world, intervals have been included as elementary types, but that has often been the extent of the concession.

The real number line is probably a fine representation of time for physicists, who are simply concerned to reason about the continuum. The continuum does not itself have any of the structure that our minds impose by virtue of being an (in our view rather special) part of it. But as soon as we are concerned in any sense with representing the way that we think and talk about the world, or with enabling machines to reason with comparable efficiency about systems as complex as the ones that we have to deal with in order to function in the world, such a representation is far too low-level. What is more, linguists in the tradition of Vendler, Dowty, Bach, and Verkuyl have good arguments for the existence of a rather rich variety of event- and situation-types that crop up fairly transparently in language after language, and which therefore seem likely to reflect more or less directly the underlying conceptualization of temporality. Artificial Intelligence researchers like Allen (1984); Allen and Ferguson (1994) have noticed that these event types seem to be rather useful in reasoning about practical action in the world, of the kind that reliably gets you around in a rapidly changing real world.

Nor is model theory necessarily the first priority. Soundness, at least, is an important property for any knowledge representation. But, having noticed that the arithmetic of real numbers appears to provide a basis for grounding temporal representations in the continuum, the linguist or artificial intelligence builder is likely to take the soundness of this component of the model for granted, as they do most of the time for ordinary arithmetic calculation, and to want to press on to develop on that basis equally sound proof-theoretic representations that are sufficiently rich to represent and support inferences about action in the world.

The proposal is that the so-called temporal semantics of natural language is not primarily to do with time at all. Instead, the formal devices we need are those related to representation of causality and goal-directed action. Chapters 1 to 4 of the book review these questions and develop a variant of the situation or event calculus of McCarthy, Hayes, Kowalski and others which I will call the Dynamic Event Calculus. Chapter 5 (which at the time of this draft remains in more preliminary form) then broadens the discussion to defend this proposal in comparison with some alternatives that have been proposed by artificial intelligence researchers, and attempts to counter some criticisms that have been mounted against modal systems in general and the situation/event calculi in particular. A brief conclusion returns to broader concerns with linguistics and human cognition.

These draft lecture notes represent work in progress. Early versions of various parts of chapters 2 to 4 appeared in substantially different formalisms as Steedman (1995, 1997, 2004, 2002b), and as the notes for a tutorial at the 2nd International Conference on Temporal Logic, Manchester July 1997, where it appeared under the Raymond Carveresque title "What we Talk About When We Talk about Time" . I am grateful to Tim Fernando, Kit Fine, Stephen Isard, Marc Moens, Charlie Ortiz, Ian Pratt-Hartmann, Matthew Stone, Rich Thomason, Bonnie Webber, and Michael White for advice and help with concepts and formalisms, and to Johan van Benthem, Pat Hayes, David Israel, Mark Johnson, Alex Lascarides, Joyce McDonough, Alice ter Meulen, Jong Park, Denise Perrett, Len Schubert, and the tutorial attendees at ICTL97 for comments and criticism at a number of stages. They are not to blame for any errors that remain. Support for various aspects of the work was provided in part by NSF grant no. IRI95-04372, DARPA grant no. N660001-94-C-6043, and ARO grant no. DAAH04-94-G0426 to the University of Pennsylvania, and by ESRC grant M423284002 and EPSRC grants GR/M96889 and GR/R02450 to the University of Edinburgh.

Chapter 1

What We Talk About When We Talk About Time

> *We don't know what we talk about when we talk about love.*
> Raymond Carver, *What We Talk About When We Talk About Love*

In thinking about the logical and computational semantics of temporal categories in natural languages, issues of temporal *ontology*, or metaphysics, must be distinguished from issues of temporal *relation*. Categories of the first kind determine the sorts of temporal entities that can be talked about—examples that are discussed below include various kinds of states and events. We shall be concerned with what Cresswell (1990), following Quine (1960), calls the "ontological commitment" of the semantics—that is, the variety of types that can be quantified over, or otherwise formally operated upon. Temporal relational categories determine the relations that may be predicated over such entities—examples to be discussed include temporal order, inclusion and overlap, together with various causal, teleological, and epistemic relations. Some of these relations depend for their identification upon inference from discourse structure and context. It follows that we must distinguish a third kind of phenomenon, that of temporal *reference*. These three distinct but interrelated kinds of phenomena are considered in turn in the three chapters that follow.

As in any epistemological domain, neither the ontology nor the relations should be confused with the corresponding descriptors that we use to define the physics and mechanics of the real world. The notion of time that is reflected in linguistic categories is only indirectly related to the common-sense physics of clock-time and the related Newtonian representation of it as a dimension comprising an infinite number of instants corresponding to the real numbers, still less to the more abstruse representations of time in modern physics.

This observation may not seem too surprising, since it is only a more extreme version of Russell and Weiner's observation of the need to distinguish between external and individual representations of time. However, the particular conceptualisation of temporality that underlies language is by no means obvious. Like the concept of an entity or individual it is confounded with practical aspects of our being in the world of a kind that physics does not discuss. In particular, it is confounded with notions of teleology that are explicitly excluded from even the most informal and common-sense varieties of physics. On the assumption that linguistic categories are fairly directly related to underlying conceptual categories (for how else could children learn them), it is to the linguists that we must turn for insights into the precise nature of this ontology.

In this connection it may seem surprising that the present paper is confined to analyses of English temporal categories. However, it will soon be apparent that we cannot analyse the categories of English without appealing to notions of underlying meaning that are closely related to a level of knowledge about events that is independent of the idiosyncracies of any particular language. The paper returns briefly to the question of the universality of this semantics in the conclusion.

Because of this psychological grounding of the natural semantics of temporality, a certain caution is appropriate in assessing the relevance to linguistic inquiry of systems of logic and computational theory that trade under names like "Tense Logic". Such logics frequently come with very minimal ontologies, restricted to states and Newtonian instants, or to the simplest kind of interval, and similarly minimal, purely temporal, relations among them. Their authors are usually careful to stress that their systems do not reflect linguistic usage. Their *raison d'être* is analogous to that of Peano's axioms in arithmetic—that is, to characterise the metamathematical properties of physical time. Such concerns are not necessarily those of the working linguist or computational linguist, who is mainly interested in performing inference. One does not calculate via proofs in Peano arithmetic.

Many properties of natural language semantics, particularly those involving the notion of discourse *context*, are most directly modelled by dynamic processes. Since computer programs are a very direct expression of procedures, many of the logical frameworks that we shall find most useful draw upon ideas from computer science and studies in artificial intelligence as frequently as from the declarative logical tradition itself. In particular, many recent formalisms invoke the computer scientist's concept of a *side-effect* or update to a database, in order to talk about the changing context of reference, including the temporal variety. This move introduces notions of non-monotonicity, of a kind discussed by Thomason (1997).

We shall combine this notion with the modal logicians' device of an *accessibility relation*, defining a structure on models, where models are databases, or *partial* models, in what has come to be called *dynamic logic*. Crucially, the accessibility relation in a dynamic logic is defined in terms of the primitive events that change one state of the world into another, rather than in terms of the states themselves. This tactic provides a basis for logics in which the structure of the space which must be searched to find constructive proofs and the structure of the proofs themselves are similar, in contrast to many other modal approaches. Such a property is essential for purposes of knowledge representation and practical inference. It is achieved by incorporating into the logic itself the inertial properties of the domain that have been exploited from earliest times in AI approaches to the problem of representing and reasoning about change. This last step uses a more recent notion from logic of a "resource-dependent" logic, as embodied in Girard's 1995 Linear Logic. (This component was not involved in the present author's earlier attempts on the problem of temporal semantics, and was advanced under the name of the Linear-Dynamic Event calculus in Steedman 2002a and 2002b.

In developing an account of the very diverse and ramifying literature that the present proposal builds on, it will sometimes be necessary to concentrate on one of these approaches, and there may be a danger of temporarily losing sight of the others. Nevertheless, they will meet up again as the book proceeds, for linguists, computer scientists and logicians are linked in this venture like mountaineers roped together during a climb. Sometimes the lead is taken by one, and sometimes another, but progress will in future, as in the past, only be made by the team as a whole.

Chapter 2

Temporal Ontology

2.1 Basic Phenomena and Descriptive Frameworks

The first thing to observe about the temporal ontology implicit in natural languages is that it is not purely temporal. To take a simple example, the English perfect, when predicated of an event like *losing a watch*, says that some contextually retrievable *consequences* of the event in question hold at the time under discussion. (Such consequences have sometimes been described under the heading of "present relevance" of the perfect—cf. Inoue (1979). In restricting the perfect to this single meaning, English differs from most other European languages, in which the perfect also acts as a past tense.) Thus, conjoining such a perfect with a further clause denying those consequences is infelicitous:

(1) I have lost my watch (# but I have found it again)

In this respect the English perfect (unlike the perfect in many other languages) stands in contrast to the more purely temporal tenses, such as the past, which make no comparable claim about the consequences of the core event:

(2) Yesterday, I lost my watch (but I (have) found it again)

Further evidence for the claim that the perfect is concerned with causal effects or consequences, and that the availability of such "contingencies" depends upon world knowledge is provided by examples like the following. Example a, below, is one in which no obvious consequences are forthcoming from the knowledge base. Example b is one in which all the obvious consequences of the core event are consequences *for Einstein*, which our knowledge tells us cannot still hold. Both examples are therefore anomalous unless supported by rather unusual contexts.

(3) a. #I have breathed
    b. #Einstein has visited New York

It is because categories like the perfect are not purely temporal that it is usual to distinguish them from the tenses proper as "aspects". Another aspect whose meaning is not purely temporal is the progressive or imperfective. The predication that it makes concerning the core event is a subtle one. While the progressive clearly states that *some* event is ongoing at the time under discussion, it is not necessarily the event that is actually mentioned. Thus in a, below, there seems to be a "factive entailment" to the effect that an event of writing actually occurred. But in b, there is no such entailment concerning an event of writing a sonnet, for b is true even if the author was interrupted before he could complete the action.

(4) a. Keats was writing $\models$ Keats wrote
    b. Keats was writing a sonnet $\not\models$ Keats wrote a sonnet

Dowty (1979) named this rather surprising property of the progressive the "imperfective paradox", and we shall return to it below. It reflects the fact that events like *Keats writing*, unlike those like *Keats writing a sonnet*, are what White (1994) calls *downward entailing*, which we can define as follows:

(5) A proposition ϕ holding of an interval t is downward entailing if it entails that ϕ also holds of all subintervals of t down to some reasonable minimum size.

The imperfective paradox is the first sign that we must distinguish various types or sorts of core event in natural language temporal ontology.

The key insight into this system is usually attributed to Vendler (1967), though there are precedents in work by Jesperson, Kenny and many earlier authorities including Aristotle. Vendler's taxonomy was importantly refined by Verkuyl (1972, 1989), and Dowty (1979, 1982), and further extended by Hinrichs (1985, 1986), Bach, Brown and Marslen-Wilson (1986), Moens and Steedman (1987), citeMitt:88, Smith (1991), Krifka (1989, 1992), Jackendoff (1991), and White (1994). The following brief summary draws heavily on their work.

Vendler's original observation was that a number of simple grammatical tests could be fairly unambiguously applied to distinguish a number of distinct aspectual categories. The term "aspectual" here refers to the intrinsic temporal profile of an event, and such categories are to be distinguished from the *sentential* aspects, the perfect and the progressive. For this reason they are often referred to under the German term *Aktionsarten*, or action-types. Vendler talked of his categorisation as a categorisation of *verbs*, but Verkuyl and Dowty argued that it was properly viewed as a classification of the *propositions conveyed by* verbs and their arguments and adjuncts—that is, of propositions concerning events and states.

We will consider just four tests used by Vendler and those who followed, although there are others. The first is compatibility with adverbials like *for fifteen minutes*. The second is compatibility with adverbials like *in fifteen minutes* and the related construction *It took (him) fifteen minutes to* ... . The third is the entailment arising from the progressive. The fourth is compatibility with the perfect.

Vendler identified a category of event such as *arriving*, *reaching the top* or *finishing a sonnet*, which he called *achievements*. These events are characterised by being instantaneous, and by resulting in a distinct change in the state of the world. They can be detected by the fact that they combine happily with *in*-adverbials, do not combine with *for*-adverbials, do not carry a factive entailment under the progressive, and combine happily with the perfect.

(6) a. Keats finished the sonnet in fifteen minutes.
b. #Keats finished the sonnet for fifteen minutes.
c. Keats is finishing the sonnet ($\not\models$ Keats will have finished the sonnet).
d. Keats has finished the sonnet.

Achievements are to be contrasted with a category of events like *walking, climbing* and *writing*, which Vendler called *activities*. Activities are extended in time, and do not seem to result in any very distinct change in the state of the world. They can be detected by the fact that they combine with *for*-adverbials but not with *in*-adverbials, that the progressive does carry a factive entailment, and that they are distinctly odd with the perfect.[1]

(7) a. Keats wrote for fifteen minutes.
b. #Keats wrote in fifteen minutes.
c. Keats is writing ($\models$ Keats will have written).
d. #Keats has written.

[1] Vendler's account leaves it unclear whether the term *activity* corresponds to past tensed expressions like *wrote* or to the corresponding progressives. I shall use the term to refer to an event type, while the corresponding progressives will denote a variety of *state* as in Steedman (1977) (cf. Webber 1978).

Both of these categories are to be contrasted with a third category of event such as *writing a sonnet* or *flying to Paris*. Vendler called such events *accomplishments*. They superficially have the same test profile as achievements:

(8) a. Keats wrote *In Disgust of Vulgar Superstition* in fifteen minutes.
b. #Keats wrote the sonnet for fifteen minutes.
c. Keats is writing the sonnet ($\not\models$ Keats will have written the sonnet).
d. Keats has written the sonnet

(See (Garrod, 1954, p.532) for some historical background to this example). However, accomplishments differ from achievements in being extended in time, like activities. As a consequence, they differ in entailments when combined with *in*-adverbials and progressives. In 8a and c it is *part of* the event (namely the writing) that respectively takes fifteen minutes and is reported as in progress. It is precisely *not* part of finishing itself that takes fifteen minutes in 6a, or is in progress in 6c. It is some other event. In fact it is presumably an event of writing, since the overall entailments of the two pairs of sentences are very similar. Because of this relation, both Verkuyl and Dowty proposed that accomplishments should be regarded as composites of an activity and a culminating achievement.

Vendler also identified a class of *states*. States are characterised syntactically by being almost the only propositions that can be expressed in English by simple present tense. (The exceptions are performatives like the following, which in all other respects are archetypal achievements):

(9) I name this ship the *Nice Work If You Can Get It*.

States differ from events in that they lack definite bounds, and are inherently non-dynamic. (The latter characterization will be fomally substantiated later.) Some lexical concepts are states, notably those expressibly using the copula, as in a, below. The progressives and perfects considered above, as well as certain predications of habitual action, are also archetypal states, as in b, c, and d:

(10) a. Keats is a genius.
b. Keats is looking into Chapman's Homer.
c. I have lost my watch.
d. I work for the union.

It should be stressed that any claim that an event like *Keats writing* is intrinsically an activity is no more than a convenient shorthand. It is true that *in most contexts* the following sentence is odd.

(11) Keats wrote in fifteen minutes.

However, as Dowty pointed out for a related example, in a discourse context in which the speaker and the hearer both believe that Keats is in the habit of writing a sonnet to time every Sunday, and the speaker knows that on the particular Sunday under discussion—say 23rd December 1816, cf. Garrod 1954:532—Keats took fifteen minutes at it, then the utterance is felicitous. Such examples show that aspectual categories like activity and accomplishment are *ways of viewing* a happening, rather than intrinsic properties of verbs and the associated propositions, or of objective reality and the external world.

The fact that the same form of words can convey more than one aspectual category, provided contextual knowledge supports this view of the passage of events, is the first clue to an explanation for the imperfective paradox. The semantics of the progressive must demand an activity as the only event type that it can map onto the corresponding progressive states. When combined with an accomplishment, as in example 8c, it must first *turn it into* an activity, by decomposing the accomplishment into its components, and discarding the culminating achievement. When combined with an achievement, as in 6c, it must first turn it into an accomplishment, identifying an associated activity from the knowledge base and the context. Then the original achievement can be discarded, Such an account would explain the fact that in normal contexts examples 6c and 8c hold of identical situations.

Events can turn into activities by turning into an *iteration* of the core event.

(12) Chapman sliced the onion (into rings)

Such iterations may themselves iterate (as in *slicing onions*), and in the progressive may be predicated of a time at which one is not performing the core event at all:

(13) I am slicing the onions

Such iterated ativities are investigated by Karlin (1988). A similar transition to a habitual state can occur if, to extend an earlier example, Keats not only writes sonnets to time, but also regularly manages it in fifteen minutes or less. Under these circumstances he can say the following on an occasion on which he is not writing at all:

(14) I am writing a sonnet in fifteen minutes (these days).

There is more to NPs like *the onions* and *a sonnet* in the above examples than may meet the eye. Verkuyl and Dowty also pointed out that some similar protean shifts in aspectual category of the event conveyed by a sentence depended upon the semantic type of the nominal categories involved as arguments of the verb, Thus *Chapman arriving* is an archetypal achievement, which happens to be resistant to combination with a *for*-adverbial, because the state that it gives rise to seems to preclude iteration, as shown by a, below. But *visitors arriving* is necessarily an iteration, as in b.

(15) a. # Chapman arrived all night.
b. Visitors arrived all night.

Such aspectual changes, which include several further varieties that cannot be considered here,[2] may compose indefinitely, especially under the influence of stacked adverbial modifiers, as in

(16) It took me two years to play the "Minute Waltz" in less than sixty seconds for one hour without stopping.

The complexities of this kind of aspectual type-shift or "coercion" are very thoroughly explored by the authors already cited. Accordingly we will pass over further details here, merely offering the chart shown in figure 2.1, by way of an informal summary. The chart divides the aspectual categories into states and events, the latter being subdivided into four sorts, based on two features representing the semantic properties of *telicity*, or association with a particular change of state, and *decomposability*. The latter property is often referred to as "durativity", but it is really to do with decomposition into sub-events, rather than temporal extent. To Vendler's three event categories we follow Miller and Johnson-Laird (1976) in adding a fourth, atomic atelic, category, here called a *point*. (They are what Smith (1991) calls "semelfactives"). These authors suggest that events like *stumbling* and *breathing a sigh of relief* may be basic concepts of this type, but the real significance of the category is to act as a way-station, where the internal structure of an event is "frozen" on the way to undergoing a further transformation, such as being iterated, or acquiring a consequent state by becoming an achievement under the influence of the perfect. The latter is the route by which accomplishments are coerced to achievements yielding the correct consequent state by the perfect, rather than a route which strips the activity out of the accomplishment. To see that this is correct, it is enough to reflect upon the narrower truth/felicity conditions of the second of the following pair, on the assumption that it concerns the achievement that

[2] In particular, a number of coercions from state categories to events, such as the shift to an inchoative seen in *At midnight, I knew that Harry was not going to show up*, are omitted from discussion at this point in the current draft.

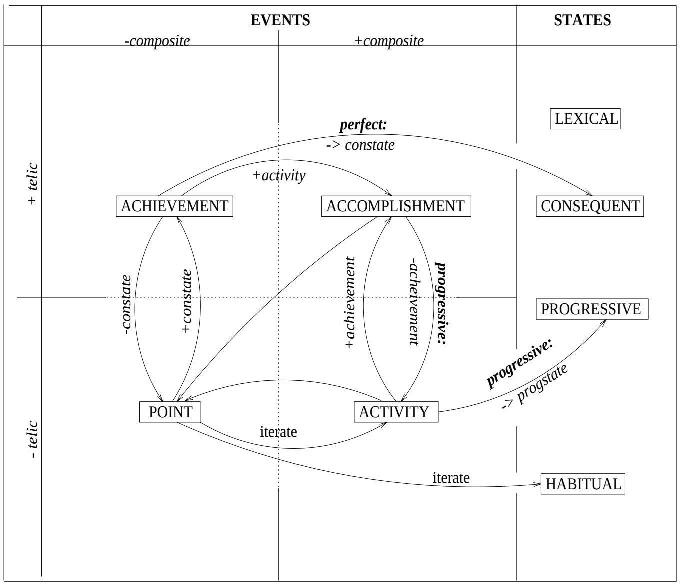

Figure 2.1: A scheme of aspectual coercion (adapted from Moens and Steedman 1988)

concludes the first:

(17) a. I have climbed Mount Everest
b. I have reached the top of Mount Everest.

The arrows in figure 2.1 indicate permissible type-transitions, with annotations indicating the nature of the aspectual change. Some of these, like iteration, are "free", provided that the knowledge base supports the change. Others, like the transition to a consequent state (*constate*) or a progressive state (*progstate*), can only occur under the influence of a particular lexical item or construction, such as the perfect or the progressive. Such restrictions are indicated by bold-face annotations. A more extensive system of coercions and lexically-based restrictions has been developed by Pustejovsky (1991).

Whether free or lexically determined, these type-changes appear to reflect a knowledge representation in which events of all kinds are associated with a *preparation*, or activity that brings the event about, and a *consequent*, or ensuing state, in a tripartite data-structure proposed by Moens and Steedman (1987) that can be viewed as in figure 2.2. This structure, or "nucleus" can be regarded as composed of the types described in figure 2.1. Thus the prepara-

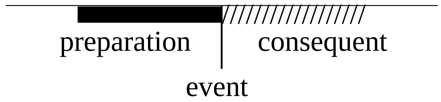

Figure 2.2: The event nucleus (adapted from Moens and Steedman 1988)

tion is an activity, the consequent is the same kind of state that the perfect gives rise to, while the event itself is an achievement. (The nucleus itself is therefore closely related to the category of *accomplishments*). The core achievement may itself be a complex event, such as an accomplishment.

These components should be thought of as "normal" preparations for or consequences of the events—that is, as "defaults" in the sense that the term is used in nonmonotonic knowledge representation (McCarthy 1977; Reiter 1978; Thomason 1997). Each of them may itself be compound. For example, the preparation may be an iteration of some kind, the consequent state may identify a chain of consequences, and the core event may itself be a complex event, such as an accomplishment. The tripartite nucleus has been adopted and used extensively in the DRT theory of *Aktionsarten* of Kamp and Reyle 1993:557-570 *et seq.*, Blackburn, Gardent and de Rijke (1993), and Gagnon and Lapalme (1995), and cf. de Swart (1998) and van Lambalgen and Hamm (2005).

2.2 Logical and Computational Approaches.

So much for the natural history of temporal ontology: how do we formalise this quite complex ontology? Simplifying somewhat, two basic approaches can be distinguished in this voluminous literature.

The first approach is to attempt to define the neo-Vendlerian ontology via quantification over more or less classical Priorian instants, or their dual, intervals. Bennett and Partee (1972), Taylor (1977), Cresswell (1974), Dowty (1979), Heinämäki (1974), Bach (1980), Galton (1984), and the computational work of McDermott (1982), Allen (1984), Crouch and Pulman (1993), and McCarty (1994). are of this kind.

This approach was extremely important in opening up the territory to include temporally extended events, which had largely been ignored in the situation calculus and modal-logic based approaches (see the discussion below). However, the recursive structure of events that follows from the ontology illustrated in figure 2.1, and in particular the problems of granularity and non-continuity in iterated events, mean that some of the definitions of *for*-adverbials and related categories in Dowty's treatment can be criticised, as he himself has pointed out (Dowty 1979, preface to second edition).

The second approach is to take certain types of events themselves as primitive, without any appeal to notions like truth of a predicate over an interval or set of instants. Such events involve a temporal extension, which for connected continuous events is an interval (or equivalently a pair of points), but modifiers like "slowly" are predications of the event rather than the interval that it occupies. This then opens up the further possibility of defining relations between event-sorts in terms of various lattices and sort hierarchies. The algebraic event-based approach was pioneered by Kamp (1979); Kamp and Rohrer (1983), among others, and characterises the work of the present author 1977; 1982, Bach, Brown and Marslen-Wilson (1986), Link (1987), Hinrichs (1985, 1986), ter Meulen (1984, 1986), Dowty (1986), Moens and Steedman (1987), Krifka (1990), Eberle (1990) and White (1993, 1994), and builds upon Carlson (1977), Link (1983) and Landman's 1991 accounts of the ontology of entities. The work of Davidson, as developed in Parsons (1990), and of Jackendoff (1991) as formalised by Zwarts and Verkuyl (1994), can also be seen as belonging to this school.

The latter approach can be seen as a logical continuation of the earlier work, for Dowty (1979) had observed the parallel betweeen the telic/atelic distinction in the event domain, and the count/mass distinction in the entity domain. Not only is the downward-entailing property characteristic of both activities and mass terms: the involvement of mass or count terms as arguments can also determine the event type of a proposition, as in the following minimal pair.

(18) a. Chapman drank beer (for an hour/#in an hour).
b. Chapman drank two pints of beer (#for an hour/in an hour)

The technicalities involved in these different accounts are considerable, and somewhat orthogonal to the present concerns. We will pass over them here, referring the interested reader to Dowty, Krifka, Link and others cited above, and to (White, 1994, ch. 2) for a recent comprehensive review and one of the few extensive computational implementations of a system of this kind.

The way that the progressive coerces an achievement first to an accomplishment and thereby to an activity whose progressive is directly provable as a proposition of the form $in\_progress(activity)$ can be captured in the following rules, in which *accomplishment* is the relation between an activity and its normal or default concluding achievement and p is a variable over predicates.

A standard logic programming convention is used, whereby all variables in the consequent are implicitly universally quantified and all other variables are implicitly existentially quantified. An activity is in progress if the database says it is. The progressive coerces its argument to be an activity, and is true if the activity is in progress.

(19) a. $activity(P) \Rightarrow coerce(P,P,activity)$
b. $activity(Q) \wedge accomplishment(Q,P) \Rightarrow coerce(P,Q,activity)$

(20) $coerce(P,Q,activity) \wedge in\_progress(Q)$
$\Rightarrow in\_progress(P)$

While we will defer discussion of temporal inference until a later chapter, it should be clear that the coercion of the progressive of an achievement to that of the associated normal preparatory activity means that its truth does not hinge on whether the original achievement actually occured or not.

For example, consider the task of proving that in a given situation Keats is completing the sonnet *In Disgust of Vulgar Superstition*:

(21) $in\_progress(complete(keats,in\_disgust))$

We must assume that the knowledge base makes explicit the relation between *completing a sonnet* and a characteristic preparatory activity, as in the nuclear relations of figure 2.2. To simplify the example, we assume that the preparation for completing a sonnet is just writing it:

(22) $\models sonnet(y) \Rightarrow accomplishment(write(x,y),complete(x,y))$

Since *in_disgust* is a sonnet, axiom (20) allows us to set up the subgoal of showing that

(23) $coerce(complete(keats,in\_disgust),Q,activity) \wedge in\_progress(Q)$

which (19b) coerces to:

(24) $in\_progress(write(x,in\_disgust))$

This in turn is true by (19b) true just in case the database contains the following stative fact:

(25) $in\_progress(write(keats,in\_disgust)$

This proof did not involve any subgoal of showing that Keats ever actually completed the sonnet. Indeed the proof would be quite consistent with adding the denial of that fact, capturing the imperfective paradox of Dowty and others. The fact that the activity in question is specified as $write(x,in\_disgust)$ goes some way towards capturing the fact that the mere writing of lines from that sonnet is not sufficient to model the progressive: it has to be *writing those lines with the intention of writing the whole thing*.

Many verbs have meanings that can be defined in terms of similar coercions—cf. Thomason (1997, 1999). For example, *trying* to do something resembles the progressive in coercing an achievement or accomplishment to the corresponding preparatory activity, and its truth similarly does not hinge on the actual attainment of the achievement:[3]

(26) $\models accomplishment(p,q) \wedge intend(q) \wedge do(p)$
$\Rightarrow try(q)$

Failing to do something and *managing* to do it are similar, except that they involve explicit assertion or denial of the attainment of the achievement:

(27) $\models accomplishment(p,q) \wedge fail(q) \wedge intend(q) \wedge do(p)$
$\Rightarrow fail\_to(q)$

Many such lexically-governed coercions are derived from nouns. Thomason 1997:820 points out that identifying the meaning of phrases like "Hammer the metal flat" with that of "causing the metal to become flat by hitting it with a hammer" overgeneralize to models such as those in which the hammering merely signals to a third party that they should put the metal though a hydraulic press. He suggests an analysis paraphrasable as "using a hammer in the normal way for metal to make the metal flat" He points out that "the information about normalcy that is needed in such examples is exactly the sort of information that is needed for practical planning". In what follows, we shall have reason to link this observation to what Gibson (1966, 1979) called the "affordances" of objects—that is, the events made possible by objects such as hammers, such as beating metal, and the consequent states of those events, such as the metal in question being flat.

Many novel denominal verbs depend on coercions involving the "normal" relations between entities and the events that result in such relations, as in the following example (adapted from P.G. Wodehouse):

(28) Keats swiftly trousered the £5 note.

We will return to such examples in formal terms below, but it is worth noting in passing that such affordance-based coercions are extremely specific. Keats'

[3] In fact of course, an utterance about *trying* to do something generally conversationally implicates *non*-attainment. That this is a pragmatic implicature rather than part of the sense-meaning is evident from the fact that the implicature can be explicitly cancelled.

trousers afford his putting £5 notes in the pockets, in order to secure them for himself. Other people's trousers do not have the same affordance for Keats, so (28) cannot mean that he put the money in someone *else's* trouser pocket.

A further indication of the central role that Gibsonian affordances play in the cognitive abilities that underlie natural language comes from North American Indian languages, such as the Athabascan group that includes Navaho. Such languages are comparatively poorly off for nouns. In particular, many nouns for artefacts are morphological derivatives of verbs. For example, "towel" is *bee 'ádít'oodí*, glossed as "one wipes oneself with it", and "towelrack" is *bee 'ádít'oodí bąąh dah náhidiiltsos*—roughly "one wipes oneself with it is repeatedly hung on it" (Young and Morgan 1987).

Such languages thus appear to lexicalize nouns as a default affordance. (Of course, no particular inferences should be drawn concerning Navaho-speakers' abilities to reason about objects. Though productive, these lexicalizations are as conventional as our own. Navaho-speakers probably find the propensity of English to productively allow denominal verbs, like "hammer" and "trouser" equally exotic. We shall return to this question.)

Chapter 3

Temporal Relations

3.1 Basic Phenomena and Descriptive Frameworks

Having established an ontology, or taxonomy of temporal types, we turn to the relational apparatus. The linguistic system that conveys temporal relations between individuals of these different sorts comprises in English the subsystems of tense, (progressive and perfect) aspect (which we have so far only treated in in terms of their effect upon ontological type), and modality.

3.1.1 Tense

The most fundamental of these systems is tense. In the case of tense, as in the case of propositional aspect or *Aktionsart*, there is one early modern piece of insightful descriptive work which most theories build upon, and which those who ignore seem doomed to reconstruct. This work is contained in two short and highly elliptical sections in Reichenbach's *Elements of Symbolic Logic* ((Reichenbach, 1947, Chapter VII, sections 48 and 51)). (Again there are direct precedents in work by Jespersen and Cassirer).

Reichenbach can be read as making two points about temporal expressions. The first is that there is a referential or extensional relation between propositions and facts or events, expressible by the inclusion of events or times as values of bound variables. This observation is the direct antecedent of Davidson's theory (cf. (Davidson, 1967, p.115-116)) and much subsequent work in formal semantics (cf. (Parsons, 1990, p.5)), and is less directly related to the situation calculus of ((McCarthy and Hayes, 1969, cf. p.498-500)) and much subsequent work in artificial intelligence and computer science, discussed below.

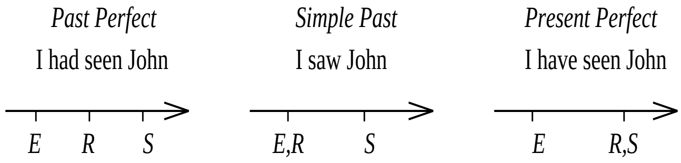

Figure 3.1: Past vs. Perfect (from Reichenbach 1947)

Reichenbach's second point is more specifically linguistic. He argued that the tense system could be understood as a predication not over two times, "now" and "then", but rather over *three* underlying times. These times he called *S*, (speech point), *R*, (reference point), and *E* (event point). *E* can be thought of as the temporal extension of the proposition itself— essentially the Davidsonian **e**, or its modern equivalent, generalised to cope with the kind of ontological questions that concerned us in the last chapter, as for example in work discussed earlier by Parsons (1990) and Schein (1993, 2003). *S* can, as its name suggests, be thought of as the speaker's time of utterance, (although we shall see that it must be generalised to cover embedded times of utterance and narrative point-of-view). Reichenbach's real innovation was the reference point, which can be identified with the notion "the time (or situation, or context) that we are talking about". It is easiest to convey the idea by example. Reichenbach offers the diagrams in Figure 3.1, in which the arrow indicates the flow of time, to show the distinctions between the past perfect, the simple past (or preterit) and the present perfect (all of which he includes under the heading of "tenses of verbs"). The important insight here is that the simple past is used to make a statement about a past time, whereas the perfect is used to make a statement about the present, as was noted earlier in connection with the "present relevance" property of examples like 1.

As Isard and Longuet-Higgins (1973) have pointed out, this claim is consistent with the observation that past tense, unlike the perfect, demands that the past reference point be explicitly established, either by a modifier, such as a *when* clause, or by the preceding discourse. Thus a, below, is inappropriate as the first utterance of a discourse, except to the extent that the reader *accommodates* a temporal referent, in Lewis' 1979 sense of that term— that is, *introduces* an appropriate individual in the database, as one often must at the beginning of a modern novel. But b is appropriate, on the assumption that the hearer can identify the time in the *when* clause:

(1) a. #Chapman breathed a sigh of relief
    b. When Nixon was elected, Chapman breathed a sigh of relief

(In many North American dialects of English, the past tense does double duty for the perfect. I am assuming that this reading is excluded in this case by the most readily accessible aspectual category of *breathing a sigh of relief*.)

The fact that the discourse can establish the "anchor" for the reference point has led a number of authors, including McCawley (1971), Partee (1973, 1984), Isard (1974), Bäuerle (1979), Hinrichs (1985), Webber (1988), Song and Cohen (1988), and others to identify tense, and by implication *R*, as "pronominal" or otherwise anaphoric in character.

We should distinguish this referent-setting function of such adverbials from the aspect-setting function that we encountered in chapter 2, concerning *Aktionsarten*. The adverbials like *in fifteen minutes* and *for fifteen minutes* were there predicated over the *event* point *E*. In the cases to hand, they are predicated over *R*. Many of the adverbials that relate two propositions temporally, particularly *when* clauses, do so by identifying or predicating a relation over the reference points of the two clauses, via what Reichenbach called the "positional use of the reference point". The following are all cases of this kind.

(2) a. In ten minutes, I looked at my watch.
    b. When Chapman arrived, the band was playing *Nice Work If You Can Get It*.
    c. After Einstein arrived in Princeton, he may have visited Philadelphia.

We return to the anaphoric role of tense in chapter 4.

With the benefit of the discussion in earlier chapters, we can go a little further than Reichenbach, and say that the predication which the perfect makes about the reference point, present or past, is that the consequent state that is contingent upon the propositional referent *E* holds at the reference point *R*.

Reichenbach extended his account of tense and the perfect to the progressives and futurates, including habituals, and to sequence of tenses in compound sentences. Some of the details of his original presentation are unclear or incorrect. For example, the exact relation of *E* and *R* in the past progressive is unclear, possibly because of a typographical error. The account of the futurates does not correctly separate the respective contributions of tense and modality. The account of sequence of tense in compound sentences omits any discussion of examples with subordinate complements requiring more than one *S* and/or *R*, such as *He will think that he has won*. He similarly seems to have failed to notice that there is a *second* "narrative" pluperfect, involving an embedded

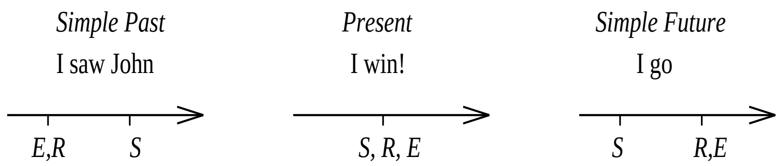

Figure 3.2: The tenses

past tense, relative to a *past* speech point, distinct from the true past perfect. It is the only realisation that English affords for the past tense of indirect speech, or *oratio obliqua*, exemplified in examples like the following:

(3) I had arrived in Vermilion Sands three months earlier. A retired pilot, I was painfully coming to terms with a broken leg and the prospect of never flying again. ...

This pluperfect cannot be the past tense of a perfect, as perfects like *#I have arrived in Vermilion Sands three months ago* are infelicitous (for reasons discussed by Moens and the present author 1988). It is rather a past tense *of a past tense*, identifying the proposition *I arrived in Vermilion Sands three months before now* as uttered by a narrator with their *own* now. (Most of Reichenbach's own examples of the pluperfect are in fact of this other kind).

For these and other reasons the following account is something of a reconstruction of Reichenbach's theory. (See Hornstein (1977, 1990), Enç (1981, 1987), Kamp and Rohrer (1983), Caenepeel (1989), (Smith, 1991, Ch. 5), Declerck (1991), Spejewski (1994), Mittwoch Mittwoch (1995), Palmer et al. (1993), Crouch and Pulman (1993), Abusch (1997b), and Hitzeman (1997) for related proposals. See also the discussions of Lascarides and Asher (1993a), and Kamp and Reyle (1993), below).

According to this view, English and presumably other languages can be seen as having three tenses in the narrow sense of the term—the familiar past, present, and future tenses, in all of which the Reference point R and the event point E coincide. The past tense is, as we have seen, one in which the pair R,E precedes S. The present tense (which we noted earlier is in English restricted as far as events go to performative acts like naming and promising) is one in which all three coincide. The true future tense in English (as opposed to other languages) is realised by the syntactic present tense, as in *I go to London (next Tuesday)* and is symmetric to the past tense, with the pair R,E later than S, as in Figure 3.2 (cf. Hornstein (1977, 1990)). Here I depart from Reichenbach himself, and Bennett and Partee (1972), who regarded the future as not merely

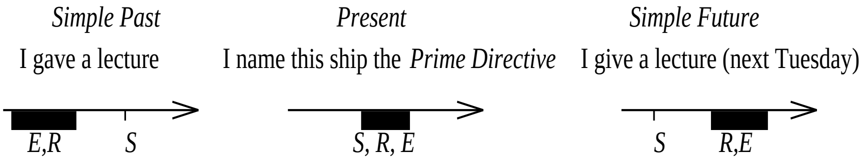

Figure 3.3: The tenses for composite events

the mirror image of the past tense, but as combining the characteristics of a tense and a futurate aspect, mirroring the perfect. (Smith, 1991, p. 246) also regards what is here called the simple future as having a *present* reference point. Nevertheless, the claim that it is a pure tense, with *R* co-temporal with *E*, is supported by the observation that the futurate is anaphoric, like the past, with exactly the same need for an "anchored" reference point.[1] Hence a, below, is inappropriate when discourse-initial, whereas the anchored b is fine. (Cf. 1. A related analysis is implicit in Abusch (1997a).)

(4) a. #Harry moves to Philadelphia.
    b. Next Tuesday, Harry moves to Philadelphia.

The modal future, *I shall go* should be understood as identical to the simple future as far as Reichenbach's underlying times are concerned, and as having a future reference point with the modal itself contributing meaning of a quite orthogonal kind, which we shall discuss in a separate section below.

The ontology of events discussed in chapter 2 should be viewed as an ontology of the Reichenbachian *E*, so that the past and (simple or modal) future tenses can be applied to durative or composite events, as in figure 3.3. (On the assumption that the performative achievement performed in saying "I name this ship the *Prime Directive*" lasts at least as long as the utterance, the present too can be regarded as having an extended *R*). With the simple tenses, as opposed to the sentential aspects considered below, the reference point *R* continues to be *coextensive* with *E* for durative or composite events.

The reference point *R* itself is nevertheless distinct from *E*, and not a part of this ontology. Davidsonians accordingly distinguish it from the Davidsonian **e** ((Parsons, 1990, p. 209) uses **I** for essentially this purpose in discussing tenses and temporal adverbials.)

We noted earlier that past tense has a second meaning in English that is

[1] In fact, for similar reasons, Reichenbach himself regarded the future tense as ambiguous between these two readings—see Hitzeman (1993, 1997) for further discussion of these and other alternatives.

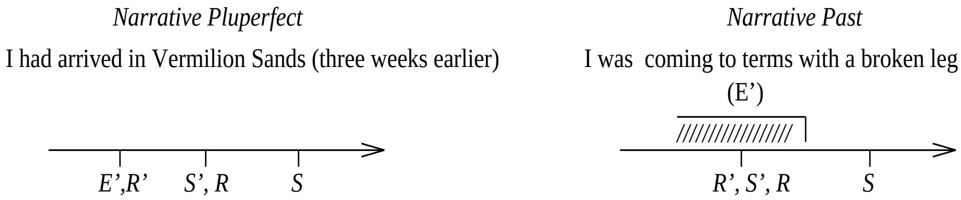

Figure 3.4: The narrative tenses

predicated of propositions in which the speaker's reference point *R coincides with* an epistemic point of view S' that is not the same as the speaker's present S, in the novellistic device of *oratio obliqua*. The syntactic past and pluperfect in the earlier example 3 are therefore represented by the diagram in figure 3.4. This analysis is related to one proposed by Kamp and Rohrer (1983) (cf. (Kamp and Reyle, 1993, p.593); Nelken and Francez (1995)), and by Hwang and Schubert (1992), all of whom postulate multiple reference points to cope with related observations. The present account differs only in preserving Reichenbach's insight that for each reference point R there is an S.

The existence of these narrative or quotational tenses in English may explain the phenomenon of "sequence of tense", in which complements of tensed verbs like *said* and *thought* tend to "inherit" the tense of the matrix verb. As (Hornstein, 1990, ch. 4) points out, this phenomenon is naturally captured in a Reichenbachian framework by similarly assuming that each tensed clause has its own S,R,E triple. The embedded S', which is naturally thought of as an embedded utterance point, or (more generally) an embedded epistemic point of view, is then coincident with the matrix event E, the event of utterance or epistemic consciousness. However, in the grammar of English, embedded clauses are specified to be semantically like quoted present tensed utterances, with past tense denoting the structures in figure 3.4. S and R in these relational structures then coincide with S and R in the matrix clause. Thus a and b, below, mean that Chapman said something like "I arrived in Vermillion Sands three months ago", and "I am painfully coming to terms with a broken leg", just like the narrator in the following examples:

(5) a. Chapman said that he had arrived in Vermillion Sands three months earlier.
b. Chapman said that he was painfully coming to terms with a broken leg.

The fact that English complement verbs specify only quotational complements

Past Counterfactual

If I had won (I would now be rich)

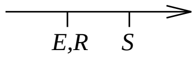

Figure 3.5: The counterfactual pluperfect

is what makes English a relatively strict “sequence of tense (SOT) language”. However, this is a syntactic convention, rather than a semantic necessity, and other languages (such as ancient and modern Greek) may allow (or insist upon) the basic tenses in these contexts.

One further remark about quotational and complement pluperfects is in order. They are in fact ambiguous in English, since besides the narrative pluperfect illustrated in figure 3.4a, they may denote the narrative past of a perfect, obtained by replacing the progressive state in 5b by a perfect, or consequent, state, as in the following variant:

(6) Chapman said that he had just broken his leg.

Such an account of sequence of tense phenomena is essentially equivalent to the accounts of Enç (1981) and Dowty (1982), who invoke related notions of “anchoring”.

We shall return later to the fact that past tense is also used in English to mark *counterfactuality* of the core proposition with respect to the reference point, as in the following conditional sentence.

(7) If he were taller, he could reach the book himself.

(Some languages have a distinct subjunctive mood for this purpose. English retains a distinct subjunctive only for the copular verb *be*). When the reference point itself is past, this means that counterfactuals also surface as pluperfects. We shall have more to say about the counterfactual relation of *E* to *R* below. However, as far as purely temporal relations go, their temporal profile is the same as the past tense, as in figure 3.5. Because of this multiplicity of functions of English past tense, Isard (1974) and Lyons (1977) suggest that syntactic past tense should be identified with a non-temporal semantic primitive *REMOTE*, rather than a time as such.

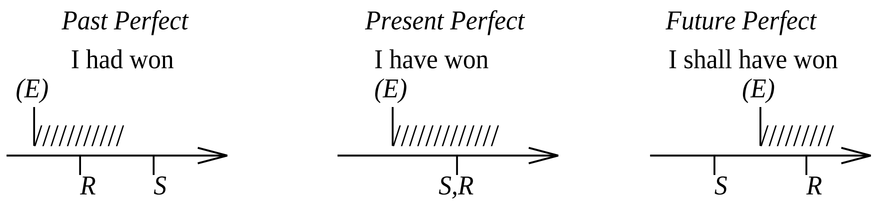

Figure 3.6: The perfect

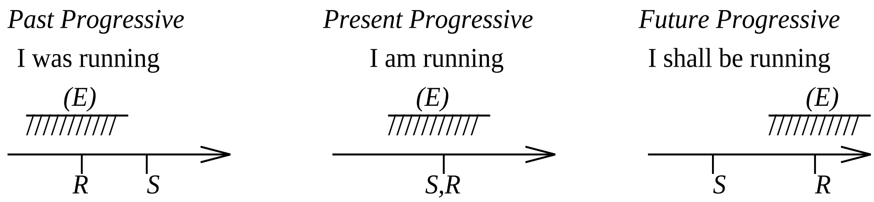

Figure 3.7: The progressive

3.1.2 The Perfect and the Progressives

With the tenses established as in Figure 3.2, we can see that the perfect and the progressive (both of which we saw earlier to be states, rather than events) compose correctly with tense, as in Figures 3.6 and 3.7. In the case of the former, the reference point *R* lies within a Consequent State, derived from the original event *E*, which must in the terminology of chapter 2 be an achievement. In the case of the progressives, *R* lies within a Progressive State, derived from the original event *E*, which must in the terminology of chapter 2 be an activity. In neither case does *E* in the sense of the event directly figure in the representation. It is the (progressive or consequent) state derived from the event *E*, here indicated by hashing, that is predicated of *R*. Unlike the tenses with *E*, *R* is not coextensive in temporal terms with such states, but temporally included within them. The position of the event *E* relative to *S* and *R* is not in fact fully determined by the perfect and the progressive—hence its appearance in brackets in the figures. This becomes important in the case of the future perfect, in which the relation of *E* to *S* may be either prior or posterior. (Here we depart slightly from standard Reichenbachian accounts such as Hornstein (1990).)

Both in the tenses and the aspects the core event *E* may be derived from a different event category E', via type coercion. For example, the achievement of *winning the race* can turn into a corresponding accomplishment, by the knowledge-based association of a characteristic preparatory activity, such as

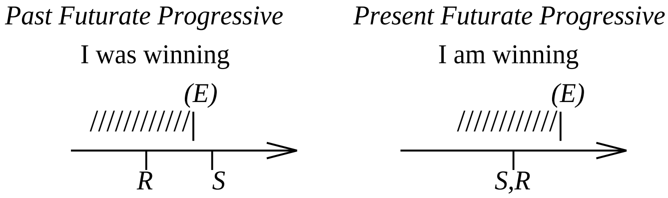

Figure 3.8: The futurate progressives

running. The progressive can then strip off the original achievement, to leave the bare activity, which is then mapped onto the corresponding state, which is predicated of *R*, the time under discussion. This explains the possibility of "futurate" progressives like a, below:

(8) a. I am winning!
    b. I was winning.

As Smith 1991:247 reminds us, (8a) is not really a predication about a win as such. It is simply a present progressive of an activity normally culminating in winning, which in Reichenbachian terms looks like figure 3.8. Since *E*, the original achievement of winning, is not predicated of any underlying time, we seem to be even closer to a resolution of the imperfective paradox, which applies to both present (including futurate), and past, progressives. However, to get to that point we must consider the third temporal-relational system, that of modality.

3.1.3 Epistemic and Deontic Modality

The modal verbs of English, such as *will, must*, and *may*, like those of many other languages, carry two distinct senses. The first concerns such notions as necessity, possibility, inferability, or predictability of the core proposition, and is usually referred to as "epistemic" modality. The following are some examples for which this is the only readily accessible interpretation:

(9) a. It must have died.
    b. That will be the mailman.
    c. She may be weary.

The other set of senses concerns notions like feasibility and permissibility of the core proposition, and the desires, abilities, and obligations of the agent, and is usually refered to as "deontic" modality. Some relatively unambiguous examples are the following:

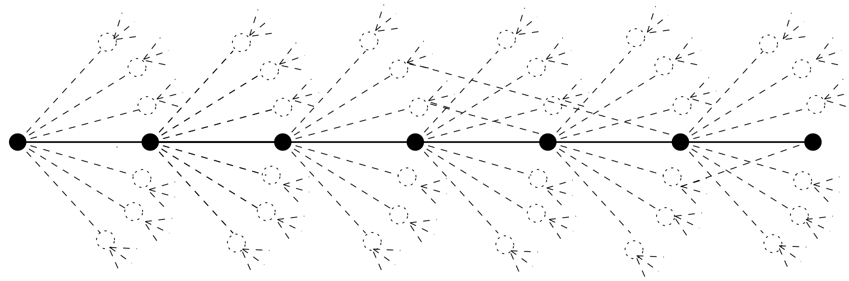

Figure 3.9: Modal Temporal Structure

(10) a. You must sit down.
b. You may smoke.
c. I can do the boogaloo.

While the pairs of senses subsumed under a verb like *must* are clearly related, the relation is indirect and appears to be somewhat arbitrarily conventionalised. While many of the deontic modals can be viewed as creating or explaining the corresponding epistemic state, there are a number of complications and lacunæ in the system as a whole. For present purposes we shall consider the deontic modals as essentially distinct from the epistemic modals.

Because of their involvement with necessity and possibility, the epistemic modals differ from the systems of tense and sentential aspect in requiring us to consider more than one domain of reference or classical model. It was possible to capture the semantics of *Aktionsart*, tense, and sentential aspect in terms of a single deterministic world history, represented informally as a time-line in Reichenbach's diagrams. Instead we must think of the flow of time as a graph, so that any particular history (such as that of the real physical universe) becomes a path of branching points in a Kripke model or discrete graph of situations or partially specified worlds (equivalently, sets of possible worlds), each of which gives rise to alternative continuations, which themselves branch into alternatives. Any given state may be reachable by more than one route. Such models are used in all logical and computational accounts of temporality and modality. Such a graph can be pictured as in figure 3.9. We use bold dots and lines to indicate the states and state transitions of actual time.[2]

[2] This informal notation is related to Morgenstern's 1994 *Real-occurs*, and to related constructs in Peleg 1987 and Lehman and Shelah 1983. Morgenstern presents a version of McDermott's 1982 calculus in which an event *Real-occurs* if it *Occurs* and is *Real*—that is, on the sequence or chronicle of actual events. A related idea is implicit in much earlier work by Isard (1974).

It should be noted that this representation does not distinguish the future history from the past in this respect. This reflects the fact that the simple future tense, which in English we have seen is realised as the present, treats the future as determinate. Of course, in actual fact, our access to past history is different in kind to our access to the future. There is a privileged set of past states which are distinguished as the actual history of the world, and we can only make more or less well-informed guesses about which states will turn out to be actual in future. We shall return to the consequences of this observation in the later section on modality.

We shall see below that this structure is closely related to the modal logician's notion of a Kripke structure, defined in terms of an *accessibility relation* over possible worlds (although the logicians frequently regard such "worlds" as including entire histories—that is, of comprising many states). It will be important to ask then how states should be represented, and what defines this relation. (For the present purpose, as in other computational applications of modal logic (cf. Goldblatt 1992), the accessibility relation is the central construct in a modal logic). However it is important first to see that the modal verbs, seen as predications over the elements in such structures, are straightforwardly compatible with the Reichenbachian view of tense and aspect.

First, we must be clear that such structures are different from the continuous temporal dimension that is implicit in the earlier figures. We must now think of time as a (partial) ordering on discrete states corresponding to instants at which changes to a model occur (or can occur).

We could in principle think of such states as densely packed, mapping to the real numbers. When modal logics of the kind discussed below have been used to model physical time according to the special theory of relativity, they have represented time in this way—cf. van Benthem (1983), and Goldblatt (1980). (The latter achieves the *tour de force* of axiomatising the Minkowski chronsynclastic infundibulum as a modal logic—see van Benthem 1995 for discussion). However, for linguistic and computational purposes, we shall invariably be interested in much sparser temporal structures. Sometimes (particularly when thinking about the theory of digital computation) states in these structures correspond to the cycles of a clock—that is, to the integers rather than the reals. In linguistics and related AI tasks like planning, we may be concerned with even sparser representations, in which only partially-ordered *changes* of state are represented.

We will continue to defer the discussion of how this is to be done formally. We may note however that in the latter case, transitions between points in the structure 3.9 are naturally associated with *events* that precipitate those changes. For example, this would be a natural way of representing the history of a board game of simple moves like *W:P-K4*, as Isard (1974) does. Doing so in effect transforms the standard modal logician's Kripke structure with a single accessibility relation over possible worlds into a structure in which the arcs are labeled according to the events which change one state into another. We shall see later how to generalize this representation to durative or composite events.

The Reichenbachian underlying times S and R can provisionally be identified with points in this structure, which will now support modality in the following way (again we go beyond anything specifically claimed by Reichenbach here).

We saw earlier that the possibility of present epistemic modal statements like 11a, below, is most naturally captured by assuming that the models or databases representing nodes or worlds in the structure specify *partial* information.[3]

This observation can be extended to the domain of temporal relations when we observe that modals and conditionals are essentially predications about entire Reichenbachian tensed propositions.[4] In a, below, the modal is predicated about a present proposition, where $S = R$. Example b is predicated of an past tensed proposition.

(11) a. She may be weary.
b. Einstein may have visited Philadelphia.

(Being the complement of an auxiliary, and hence infinitival, this past shows up as a perfect. However it is clear that the proposition concerns a past reference point, rather than a present perfect, because the corresponding declarative perfect, below, is pragmatically anomalous, for reasons discussed in chapter 2.)

(12) #Einstein has visited Philadelphia

[3] This proposal is subtly different from the usual use of partial information in model-theoretic semantics—cf. Kripke (1965), Turner (1981), Veltman (1984), Landman (1986), and related notions in DRT and situation semantics—see Kamp and Reyle (1993), Barwise and Perry (1983), and Cooper (1986).) (Hintikka 1962, van Fraassen 1971, Kratzer 1977, and Cresswell 1985, ch.5, 1988 offer other notions of partial information.) The present use of partial information to represent the temporal/modal dimension itself follows work in AI like that of Isard (1974), Moore (1980, 1985), Morgenstern (1988, 1994), Scherl and Levesque (1993) and others.

[4] The following material departs from the analysis in the earlier paper, which I now regard as wrong. I am grateful to Matthew Stone for help on this point.

Such "modal pasts" do in general require the reference point to be previously established or accommodatable. This can be seen in the fact that they are compatible with temporal adverbials like *yesterday*, which present perfects in general are not, as we have noted:

(13) a. She must have visited Philadelphia yesterday.
b. #She has visited Philadelphia yesterday.

One way to capture the above facts is in terms of partial models or databases of the kind discussed above. Modal pasts like *may have* assert that the speech point is is *accessible from* some such past reference point. Modals like *must* assert that the speech situation is *only* accessible from such points, and implicates the past proposition under some argument or line of reasoning, as proposed by Kratzer (1991) and Stone (1994); Stone and Hardt (1997).

I have already argued that in connection with the non-modal future tense that English and other languages treat the future part of the structure in figure 3.9 symmetrically with the past, as having a determined set of states constituting actual future history. Of course, our information about future actuality is in fact limited, and our knowledge merely probabilistic. Because of this practical fact of human existence, the most common kinds of statement about the future are modal, so it is not too surprising that the modal system is in English somewhat confounded with the future component of the tense system. Nevertheless, in sentences like the following, we should clearly distinguish the contribution of the modal in a, below, from the fact that it is predicated of a future reference point, as was pointed out by Boyd and Thorne (1969) in connection with examples like b, below, where the same modal is predicated of a present reference point:

(14) a. You will marry a tall dark stranger.
b. It's late. Your mother will be worrying about you.

This point was also a subject of lively debate among 19th century linguists, as Verkuyl (1989) has shown. The ambiguity between present and futurate readings holds for other modals:

(15) a. You may marry a tall dark stranger.
b. It's late. Your mother may be worrying about you.

To fully capture the meaning of epistemic modal predications over events, as in the following example, we must generalise the above apparatus.

(16) (If you take my queen), you may win the game.

This implies that the reference point must include or give access to the entire accessible subgraph of possible futures after the core event. This observation in turn suggests that the reference point is more like the nucleus of figure 2.2, with the consequent state capturing the notion of accessible subgraph, than it is like a situation or a time. We shall return to this point below.

3.1.4 Counterfactuals

The system of linguistic modality is closely related to that of counterfactuality in conditionals, which in English is marked by past tense, and which will turn out to be central to the resolution of the imperfective paradox.

In order to capture the meaning of counterfactual implication, and hence causation, Lewis suggested that the meaning of counterfactuals in sentences like the following depends on the notion of *similarity* between possible worlds in a modal structure like figure 3.9

(17) If you had taken my queen when you took my rook, you would have won the game.

The proposal was that *P (your taking my queen)* in the situation *W* when you took my rook counterfactually implies *Q (your winning the game)* (written $P \Box\!\!\rightarrow Q$) if among all the worlds accessible from *W* satisfying *P*, all the ones that are most similar to *W* also satisfy *Q*. (It should be observed here that "worlds" are entire world histories, not the transitional states of the situation calculus).

This works well for the example to hand, because the only counterfactual world is the one that results from your taking the queen instead of making the move you actually made. By definition it is the most similar counterfactual world, so provided all continuations of the game from that world result in your winning the game, the claim is true.

However, as Fine (1975) pointed out, not all actions are like this. His example was the following:

(18) If Nixon had pushed the button, there would have been a nuclear war.

This statement might well be true, despite the fact that worlds in which that least monotonic of presidents pressed the button, but war did not ensue, seem to be more similar to the actual world, on the reasonable assumption that nuclear war changes just about everything. Thomason and Gupta (1980) point out that Lewis' account is compatible with an alternative notion of closeness over worlds, defined in terms of causality, a suggestion that we shall return to below, in discussing the situation calculus of McCarthy and Hayes, and its extensions.

The problems of modality and counterfactuality are closely related to the imperfective paradox, which will be recalled as arising from the existence of occasions of which it can be claimed that *Keats was crossing the road*, in spite of the fact that he was hit by a truck before the action could be completed. The problem for possible worlds semantics is precisely the same as the problem of counterfactuals, namely to specify the worlds which are most similar to the actual one, differing only in relevant respects. To specify this in terms of worlds themselves is very difficult: as Vlach (1981) pointed out, there are a great many world-histories that differ in minor respects from the actual one, but where Keats is still hit by the truck. As Landman (1992) has pointed out, there are cases of world-histories which differ from the actual world *only* in that Keats is not hit by the truck, but in which Keats would nevertheless not have succeeded in crossing the road—as when there is a second equally inattentive truck right behind. Even if there were an infinite number of such trucks, requiring an infinitely different world for Keats to succeed, it still seems true that *Keats was crossing the street* if that is what he intended, and if our knowledge of the world supports no other obstacle to the causal realisation of that intention. Even more strikingly, (to adapt another of Landman's examples), it seems *not* to be true in any of these situations to make this claim if Keats did *not* have that intention, or if there is some other obstacle to its realisation. If he knew perfectly well that he could not possibly get to the other side, and set out with suicidal intentions, or if he intended to turn around just short of the opposite kerb and come back again, or if he fully intended to cross but was unaware of a glass wall in the middle of the road, then the claim is false. Yet, apart from the intention itself, and its consequences for Keats' projected future actions, the counterfactual worlds are all identical.

Because of these difficulties, most theories of the progressive have invoked a function mapping possible states onto *relevant* continuations. Dowty (1979, p.148) "reluctantly" assumed a primitive function *Inr*, mapping world-time indices onto "inertia worlds". Landman (1992) defines a function *C* which maps an event *e* and a world index onto their "continuation branch", invoking a primitive function *R* which maps such pairs onto event-indexed inertia worlds or "reasonable options". Some related ideas have been invoked within the DRT camp (cf. Roberts (1989)).

However, both *Inr* and *R* are unanalysed, and the involvement of intention makes it seem unlikely that there could be any definition other than one in terms of an action-based accessibility relation.

3.2 Logical and Computational Approaches

So much for the natural history of temporal relations: how do we formalise them? We should at this point distinguish two kinds of question that are somewhat confounded in the computatioinal literature. One is the use of abstract computations to do the same job as a traditional model theoretic semantics. The other is the efficient implementation of such a semantics, to minimise costs such as search. In this section, we shall first develop a Kripke-like semantics including a representation of states and the accessibility relation. We shall then consider an efficient representation of this semantics, which builds in certain "inertial" properties of the world as it is conceptualised by human beings, via a constrained use of defaults. Finally we shall consider a reformulation of this system in terms of dynamic logic.

We noted a resemblance between the structures like Figure 3.9 and the notion of a *frame* in the standard semantics of Kripke (1972) for a modal logic. A frame is a structure defined in terms of *a set of worlds or states W* and an *accessibility relation* ρ over them. For the present purpose, the worlds or states can be thought of as classical models of the kind used in first-order predicate calculus (that is, sets of individuals and relations, possibly typed or sorted), except that we shall assume that states which happen to have the same individuals and relations may nevertheless be distinct. One can then define $\Box p$ ("necessarily p"), to mean that p necessarily holds in a state $s \in W$, just in case p holds in every state accessible from s. Similarly, $\Diamond p$, ("possibly p"), can be defined to hold in s if p holds in at least one state accessible from s under ρ. In most modal logics, these operators are duals, interdefinable via negation.

Possible worlds are generally assumed by modal logicians to include entire histories of the universe of discourse through many different states. However, this assumption is based on a view of time that is not the one pursued here, and for present purposes it is more useful to identify the elements under the accessibility relation with single states, as the computer scientists tend to. The accessibility relation can be any relation whatsoever, but for present purposes it is appropriate to think of it as defining the ways in which one state of the world can lawfully turn into others.

In taking advantage of this affinity between the linguistic phenomena and modal logic, we must be careful to avoid being distracted by two related concerns that have greatly occupied modal logicians. One is an interest in distinguishing between necessary propositions, such as theorems of arithmetic, and contingent ones, such as the fact that this sentence happens to have been written at 5.25 p.m on an October evening. This notion is naturally captured in a logic in which the accessibility relation is *reflexive, transitive*, and *symmetric*—that is, an equivalence relation under which all worlds in W are accessible to all others. (This is the modal logic known as *S5*.) However, this distinction may not be particularly relevant to everyday reasoning, which typically concerns an uncertain world. It does not appear to be reflected in the linguistic ontology.

The second is the representation of physicists' notions of time and causality. The mere fact that quantum theory discusses processes which reverse the arrow of time and causality does not entail that a theory of the knowledge involved in linguistic semantics should do the same. The logics we shall consider have an accessibility relation which is *asymmetric*, reflecting the directionality of the flow of time and causality. (They are therefore somewhat more like the modal logic known as *S4*, although in fact their accessibility relation will turn out to be more restricted still.)

3.2.1 The Situation/Event Calculus

While modal logics offer an elegantly terse notation for quantifying over states or models, the efficient use of most such logics as planners via theorem-proving is limited by the fact that the proof space that must be searched bears no clear relation to the space of possible plans (Stone 1998; Bacchus and Kabanza 2000:125). The latter authors, for example, overcome this problem by making plan operators and plan search strictly extra-logical, using a Temporal Logic $\mathcal{LT}$ only to represent domain knowledge used to guide the planning search.

An alternative is to try to construct the logic itself to have a proof search-space isomorphic to the space of plans implicit in the S4 Kripke model of Figure 3.9, much as the search-space of logic-programming languages like Prolog is isomorphic to the proofs themselves.[5] This is the alternative originally proposed by McCarthy and Hayes (1969) under the name of the "Situation Calculus". The Situation Calculus and its descendants take the states of the S4 model as individuals that can be quantified over, represented either by constants t_i, t_j etc, or by more complex terms, using a technique of "reifica-

[5] Of course, such a move does not solve the planning problem itself. Finding a plan to make a stack of twenty distinguished blocks in a prescribed vertical order still yields a search space that will overwhelm any planner or theorem prover that is not equipped with domain-specific heuristics of the kind that Bacchus and Kabanza (2000) discuss, such as that it is a good idea to start from the bottom. The point is that these such knowledge can be represented in the *same* logic as the action representation.

tion".[6]

The Situation Calculus of McCarthy and Hayes (1969) was developed within a computational framework for reasoning about actions, and is interesting from the point of view of our earlier assumption that the temporal linguistic categories need to be based in a theory of action rather than of time. One of the most useful and attractive features of the situation calculus was the use of *terms* like $result(arrive(person), s)$ as individuals denoting situations or states as functions of other situations. Functions like *result* were called *situational fluents* by McCarthy and Hayes. Such terms can be used in rules like the following to transparently capture the notion that a person is present in the situation that results from their arriving:

(19) $\forall s, \forall person, present(person, result(arrive(person), s))$

This particular logic (which is, as McCarthy and Hayes point out, quite closely related to von Wright's 1964; 1967 "logic of action") embodies only the most minimal ontology of states (represented by predicates that hold over situations, such as $present(person, s)$) and atomic actions (represented by expressions like $arrive(person)$). We shall look in a moment at some descendants of the situation calculus which attempt to include richer ontologies.

McCarthy and Hayes were interested in the use of such rules to construct plans of action, via inference. For example, given the following rules, one might expect to be able to infer a successful plan for bringing about a situation s in which three blocks satisfy the condition $on(a, b, s) \wedge on(b, c, s)$:

(20) a. $clear(a, s_0) \wedge clear(b, s_0) \wedge clear(c, s_0)$

b. $\forall x, \forall y, \forall s, clear(x, s) \wedge clear(y, s) \wedge x \neq y$
$\Rightarrow clear(x, result(puton(x, y), s))$
$\wedge \neg clear(y, result(puton(x, y), s))$
$\wedge on(x, y, result(puton(x, y), s))$

The formulæ say, first, that everything is *clear* in a particular situation s_0, and second, that if two distinct things x and y are clear in a situation s, then in the situation that results from putting x on y in that situation, x is on y, x is clear and y is no longer clear. (The rule embodies the idea that only one thing at a time can be manipulated, in stipulating that y is no longer clear).

[6] Reification using constants requires constraints or predications over those constants to be derived in theorem proving, which may in turn require the use of equational theories to determine when times are equal. Moore (1980, 1985), Jackson and Reichgelt (1987) and Frisch and Scherl (1991) take the former approach, while Smullyan (1973), Wallen (1990), Ohlbach (1991), and Stone (1998) take the latter.

Using standard inference rules of conjunction elimination, modus ponens, etc., we might expect to be able to prove the following, in which the situational terms neatly describe the sequence of putting *b* on *c*, then putting *a* on *b*:

(21) $on(a,b,result(puton(a,b),result(puton(b,c),s_0)))$
$\wedge\, on(b,c,result(puton(a,b),result(puton(b,c),s_0)))$

As yet, this doesn't quite work. While we can prove the intermediate result $on(b,c,result(puton(b,c),s_0))$ (which looks useful) we cannot go on to prove the first conjunct, because the formulæ in 20 do not capture the fact that *a* remains clear after putting *b* on *c*. Nor can we prove the second conjunct, because the same formulæ fail to capture the fact that *b* remains on *c* after putting *a* on *b*.

McCarthy and Hayes point out that we can fix this by adding further "frame axioms" to the effect that if *u* is on *v* in a situation *s*, then *u* is still on *v* in the situation that results from putting something *x* on something *y*, so long as *u* is not the same as *x*. Similarly, if *u* is clear in *s*, it is still clear after putting something *x* on something *y*, so long as *u* is not the same as *y*:

(22) a. $\forall u, \forall x, \forall y, \forall s, clear(u,s) \wedge u \neq y \Rightarrow clear(u, result(puton(x,y),s))$
b. $\forall u, \forall v, \forall x, \forall y, \forall s, on(u,v,s) \wedge u \neq x \Rightarrow on(u,v,result(puton(x,y),s))$

The addition of these rules allows the proof (which is suggested as an exercise) to proceed to completion.

Such a system, whose affinities to von Wright's logic of action we have already remarked upon, seems to offer a very natural expression for states and the accessibility relation between them. However, as McCarthy and Hayes were aware, for computational purposes, this logic seems cumbersome. If we want to represent a less trivial universe with more state predicates and more actions or action sequences, we shall need a frame axiom pairing every predicate with every action.

McCarthy and Hayes christened this the "Frame Problem". It is not always appreciated that there are two aspects to this problem, which Shanahan (1997) distingushes as the "representational" and "computational" or inferential versions.

The representational problem is that frame axioms like (22) somehow miss the point as a representation of action. The way we think of actions is precisely as local operations that affect just a few properties, leaving most facts unchanged. For example, *my eating a hamburger* is an event whose effects are confined to the hamburger and myself, and leave countless other facts about the situation, such as the color of the walls, unaffected. There seems to be something wrong with a notation that would make it no more inconvenient to define a highly distributed event which inverted the truth value of every fact about the world. Even the action of dropping a hydrogen bomb doesn't do *that*. Such redundancy exacerbates the search problem for the computational purposes that originally motivated the situation calculus.

The computational or inferential version of the frame problem says that even if we achieve a terser knowledge representation that minimizes the number of tedious frame axioms (say by quantifying over propositions and "circumscription" of affected fluents, of a kind discussed below), we may still have to do computational search to find out what the colour of the walls is after five events of my eating a hamburger.

McCarthy and Hayes discuss a number of possible solutions, including one which they attribute to Rescher (1964), which was to assume that *all* facts that held at the start of an action held in its result, and then to eliminate any inconsistencies via what would now be recognised as a "Truth Maintenance" system (Doyle 1979; deKleer (1984)). However, they did not in this early paper offer a definitive solution. The search for a solution has engendered much research, not least their own (see for example Hayes (1971); McCarthy (1977)).

A solution to both the representational and computational frame problems that was related in spirit to Rescher's was nevertheless at hand in work that was being done contemporaneously in robot planning. The idea was to build into the model itself the "inertial" property just identified. The simplest way to do this is to specify actions in terms of the facts about the world that become untrue and the new facts that become true when they occur. One computationally convenient way to do this is to represent the starting state of the world as a collection of facts, and to represent actions in terms of a triplet. Each such triplet consists of 1) a list of *preconditions* that must hold if the action is to apply in the current state, 2) *deletions* or facts that become false in the state that results from the action, 3) *additions* or facts that become true in that state.

The history of an episode up to any given state can then be determined from the current state and the sequence of actions that led to it. Any earlier state can be fully determined by running the sequence of additions and deletions backwards to the relevant point.

It is also natural for this purpose to further economise by representing the state of the world solely in terms of *positive* truths, and to represent the (generally much larger) set of negative facts via a "closed world assumption" (Reiter (1978)), according to which any fact that cannot be proved true is assumed by default to be false. (It should be noted that this move demands that everything true be provable, if consistency is to be maintained.)[7]

It is not clear who first proposed this idea, because its transparent representation in terms of "assignment", database "updates" and other computational side-effects makes it almost the first thing a computer scientist would think of as a representation for action. It usually goes by the name of the "STRIPS solution" to the frame problem, because it was first made explicit in the context of a robot action planner by that name (Fikes and Nilsson 1971).

For example, a STRIPS operator *eat* can be defined as follows for a simplified world in which hamburgers are the only edible things and the agent of the eating is suppressed:

(23)	PRECONDITIONS:	$hamburger(x)$
		$here(x)$
		$hungry$
	DELETIONS:	$here(x)$
		$hungry$
	ADDITIONS:	$thirsty$

Such representations were initially derided by logicians because of their procedural expression and nonmonotonic character.

3.2.2 A Declarative Solution to the Frame Problem

Although the STRIPS representation of actions was originally thought of in non-declarative terms, (Kowalski 1979, circulated in 1974) showed that the representational aspect of the STRIPS solution could be elegantly realized in entirely declarative terms, via the introduction of the closed world assumption and a more radical use of reification to simulate modal quantification. (See Nilsson 1980, p. 308-316 for a more extensive discussion of Kowalski's proposal). He proposed a predicate *holds*, which applies to a proposition, represented as a term, and a state. The earlier state 20a can therefore be written as follows:

(24) $holds(clear(a), s_0) \wedge holds(clear(b), s_0) \wedge holds(clear(c), s_0)$

The action of putting x on y can be represented as a STRIPS rule, as follows. The preconditions are defined by the following rule which says that if you can get at x and you can get at y, the preconditions for putting x on y hold:

[7] There is also a potential conflict with the earlier suggestion that modal possibility should be represented as consistency with a partial model. See Gelfond and Lifschitz (1993) on closed world assumption in the situation calculus.

(25) $holds(clear(x),s) \wedge holds(clear(y),s) \wedge x \neq y$
$\Rightarrow preconditions(puton(x,y),s)$

(In this rule, and henceforth, we again adopt the logic programming convention whereby variables in the result are implicitly universally quantified over and all other variables are implicitly existentially quantified over.) The new facts that result from the action of putting x on y can be defined as follows:

(26) a. $holds(on(x,z),s) \Rightarrow holds(clear(z), result(puton(x,y),s))$
b. $holds(on(x,y), result(puton(x,y),s))$

Kowalski assumes negation as failure, and so avoids the need to state explicitly that y is no longer clear. This fact is implicit in the following frame axiom, which is the *only* frame axiom we need for the action of putting x on y. It says that any fact which holds in s holds in the result of putting x on y in s *except* the fact that y is clear, and the fact that x was on something else z (if it was).

(27) $holds(p,s) \wedge p \neq clear(y) \wedge p \neq on(x,z)$
$\Rightarrow holds(p, result(puton(x,y),s))$

The use of inequality here involves an implicit quotation of terms like $clear(y)$ and $on(x,z)$, once the variables have been bound. There is also an assumption implicit in the use of inequality (rather than a related notion involving implication) that p is a positive literal rather than a formula like $graspable(y) \wedge on(x,y)$. This assumption is in effect a restriction to Horn logic, in which the consequent may not include conjunction, disjunction, implication or negation, modulo the reification.

Kowalski's proposal was followed by much work on tense using reified calculi (Allen (1984); McDermott (1982); Kowalski and Sergot (1986); Kowalski and Sadri (1994); Robert and Sadri (1997)). It was also closely related to the notion of "circumscription of qualifications"—see McCarthy (1977), and much other subsequent work, collected and reviewed in Ginsberg (1987). In particular, Reiter (1991) shows how the restricted frame axioms or "successor state axioms" can be derived automatically. We can now define a predicate *poss*, closely related to the familiar modal operator $\Diamond$, over the set of possible states, via the following rules, which say that the start state s_0 is possible, and the result of an action in a state is possible if its preconditions hold:

(28) a. $poss(s_0)$
b. $poss(s) \wedge preconditions(action,s) \Rightarrow poss(result(action,s))$

These axioms in effect define an accessibility relation *poss* in terms of the elementary actions of the blocks world.

The earlier goal of stacking *a* on *b* on *c* can now be realised as the goal of finding a constructive proof for the following conjunction

(29) $poss(s) \wedge holds(on(a,b),s) \wedge holds(on(b,c),s)$

These rules can be very straightforwardly realised in Prolog, and can be made to yield a proof (although the search problem of finding such proofs automatically remains hard in general) in which

(30) $s = result(puton(a,b), result(puton(b,c), s_0))$

This technique restores declarativity to the logic embodying the STRIPS solution.

However, it important to notice that, while solving the representational component of the frame problem, it leaves the computational problem unsolved. If we want to know the color of the walls in the situation that results from a sequence of *n* instances of *puton* actions, we still have to do *n* potentially costly steps of inference involving search.

There is a sense in which—despite the involvement of the closed world assumption—Kowalski's technique also restores monotonicity, for so long as we do not add new facts (like some previously unsuspected object being present, or a familiar one having fallen off its support) or some new rule or frame axiom (say defining a new action or stating a new precondition on an old one) then we can regard negation-as-failure as merely efficiently encoding classical negation.

Of course, in the real world we *do* learn new facts and rules, and we encounter exceptions to the closed world assumption of complete knowledge. These problems are known in AI as the *ramification problem* (that is, that actions may interact with other actions in complicated and unforeseen ways with consequences that our default model does not predict) and the *qualification problem* (that actions may have indefinitely many preconditions that our default model does not anticipate). In many recent papers, the frame problem is assumed to include these further problems.

One standard example of the ramification problem is the following. The frame axiom (31) of a form commonly assumed in situation calculus accounts assures me that I shall be the same colour after moving as I was before.

(31) $holds(p,s) \wedge p \neq at(x, place_1) \Rightarrow holds(p, result(move(x, place_1, place2), s))$

However, if the path of my movement takes me through a paint spray my color may change, engendering inconsistency.

The answer to this problem is simply to fix the representation. The logic must represent the fact that moving takes time, and that other things can happen during your move, in a way that is modular and independent of the specification of other events in the knowledge domain.

Similarly, if our logic has difficulty with the fact that if a cup is on a saucer, and the saucer is moved, then the cup moves with it, whereas if the cup is moved, then the saucer stays behind, then we have not done a good job of representing knowledge about cups and saucers (Hayes 1971). Our logics should similarly be able to represent such predictable ramifications as that if a car is in Edinburgh and a pizza is in the car and the car moves to Auchterarder, then the pizza will be in Auchterarder, but also that if the pizza is thrown from the car while the journey is in progress, the pizza will *not* be in Auchterarder. (It will be convenient to refer to this example, which is isomorphic to other versions of the "Temporal Projection Problem" in the literature, as "the Auchterarder Ramifying Pizza Delivery Scenario"—cf. Sandewall (1994).) We will see in the next section that the situation calculus can be readily adapted to this purpose, using the temporal ontology discussed in the last chapter. The Horn clause form and closed world assumption of the declarative situation calculus also support logic programming in languages like Prolog, and offer some hope of computational efficiency (cf. Levesque (1988)). (This is not of course to claim that in itself it solves the explosive search problem implicit in *finding* plans.)

The following is the standard example of the qualification problem, due to McCarthy. Suppose I am in my car and I want to drive. I know that for me to drive the engine must be running, and that for it to be running I must switch it on and that the precondition for switching on is to have the key in the lock and the way to do that is to put it there, for which the only precondition is my being in my car.

Most of the time this works fine, but the space-time continuum is treacherous. Actually there are several, possibly infinitely many, further preconditions. For example, it is well-known (?) that if there is a potato on the tailpipe of a car, switching on will not result in the engine running. By assumption, this practical precondition is not explicit in the knowledge representation that I am using. What do I do?

As far as inference using this logic goes, there is simply nothing that I can do. All I can do is seek another knowledge base. I might start listing hypotheses, and doing experiments, and reasoning from first principles (slow). Or I might read the manual (but it may have the same problem with potatoes and tailpipes). Or (most likely) I may just wander around helplessly kicking the tyres until I notice something out of the ordinary (What's that potato doing there?). This last tactic is in the spirit of "reactive" planning (Agre and Chapman (1987).

However, if we are in possession of an efficient default model which works reasonably well most of the time, it may well be wiser to regard the problem of coping with new information as residing outside the logic itself, in the truth-maintenance or "housekeeping" system. Rather than coping with genuinely unforseeable ramification and qualification in the logic itself, we should think in terms of a system of truth-maintaining transitions between entirely monotonic logics.

Related techniques and their relation to ramification and qualification in a narrower sense of *known* or anticipated causal or contingent relations between events are further explored by by Schubert (1990, 1994), and Reiter (1991, 1993, 2001) and below.

3.2.3 The Event Calculus

The STRIPS version of the situation calculus, with one class of situational fluent and one class of state-predicate, did not embody any of the ontological richness discussed in earlier chapters. However, a number of systems subsequently generalised the situation calculus to deal with richer event ontologies. Allen (1984) was the first to do so, defining a number of reifying predicates of which the most basic were *HOLDS* and *OCCURS*, respectively relating properties and events to intervals. Events could be events proper, or processes, after the Vendler-like scheme of Mourelatos (1978), those events that in other systems are points or instants being represented as very short intervals, unlike instants in the related extension of the situation calculus proposed by McDermott (1982), which introduced further temporal types such as "chronicles". Galton (1990) proposes an elegant revision of Allen's theory in these respects, according to which Allen's processes correspond to progressive states in the terms of chapter 2. A number of causal and temporal predicates allowed events and times to be related in ways that permitted a treatment of phenomena such as inaction, propositional attitudes and interacting subgoals in plans. The present author (1982) defines the durative or composite event categories in terms of instants of starting, stopping and (in the case of accomplishments) culminating, using a STRIPS-like representation to handle the related progressive states, which (like perfects, habituals etc.) are treated as inertial properties of intervening times. (A similar approach has been advocated by Lin (1995) and Lin and Reiter (1995), and in ter Meulen (1995), in which progressive states are represented by "stickers," an attractive Post-It-Note-like metaphor for fluents or properties of states). All of these approaches are closely related to the "event calculus" of Kowalski and Sergot (1986), and Shanahan (1997,b), itself a descendant of Kowalski's earlier work on the situation calculus, although their own ontology of events was even more minimal than these other approaches.

If possible states are defined as databases of facts exploiting a closed world assumption, then the above definition of the accessibility relation in terms of actions is essentially identical to the branching modal frame identified in figure 3.9 in the last section. We noted there that in order to capture linguistic modality we seemed to need *partial*, or underspecified, states.

The accessibility relation in question is (in a sense) *transitive*, but it is *asymmetric* (and therefore *irreflexive*)—that is, a partial order. This relation defines an even weaker logic than *S4*, which has a transitive and *antisymmetric* accessibility relation.

We have already noted the similarity to the system of von Wright (1967). The states in this structure are *counterfactual* in the sense that the actions or action sequences that take the place of situational fluents, generating the successors of any given state, are disjunctive, and only one of them can correspond to the actual history of events. In this respect, our system has some affinities to proposals by Stalnaker (1968, 1984), Thomason (1970), and Lewis (1971, 1973), and the related computational work of Ginsberg (1986). However, the definition of accessibility and counterfactuality in terms of events rather than states avoids the problem with some of these accounts that was noted earlier in connection with Fine's example 18, repeated here:

(32) If Nixon had pushed the button, there would have been a nuclear war.

According to the event-based system, there is exactly one counterfactual world in which Nixon pressed the button, rather than doing whatever else he did, and its accessibility is defined by the action itself. (Cf. (Stalnaker, 1984, ch. 7, esp. p.133-134)).

Such a system similarly resolves the imperfective paradox. We noted earlier that *Inr* and *R* were unanalysed in the systems of Dowty and Landman. However, in the present system they can be identified with the event-based accessibility relation proposed here. Since that appears to be the *only* accessibility relation that we need, the event-based account appears to have an advantage. In fact, to make the identification of these inertial functions with the accessibility relation itself seems to be a very natural move within all of these theories, particularly in view of the close relation in other respects that Dowty notes between his theory and the logic of action. (Cf. (Dowty, 1979, p.144)). (Thomason, 1991, p.555) also notes the close relation between inertia worlds and the situation calculus, and suggests a rather different analysis in terms of defaults in a non-monotonic logic.

This is not of course to claim that the situation calculi described above solve the problem of representing causality, agency, and the like, (although here too von Wright (1967) made a start). Shoham (1988),Morgenstern and Stein (1988) and (Stein, 1991, p. 117-8), in contrast to Ginsberg (1986) and Ortiz (1994, 1999b,a), eschew counterfactuals *and* situation-calculus like systems in favour of some more general notion of the accessibility relation based on *causality* (usually on several distinct causal operators, including as *enabling, generating* and *preventing* as well as simple causation). Schank (1975), Wilensky (1983), and Lansky (1986) are also important in this connection.

3.2.4 The Calculus of Affordance

There is another way of looking at all of these variants of the situation/event calculus. To the extent that the accessibility relation is defined in terms of a number of different events or causal primitives, possibly a large number, it is possible to regard each of these as defining its own distinct accessibility relation, possibly differing from others in properties like transitivity. Such systems can then be viewed as instances of the "dynamic" logics that were developed in the first place for reasoning about computer programs—see Pratt (1979), Harel (1984), Goldblatt (1992) and Muskens, van Benthem and Visser (1997). The application of various forms of dynamic logic in knowledge representation and natural language semantics has been advocated by Moore (1980), Rosenschein (1981), Webber (1983), Pednault (1989), and Scherl and Levesque (1993). (It should be noted that this original notion of dynamic logic is not the same as the "dynamic predicate logic" (DPL) of Groenendijk and Stokhof, which is briefly disussed below).

Dynamic logics relativise the modal operators to individual actions, events, or programs. For example, if a (possibly nondeterministic) program or command α computes a function F over the integers, then we may write the following:

(33) $n \geq 0 \Rightarrow [\alpha](y = F(n))$

(34) $n \geq 0 \Rightarrow \langle\alpha\rangle(y = F(n))$

The intended meaning of the first of these is "for $n \geq 0$, after every execution of α that terminates, $y = F(n)$". That of the second is (dually) that "for $n \geq 0$, there is an execution of α that terminates with $y = F(n)$".

In what follows, we shall be exclusively concerned with the deterministic modality $[\alpha]$. While the world is a surprising place, in which doors do not always open when pushed, and cars do not always start when the ignition key is turned, the claim will be that the knowledge representation which we use to plan actions, and which underlies the semantics of temporality in natural language, is deterministic. When the world turns out to be inconsistent with such representations, plan repair or replanning is required. However, this is generally treated in the AI literature as a completely different kind of extra-logical "truth maintenance" problem (see Bratman, Israel and Pollack 1988; Horty and Pollack 2001. Even when we come to consider probabilistic action representations, we shall use the $[\alpha]$ modality, associating probabilities with the states of the model that result from actions, rather than the actions themselves.

While all of the calculi that we have considered so far are ones in which the elementary programs α are *deterministic*, dynamic logics offer a framework which readily generalises to concurrent and probabilistic events, offering a notation in which all of the theories discussed here can be compared. (In some of these, the modal operators $[\alpha]$ and $\langle\alpha\rangle$ are no longer interdefinable—cf. Nerode and Wijesekera (1990)).

The particular dynamic logic that we are dealing with here is one that includes the following dynamic axiom (the operator ; is *sequence*, an operation related to composition, and to von Wright's T):

(35) $[\alpha][\beta]P \Rightarrow [\alpha;\beta]P$

In this we follow (Moore, 1980, ch. 3) and Rosenschein (1981). The situation calculus and its many variants can be seen as reified versions of this dynamic logic.

We achieve an immediate gain in perspicuity by replacing the reified event calculus notation in a, below, by the equivalent dynamic expression b.

(36) a. $holds((on(a,b) \wedge on(b,c)), result(puton(a,b), result(puton(b,c), s_0)))$
b. $[puton(b,c); puton(a,b)](on(a,b) \wedge on(b,c))$

Kowalski's "vivid" version of STRIPS can be very simply represented in this logic.[8] The initial state of the world is that it contains three blocks, none of

[8] Bibel (1986) proposes a related representation.

which is *on* another:

(37) $block(a) \wedge block(b) \wedge block(c) \wedge on(a, table) \wedge on(b, table) \wedge on(c, table)$

The axiom defining the preconditions of $puton(x,y)$ is now directly definable in terms of a predicate closely related to Kowalski's *possible*, which we will call *affords*. The use of this Gibsonian term reflects the fact that it applies to *events*, rather than consequent states:

(38) $\models \; block(x) \;\wedge\; block(y) \;\wedge\; \neg on(z,x) \;\wedge\; \neg on(w,y) \;\wedge\; (x \;\neq\; y) \;\Rightarrow$ $affords(puton(x,y))$

To define the update consequences of putting something on something else we need a different kind of rule, using Girards's 1995 notion of linear implication, written $\multimap$, as follows:

(39) $\models \{affords(puton(x,y))\} \wedge on(x,z) \multimap [puton(x,y)]on(x,y)$

Linear implication, $\multimap$, treats positive ground literals or "facts" in the antecedent as consumable resources, removing them from database and replacing them by the consequent. The "exponential" "!" marks $affords(puton(x,y))$ as a nonconsumable precondition: the truth of this condition after a *puton* event is not defined by the linear implication, and is a matter for further inference, via rules like (38). Girard 1995:23 provides a model theory for linear implication, in which facts are elements of the model and $\multimap$ is defined in terms of an operation on them (Girard, Blass (1992), and Japaridze (1997, 1998, 2002) also offer a related Game-Theoretic semantics).

The resultant calculus, which it is convenient to refer to as the Linear-Dynamic Event Calculus (LDEC), is therefore a multi-implicational logic, in which the to connectives $\Rightarrow$ and $\multimap$ are "fibred", or linked together via the *affords* predicate. This calculus has the effect of eliminating frame axioms like (27) entirely: rule (39) already captures the fact that the *puton* action is "abnormal" with respect to $on(x,z)$ and $on(x,y)$.

The transitive part of Kowalski's *poss* axiom (28) is now reduced to the following:

(40) $\models affords(\alpha) \wedge [\alpha]affords(\beta) \Rightarrow affords(\alpha;\beta)$

This fragment preserves the virtues of Kowalski's treatment in a modal notation. That is, the following conjunctive goal can, given a search control, be made to deliver a constructive proof where $\alpha = puton(b,c);puton(a,b)$. (The proof is suggested as an exercise):

(41) $affords(\alpha) \wedge [\alpha](on(a,b) \wedge on(b,c))$

The suppression of state variables in dynamic logic affords some improvement in perspicuity over the related proposals of Kowalski, McCarthy, Schubert, and Reiter that it is here used to capture, and makes it easier to extend the calculus. While I have so far confined the description of states to simple first-order predicate calculus formulae, it should be noted that any other logical description could be used instead, including standard modal representations of knowledge and belief, typed logics, and/or DRSs, for example.

LDEC thus solves both the representational and computational versions of the frame problem. However, it is important to notice once again that the search problem implicit in even simple blocks world problems of this kind remains hard. The space of possible action sequences that must be searched to finding a plan to stack *n* blocks in a specific configuration is potentially infinite, and even when search is limited in depth is exponential, and in practice overwhelms most planners for quite small *n*. Domain-specific knowledge (such as that it is a good idea to start building towers from the bottom) is essential.

The dynamic axioms of LDEC can be viewed as a representation of Miller et al's **TOTE units**, or of the Behaviorists' notion of an **operant**. The "Test-Operate/Test-Exit" loop of TOTE units is necessary for the execution of the plan in the world, and can also be represented in the logic (see below). They are also reminiscent of Hoare (1969) triples, the nodes of Petri (1962) nets, and the Finite-State Transducers of Fernando (2004, 2005). What the logic adds is a formal way to plan with such dynamic units.

Some animals can make plans of this kind involving tools (Köhler 1925—see Figure 3.10).

This planning behavior can be modeled using LDEC, as follows (as usual, the axioms are simplified for exposition).

Grabbing something gets you to the state of having it, and if you were 6 ft higher than where you are you could grab the bananas (a numerical hack is used to avoid axiomatizing arithmetic):

(42) a. $\{affords(grab(bananas))\} \multimap [grab(x)]have(x)$
b. $at((here+3)+3) \Rightarrow affords(grab(bananas))$

Boxes afford climbing on them:

(43) $box(b) \Rightarrow affords(climb\text{-}on(b))$

—and if you are at a place and you climb on a box you are at a place that is higher by 3ft:

(44) $\{affords(climb\text{-}on(b))\} \wedge at(p) \multimap [climb\text{-}on(b)]at(p+3)$

Figure 3.10: From Köhler 1925

As before, if two boxes have nothing on top of them and are not the same box you can put one on the other (38), and if *x* is on something and you put it on something else then that something becomes clear and *x* is on that something else (39).

If the initial state of the world is as follows:

(45) $at(here) \wedge box(b1) \wedge box(b2) \wedge clear(b1) \wedge clear(b2)$

—then the goal (46a) gives rise to (46b) as one possible plan

(46) a. $affords(\alpha) \wedge [\alpha]have(bananas)$
b. $\alpha = [puton(b1, here); climb\text{-}on(b1);$
$puton(b2, b1); climb\text{-}on(b2); grab(bananas)]$

We have said nothing yet about the problem of *Search* implicit in identifying such plans. Such planning in animals seems to be *reactive* to the presence of the tool and *forward-chaining*, rather than backward-chaining (working from goal to tool). That is, the animal can make a plan in the presence of the tool, but has difficulty with plans that require subgoals of finding tools.

This is consistent with our Gibsonian affordance-based assumption that actions are accessed via perception of the objects that mediate them. It seems a good way for an animal to plan. If there *is* a short plan using available resources, breadth-first forward chaining will find it. Backward chaining, on the other hand, will usually engender much larger search spaces, since there are usually more possible plans than practically applicable plans, and is only really effective when you have evolved to the point of having very general devices which afford practically anything, such as credit cards and mobile phones. (The assumption of forward chaining is also explicit in the "production system"-based planners of Laird, Newell and Rosenbloom (1987) and Veloso et al. (1995).)

We can define the affordances of objects directly in terms of LDEC preconditions like (43) and (43). Thus the affordances of boxes in the above example are climbing on and putting on:

(47) $affordances(box) = \left\{ \begin{array}{c} climb\text{-}on \\ put\text{-}on \end{array} \right\}$

LDEC also offers a representation for the knowledge underlying the semantics of the causative constructions and denominal verbs discussed at the end of Chapter 2. For example, hammering metal flat can be represented by rules like the following.

A hammer affords hitting metal with it:

(48) $hammer(h) \wedge metal(m) \Rightarrow affords(hit(h,m))$

If something is bent and you hit it with a hammer it, it stops being bent and becomes flat:

(49) $\{affords(hit(h,m))\} \wedge bent(m) \multimap [hit(h,m)]flat(m)$

Similar axioms define the affordances of hammers with respect to nails. Thus, a hammer and a nail afford hitting the nail with the hammer:

(50) $hammer(h) \wedge nail(n) \Rightarrow affords(hit(h,n))$

If a nail is proud and you hit it with a hammer, it stops being proud and becomes flush

(51) $\{affords(hit(h,n))\} \wedge proud(n) \multimap [hit(h,n)]flush(n)$

These axioms represent classic examples of Miller, Galanter and Pribram's TOTE units To make the resemblance clear, we need to decompose this perfective *hammer* into a durative event, much like moving, as in the next section.

Similarly, the denominal action of trousering a £5 note identified in connection with example (28) affords securing the money. Thus we can write:

(52) $trousers(t) \wedge £5(n) \wedge \neg in(n,t) \Rightarrow affords(put\_in(n,t))$

(53) $\{affords(put\_in(n,t))\} \multimap [put\_in(n,t)]in(n,t)$

(54) $yours(t) \wedge in(n,t) \Rightarrow safe(n)$

3.2.5 Durativity in the Dynamic Event Calculus

The above examples only involve non-composite or "non-durative" events, like the original situation calculus. However, we also need to capture durative events, including the kinds of ramification exemplified in an earlier section as the Auchterarder Pizza Delivery Scenario. The following dynamic Horn clauses begin to capture the composite events involved, along the lines suggested in the previous chapter following the present author 1982, Moens and Steedman (1987) and White (1994). (As usual, the example is greatly simplified, and omits many rules needed to capture even this small domain completely).

Continuing to avoid tedious axiomatization of space-time trajectories, we will assume that location, like clock-time, is autonomously updated by the world when you are in motion. This means that propositions of the form $at(x,l)$ will only appear in preconditions or qualifications, not in updates or ramifications,

If you are somewhere, and you aren't already moving, and you intend to be somewhere else, and you know the way, then the situation affords starting to move from there to there:

(55) $at(x,l_0) \wedge \neg in\_progress(move(x,l_1,l_2)) \wedge intend(at(x,l))$
$\wedge\, path(l_0,l) \Rightarrow affords(start(move(x,l_0,l)))$

If you are somewhere, and you start to move from there to somewhere else, you stop being there and are in motion:

(56) $\{affords(start(move(x,l_0,l)))\} \multimap [start(move(x,l_0,l))]in\_progress(move(x,l_0,l))$

(Note that the activity term $move(x,l_0,l)$ includes the goal of the movement.)

If you are moving from somewhere to somewhere else, you can always stop:

(57) $in\_progress(move(x,l_0,l)) \Rightarrow affords(stop(move(x,l_0,l)))$

—and if you do so, then nothing much changes. You are wherever you are:

(58) $\{affords(stop(move(x,l_0,l)))\} \multimap [stop(move(x,l_0,l))]$

More interestingly, stopping moving to a place *when you are at that place* affords arriving:

(59) $in\_progress(move(x,l_0,l)) \wedge at(x,l) \Rightarrow affords(arrive(x,l))$

If you arrive at l, then you no longer intend to be at l, and are at l.

(60) $in\_progress(move(x,l_0,l)) \wedge intend(at(x,l))$
$\multimap [arrive(x,l)]at(x,l)$

Because a fact $at(x,l)$ is (redundantly) added by this rule, it can be used as a goal for planning purposes. This aspect of the knowledge representation has the desirable property of also not generating plans to get home by going to Toronto and stopping when you get to your house.

This representation of durative events is well behaved with respect to standard examples of the ramification problem such as the one given earlier concerning moving through a paint-spray.

Suppose the situation is $at(car,here) \wedge color(car,green)$: Axioms for events of spraying someone some color are as follows. Spraying can start anytime

(61) $affords(start(spray(y,c)))$

Spraying affects your color:

(62) $\{affords(start(spray(y,c)))\} \wedge color(x)$
$\multimap [start(spray(y,c))]in\_progress(spray(y,c))$

Spraying affords stopping spraying:

(63) $in\_progress(spray(y,c)) \Rightarrow affords(stop(spray(y,c)))$

After spraying you are the color of the spray:

(64) $\{affords(stop(spray(y,c))\} \wedge in\_progress(spray(y,c)))$
$\multimap [stop(spray(y,c))]color(y,c)$

It follows that for the situation in which $at(car,here) \wedge color(car,green)$, we correctly prove all of the following the following without encountering inconsistency:

(65) $[start(move(car,here,there)); start(spray(car,pink));$
$stop(spray(car,pink)); stop(move(car,here,there))]color(car,pink)$

(66) $[start(spray(car,pink)); start(move(car,here,there));$
$stop(move(car,here,there)); stop(spray(car,pink))]at(car,there)$

(67) $[start(spray(car,pink)); start(move(car,here,there));$
$stop(spray(car,pink)); stop(move(car,here,there))]color(car,pink)$

(68) $[start(move(car,here,there)); start(spray(car,pink));$
$stop(move(car,here,there)); stop(spray(car,pink))]at(car,there)$

This example makes it clear that this variety of ramification problem arises as an artefact of event representations that represent moving as an *atomic* STRIPS event or interval of the kind implicit in the misleading frame axiom (31). The independence of this observation from the identification of such atomic events with instants or intervals suggests that Allen's emphasis, following Taylor (1977), Dowty (1979) and Richards (1982), on the advantages of intervals as primitives over instants is not necessarily correct (although the duality between instants and intervals means that there is nothing that you can do by taking instants as primitives that you cannot do with intervals and suitable entailments). See Tichy (1985) and Galton (1990) for related critiques of interval semantics.

Another frequently-raised objection to STRIPS representations is that they insist that all intervals are sequential. However, this is an advantage, rather than a problem. If $start(spray(car,pink))$ and $start(move(car,here,there))$ could be simultaneously accessed, then we would have a contradiction in the relevant *at* and *color* entailments. But that is simply what we mean by the accessibility relation: those would be different branches of the Kripke model.

The thing to remember is—*that the accessibility relation has nothing directly to do with clock-time.* There is absolutely nothing to stop us indexing two "successive" moves with the *same* instantaneous time. For example, the following still correctly fails:

(69) $[_0 start_1(move(car,here,there)); start_1(spray(car,pink));$
$stop_2(spray(car,pink)); stop_2(move(car,here,there))]_2 color(car,green)$

(70) $[_0 start_1(spray(car,pink)); start_1(move(car,here,there));$
$stop_2(move(car,here,there)); stop_2(spray(car,pink))]_2 at(car,here)$

3.3 Coercion and the Progressive

In order to capture the semantics of the progressive in this representation, its coercive effects on aspectual and temporal types were defined as follows in Chapter 2. First, the coercion of achievements to preparatory activities was captured in the following logic program.

(71) a. $activity(P) \Rightarrow coerce(P,P,activity)$
b. $activity(Q) \wedge accomplishment(Q,P) \Rightarrow coerce(P,Q,activity)$

As before, an activity is in progress if the database says it is. The progressive coerces its argument to be an activity, and is true if the activity is in progress.

(72) $in\_progress(Q) \wedge coerce(P,Q,activity) \Rightarrow in\_progress(P)$

The Imperfective Paradox can then be captured as follows.

First, moving to someplace is the preparatory activity for arriving at that place:

(73) $accomplishment(move(x,l_0,l),arrive(x,l))$

(Recall that for an event to count as $move(x,l_0,l)$ a number of telic conditions must hold) Second, and simplifying somewhat, if a car is moving from Edinburgh to Auchterarder, the car is arriving at Auchterader:

(74) $in\_progress(move(car,Edinburgh,Auchterarder))$
$\models in\_progress(arrive(car,Auchterarder))$

However the truth of this progressive, like that of the car travelling to A., doesn't depend on the realization of the achievement itself. Certain telic conditions on being able to utter it truthfully follow from the knowledge representation.

3.3.1 Cause and Contingency in Dynamic Event Calculus

The following query asks for a plan α yielding a state where Keats is completing the sonnet *In Disgust of Vulgar Superstition*:

(75) $affords(\alpha) \wedge [\alpha] in\_progress(complete(keats,in\_disgust))$

Writing something is the preparatory activity for completing it:

(76) $accomplishment(write(x,y),complete(x,y))$

(Note that this is writing *something*, not just writing). If (skating over some ancient epistemological conundrums) you intend something to be a sonnet, then you can start writing it:

(77) $\models intend(x,sonnet(y)) \Rightarrow affords(start(write(x,y)))$

If we assume that Keats already intends *In Disgust of Vulgar Superstition* to be a sonnet, then the accessibility relation implicit in definition 40 gives rise to a proof where

(78) $\alpha = start(write(keats,in\_disgust))$

As before, the proof that generates the above plan does not involve the actuality of Keats' completing the sonnet, as expressed in a plan of the form $\alpha = start(write(keats,in\_disgust)); complete(keats,in\_disgust)$. Indeed the proof is quite consistent with denying that actuality, because the actuality of the achievement P in rule 72 is not involved in the antecedent, capturing the imperfective paradox.

The above example constructing a plan to bring about a situation in which Keats is finishing writing *In Disgust of Vulgar Superstition* is slightly artificial, because such states are extensive, and there may be several such plans. For example, consider the effect of adding the following rule defining the consequences and preconditions of other people arriving.

(79) a. $\models [arrive(x,l)]at(x,l)$
b. $\models not(present(x)) \Rightarrow affords(arrive(x))$

The accessibility relation 40 now allows

(80) $\alpha = start(write(keats, in\_disgust))$
$\alpha = start(write(keats, in\_disgust)); arrive(chapman)$
$\alpha = start(write(keats, in\_disgust)); arrive(chapman); arrive(homer)$
...

As plans, these are rather foolish, because all except the first one include redundant events. However there is no guarantee that we will find the simplest one first, although of course incorrect solutions such as the following are still correctly excluded for the goal in question:

(81) $\alpha = start(write(keats, in\_disgust)); stop(write(keats, in\_disgust)$

Part of the problem is that we are not yet distinguishing true consequences, including ramifications or causal relations among events themselves, from facts that are merely coincidentally true in the state that results, because of the inertial property of the frame axiom. Nor are we distinguishing causal relations *between* event sequences from mere temporal sequence.

We can remedy this shortcoming by distinguishing the *affords* accessibility relation from a causal or contingent accessibility relation *forces*.[9] Accordingly, we need to add some further rules parallel to 40, reflecting a relation of necessity or expectation across sequences of events, including the following:[10]

[9] Moens and the present author 1988 implicitly associate such a relation with a causal or contingent sequential operator parallel to ;, written @, because of its relation to one of Lansky's 1986 operators. (Related proposals to involve various causal primitives are made by Morgenstern and Stein (1988), Geffner (1990), Stein (1991), Lin (1995), and McCain and Turner (1997)).

[10] Steedman 1995; 1997 and certain earlier versions of the present paper give this rule with the @ sequence operator of footnote 9, as:
(i) (i) $\models (affords(\alpha) \wedge [\alpha]forces(\beta)) \Rightarrow affords(\alpha @ \beta)$
However, this complication is unnecessary under the present version of the theory.

(82) a. $\models (\mathit{affords}(\alpha) \wedge [\alpha]\mathit{forces}(\beta)) \Rightarrow \mathit{affords}(\alpha;\beta)$
b. $\models (\mathit{forces}(\alpha) \wedge [\alpha]\mathit{affords}(\beta)) \Rightarrow \mathit{affords}(\alpha;\beta)$
c. $\models (\mathit{forces}(\alpha) \wedge [\alpha]\mathit{forces}(\beta)) \Rightarrow \mathit{affords}(\alpha;\beta)$

(We will assume that formulae of the form $\mathit{forces}(\alpha)$ are like $\mathit{affords}(\alpha)$ in not being among the atomic formulae transmitted by frame axioms. We will see below that the more extended causality involved in delayed action can be modeled using consequent states and explicit mention of them in frame axioms.)

We can now add a rule saying that anyone else being present forces Keats to stop writing:

(83) $\models \mathit{present}(x) \wedge (x \neq \mathit{keats}) \wedge \mathit{in\_progress}(\mathit{write}(\mathit{keats},y))$
$\Rightarrow \mathit{forces}(\mathit{stop}(\mathit{write}(\mathit{keats},y)))$

We can now search for plans which make an event of Keats stopping writing necessarily occur, as distinct from those that merely make it possible, like $\alpha = \mathit{start}(\mathit{write}(\mathit{keats},y))$, by constructively searching for a proof for an event sequence α which the situation affords and which forces Keats to stop writing:

(84) $\mathit{affords}(\alpha) \wedge [\alpha]\mathit{forces}(\mathit{stop}(\mathit{write}(\mathit{keats},y)))$:

One such proof is the following:

(85) $\alpha = \mathit{start}(\mathit{write}(\mathit{keats},y));\mathit{arrive}(x)$

Again the examples are artificial: their usefulness for an account of tense and temporal anaphora will become apparent in the next chapter.

The *forces* relation can also be used to capture Goldman's 1970 "generation" relation between events that are so intimately contingent that one effectively "counts as" the other, as in *flipping a switch* and *turning on the light*, moving Pawn to Q7 and *giving check*, or *blowing up the bridge* and *destroying it*. For example we might have a rule saying that in any situation, flipping the switch *forces* turning on the light:

(86) $\models [\mathit{flip}]\mathit{forces}(\mathit{turnon})$

Together with the obvious preconditions, consequents, and frame axiom for *turnon*, this will allow us to construct sensible plans for getting the light to be on, via axiom 82. The difference between generation and the previous type of causal contingency (besides some purely temporal considerations that we will come to directly) lies in the fact that the previous causal relation was explicitly mediated in the logic via a property of the situation. However, it is important to be clear that, like the aspectual distinctions of chapter 2, the generation relation is *a way of regarding* the events in question, rather than a primitive relation in the knowledge representation. Still less is it a fact about the objective character of events in the physical universe. The lightswitch example could also be modelled in terms of flipping the switch causing a state in which current flows and in turn forces a turning on event. Many of Goldman's generation examples can be also be regarded as other varieties of causal relation, when viewed at other granularities.

Using the temporal indices developed earlier, we can incorporate the temporal simultaneity condition on the generation relation between events, writing 86 as follows:

(87) $\models [flip_t]_t forces(turnon_t)$

If we make the same assumption about clocktime that we made about spatial position, namely that it is autonomously updated, we can handle delayed causation, a phenomenon with which STRIPS is often claimed to have intrinsic difficulties. The standard examples of the problem involve placing arsenic in people's tea.

First we need the usual preconditions and dynamic axioms for the action of dying:

(88) a. $\models \neg dead \Rightarrow affords(die)$
b. $\models \{affords(die)\}[die]dead$

The consequent state of *drinking arsenic in one's tea at time t_1* holds if that is what you do at that time.

(89) $\models [_{t_0} drink_{t_1}(A_s)]_{t_2} consequent(drink_{t_1}(A_s))$

This consequent state persists through most actions except abnormal ones like administering an antidote or the victim dying. (Note that these are both events after which you cannot say that the victim "has drunk arsenic.")

It is well-known that if you have drunk arsenic in your tea and the time is four hours later then you necessarily die. We capture this fact of causation at a temporal distance be exploiting the consequent state fluent associated with the linguistic perfect:

(90) $\models consequent(drink_{t_1}(A_s)) \wedge clocktime(t_2) \wedge t_2 = (t_1 + 4)$
$\Rightarrow forces(die_{t_2}))$

Soif someone is alive and we want to achieve a situation *dead* we seek a sequence of actions α such that:

(91) $affords(\alpha) \wedge [\alpha]forces(die)$

Given some arithmetical equation-solving capabilities that we will not spell out here, this goal will yield a proof that you can get them to be not alive by a sequence of actions lasting four hours starting with them drinking arsenic, such as:

(92) $[_0 drink_0(A_s)]_4 dead$

(This analysis of temporal intervals in an instant-based calculus has much in common with Galton's 1990 alternative to Allen's 1984 interval-based approach).

We will return to the *forces* relation in the discussion of *when*-clauses in the next chapter, which is where it really comes into its own.

We can also reason correctly on this basis concerning the Auchterarder Ramifying Pizza Delivery problem.

Given the earlier representations of arriving (57), (58) and the accomplishment of moving to a place and arriving there, the goal goal (93a.) delivers a constructive proof (93b.) of a plan α for getting a car that is in Edinburgh to be at Auchterarder:

(93) a. $affords(\alpha) \wedge [\alpha] at(car, Auchterarder)$
b. $\alpha = start_0(move(car_1, Edinburgh, Auchterarder))$
$; arrive_2(car_1, Auchterarder)$

Slightly more interestingly, if we include an axiom describing the behavior of containers, we can prove of the situation that results from moving a car with a pizza in it to Auchterarder that the pizza is no longer in Edinburgh, but in Auchterarder:

(94) $\models in(x,y) \wedge at(y,l) \Rightarrow at(x,l)$

(95) a. $affords(\alpha) \wedge [\alpha] at(pizza, Auchterarder)$
b. $\alpha = start_0(move(car_1, Edinburgh, Auchterarder))$
$; arrive_2(car_1, Auchterarder)$

We include axioms concerning the preconditions and consequences of throwing pizzas from cars:

(96) a. $\models affords(eject(x,y))$
b. $\models \{affords(eject(x,y))\} \wedge in(x,y) \Rightarrow [eject(x,y)]$

As a consequence, any attempt to prove that the pizza is in Auchterarder after the following sequence will fail:

(97) $\alpha = [start_0(move(car_1, Edinburgh, Auchterarder))$
$; eject_1(pizza, car)$
$; arrive_2(car_1, Auchterarder)$

There is a close relation between the dynamic modalities $[\alpha]$ and von Wright's "and Next" operator T, which is often written $\bigcirc$ in other temporal logics—cf. (Goldblatt, 1992, ch.9). Von Wright's operator in turn has its origin in deontic logic (see Thomason 1981), and there is in general some affinity between deontic logic and dynamic logic. The interesting property of the system for present purposes is that it represents causally or contingently related sequences of actions.

Of course, to go beyond the simple histories discussed so far to deal with histories involving more than one instance of a given type of event such as $move(car_1, Edinburgh, Auchterarder)$, we need to beef up our axioms with Davidsonian event variables and inequalities to manipulate these temporal indices correctly, capturing fact such as that the *start* of any given event of spraying must precede its *stop*. Allen (1984), Allen and Hayes (1989), and Hayes (1995) provide an (interval-based) calculus and axioms for reasoning about the inequalities here established by the STRIPS instants. That calculus seems to transfer essentially unchanged.

While we have eschewed the standard axiomatic approach to modality in the above account, such a logic can be captured in the kind of axiomatisation that is standard in the literature in on dynamic logic—Harel 1984:512 *et seq.* provides a model which can be adapted to the more restricted deterministic logic that is implicit here by omitting some axioms and adding a further axiom of determinism. (See p. 522 et seq. I am indebted to Rich Thomason for suggesting this approach.) However, such a toy needs considerable further work to make it into a linguistically interesting object. In particular, it might benefit from the inclusion of a type system of the kind developed by Naumann (1995, 1997, 2001) to capture temporal ontologies like the one discussed in chapter 2 (although we have seen that entities like accomplishments can be captured in first-order terms like (73)). We must also extend it to capture the fact that events may be embedded within other events. We may find ourselves needing to express further varieties of distinct causal and modificational sequential relations between events, as Shoham (1988) and other authors cited earlier have suggested (although in the case of one prominent candidate, the "generation" relation of Goldman (1970) we have seen that this is not necessary). We may wish to introduce some coindexing device equivalent to Davidsonian **e** event variables to individuate distinct tokens of the same type of action. We also need to relate the contingent sequences to clock-time. Some of these extensions are touched on in the next chapter, which considers how to bring this apparatus more appropriately under the control of language, by making it refer to an actual chronicle or historical sequence of events.

Chapter 4

Temporal Reference

4.1 Basic Phenomena and Descriptive Frameworks

In the discussion so far, we have largely ignored the question of how the Reichenbachian reference point is represented and accessed, and the anaphoric nature of tense. Several logical and computational approaches have explored this possibility.

Temporal anaphora, like all discourse anaphora and reference resolution, is even more intimately dependent upon world knowledge than the other temporal categories that we have been considering. In order to control this influence, we will follow the style of much work in AI, drawing most of our examples from a restricted domain of discourse. We will follow Isard (1974) in taking a board game as the example domain. Imagine that each classical model in the structure of figure 3.9 is represented as a database, or collection of facts describing not only the position of the pieces in a game of chess, and the instantaneous moves at each frame, but the fact that at certain times durative or composite events like *exchanging Rooks* or *White attacking the Black Queen* are in progress across more than one state.

Consider the following examples from such a domain

(1) a. When I took your pawn, you took my queen.
    b. I took your pawn. You took my queen

The *when*-clause in (1a) establishes a reference point for the tense of the main clause to refer to anaphorically, just as the definite NP *Keats* establishes a referent for the pronoun. Indeed, as Webber (1988) pointed out, the *when*-clause itself behaves like a definite, in that it seems to presuppose that the event of *my*

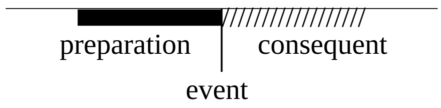

Figure 4.1: The nucleus of chapter 2

taking your pawn is identifiable to the hearer. (Of course, the reader will have effortlessly accommodated this presupposition.) The first sentence in (1b) behaves exactly like the *when* clause in setting the reference point for the second. The only difference is that the simple declarative *I took your pawn* itself demands a previously established reference point to be anaphoric to, whereas the *when* clause merely presupposes that you know enough to establish such a reference point from scratch.

As has been frequently noticed, the state to which the tense in *you taking my queen* refers in (1a) above, is not strictly the state in which *I took your pawn*. It is the state that *resulted from* that action. However, it is not invariably the case that the temporal reference point moves on in this way. Most obviously, a stative main clause is primarily predicated of the original reference point of the *when*-clause:

(2) When I took your pawn, I did not know it was protected by your knight.

(Presumably, the ignorance in question may have ended with that very move.) Events also may be predicated of the original reference point, rather than moving the action on:

(3) When I took your pawn, I used a rook.

In fact, as Ritchie (1979) and Partee (1984) have pointed out, in strictly temporal terms, we can find main clauses that *precede* the reference point established by a *when* clause:

(4) When I won my only game against Bobby Fischer, I used the Ruy Lopez opening.

These phenomena arise because the temporal referent is *not* strictly temporal. Rather than being a time or an interval, it is (a pointer to) an event-nucleus of exactly the kind that was used in chapter 2 to explain the aspectual sort hierarchy and possible coercions among the *Aktionsarten*. That is, it is a structure of the kind shown in figure 2.2, repeated here as figure 4.1. It will be recalled that the preparation is an activity, the consequent is a (perfect) state, and that the core event is an achievement. (Recall that any event-type can turn into an achievement via the sort-transition schema in figure 2.1).

In the terms of our modal frame, the preparation of an event is the activity or action that led to the state in which that achievement took place. The consequent is the consequent state, and as suggested in the earlier discussion of the modals, includes the entire subgraph of states accessible from that state. The referent-setting effect of a *when*-clause can then be seen as identifying such a nucleus. The main clause is then temporally located with respect to the nucleus. This may be by lining it up with the core event itself, either as a property of the initial state, as in example 2, or as a property of the transition itself, as in 3. Alternatively, since accessibility is defined in terms of the subsequent actions, the actual subsequent action is a possible main clause, as in 1. Or the main clause may be located with respect to the preparation, as in 4. Which of these alternatives a given example gives rise to is a matter determined by the knowledge representation, not by rules of the semantics.

On the assumption that the consequent in the nuclear referent includes the entire subgraph of future states, the information needed by conditionals, modals, and other referent-setting adverbials will be available:

(5) a. If you take my queen, you may win.
b. If you had taken my queen, you might have won.
c. Since you took my queen, you have been winning.

All of this suggests that states or partial possible worlds in a logic of action deriving ultimately from von Wright and McCarthy and Hayes, with a much enriched ontology involving a rather intimate connection to the knowledge-base, are appropriate candidates for a Reichenbachian anaphoric account of tense and temporality. But this does not tell us how the temporal referent is set up to act as a referent for anaphora.

4.2 Logical and Computational Approaches

It is possible in principle to embody a Reichenbachian account in a pure modal logic, say by developing "multi-dimensional" tense logics of the kind used by Nerbonne (1984) (see van Benthem (1991a,b), and (van Benthem, 1995, section III.3)). However, the dynamic event calculus of chapter 3 offers a promising candidate for a representation of Riechenbach's reference point R, in the form of deterministic event sequences $[\alpha]$. This opens up the possibility of applying the general modal apparatus developed so far, not only for quantifying over states, but to act as the temporal link between sentences and clauses, as in *when*-clauses and multi-sentence discourse.

Most computational approaches have equated sentential temporal anaphora with *discourse* temporal anaphora, rather than any structurally bound variety. Thus Winograd (1972), Isard (1974), and the present author treated the establishment of temporal (and pronominal) referents as temporary side-effects to a single STRIPS-like database. A reference-point establishing *when*-clause or conditional had the effect of setting the database to the state of the world at the (in Isard's case, possibly counterfactual) time in question. The way this was actually done was to "fast-forward" (or -backward) the world to the situation in question, using the history of events to carry out the sequence of updates and retractions necessary to construct the state of the world at the reference point.

Within the situation calculus and its descendants including the dynamic event calculus, this strategem is unnecessary.

The history of events is a sequence such as the following:

(6) $start(write(keats, in\_disgust)); arrive(chapman)$
$; stop(write(keats, in\_disgust)$

The referent of a *when*-clause, such as *When Chapman arrived*, is simply the sequence up to and including $arrive(chapman)$, namely:

(7) $start(write(keats, in\_disgust)); arrive(chapman)$

To identify this referent in the history we need the following definition of a relation we might call *evoke*. This is merely a logic-programming device which defines a search for a deterministic event sequence of the form $[\alpha;\beta]$ over a history in which the sequence operators are "left-associative":

(8) a. $\models evoke((\alpha;\beta),(\alpha;\beta))$
b. $\models \beta \neq \delta \wedge evoke((\alpha;\beta),\gamma) \Rightarrow evoke((\alpha;\beta),(\gamma;\delta))$

(8a) says that if you are looking for a sequence of events ending in an event β and the whole history ends in β then the history is the result you want. (8b) says that if the history ends with an event δ that isn't β, then you should try to recursively evoke a history ending in β from the history before δ, namely γ. Evokable α are by definition possible, even for counterfactual histories.

The referent-setting effect of *when* can now be captured to a first approximation in the following rules, which first find the current history of events, then *evoke* a suitable reference point, then test for the appropriate relation *when*. (Again this is a logic programming hack which could be passed over):

(9) a. $\models state(\gamma) \wedge S(history) \wedge evoke((\alpha;\beta), history) \wedge [\alpha;\beta]\gamma \Rightarrow when(\beta,\gamma)$
b. $\models event(\varepsilon) \wedge S(history) \wedge evoke((\alpha;\beta), history) \wedge [\alpha;\beta]forces(\varepsilon)$
$\Rightarrow when(\beta,\varepsilon)$

The predicate S determines the Reichenbachian speech point, which is an event or sequence of events. $S(history)$ is assumed to be available in the database, as a fact. The first rule, a, applies to *when* sentences with state-type main clause propositions, and says that $when(\beta,\gamma)$ is true if γ is a state, and you can *evoke* an event sequence ending in β after which γ holds. The second applies to *when* sentences with event-type main clauses, and says that $when(\beta,\varepsilon)$ is true if ε is an event and you can *evoke* an event sequence ending in β after which $forces(\varepsilon)$ holds. The question a, below, concerning the ensuing state, therefore translates into the query b:

(10) a. When Chapman arrived, was Keats completing *In Disgust of Vulgar Superstition*?
b. $when((\alpha; arrive(chapman)),$
$progressive(achievement(complete(keats, in\_disgust))))$

The progressive is, it will be recalled, a state, so in our greatly simplified world, this is true, despite the fact that under the closed world assumption Keats did not complete the poem, because of the earlier elimination of the imperfective paradox.

A *when*-question with an event in the main clause, as in a, below, translates as in b:

(11) a. When Chapman arrived, did Keats stop writing *In Disgust of Vulgar Superstition*?
b. $when((\alpha; arrive(chapman)), stop(write(keats, in\_disgust)))$

In the case to hand, this last will yield a proof with the following constructive instantiation:

(12) $when((start(write(keats, in\_disgust)); arrive(chapman)),$
$stop(write(keats, in\_disgust)))$

In either case, the enduring availability of the Reichenbachian reference point for later simple tensed sentences can be captured on the assumption that the act of *evoking* a new referent causes a sideffect to the database, causing a new fact—say of the form $R(\alpha)$—to be asserted, after any existing fact of the same form has been removed, or retracted. (We pass over the formal details here, merely noting that for this purpose a blatantly non-declarative STRIPS-like formulation seems to be the natural one, although we have seen how such non-declarativity could in principle be eliminated from the system. Lin and Reiter (1995) show how such a process of "progressing" the database can be defined on the basis of the declarative representation.)

The representation captures via the *forces* predicate the fact that Keats stopped writing the poem *because* Chapman arrived, whereas Chapman merely arrived *after* Keats started writing, not because of it. The importance of this fact is that it is really only under the assumption of a causal relation of some kind that the answer to the question should be affirmative. For the chronicle that we have been considering sentence a, below, is presuppositionally infelicitous, unlike the corresponding concatenation b:

(13) a. #When Keats started writing *In Disgust of Vulgar Superstition*, Chapman arrived.
b. Keats started writing *In Disgust of Vulgar Superstition*. Chapman arrived.

This suggests that a relation of contingency between main and subordinate clause is an intrinsic component of the semantics of *when*, rather than a mere implicature arising from relevance assumptions as assumed in Kamp and Reyle's closely related account (Kamp and Reyle 1993).

The definition of the semantics of *when*-clauses with event main clauses in terms of the *forces* relation, and the possibilities set out in chapter 3 for defining temporal relations over sequences of events, also captures the conditions under which the main clause event *follows* the event of the *when*-clause in time, as in a, and when it is temporally coincident with it, as with events that have been held to be related by "generation", as in b:

(14) a. When Chapman arrived, Keats stopped writing.
b. When Chapman flipped the switch, he turned on the light.

If the contingency relation is causal, then the main clause follows at whatever interval is supported by common-sense knowledge about the events in question. If the relation is generational, then the two are simultaneous.

Of course, it will be clear from the earlier discussion that such a system remains oversimplified. Such sentences also suggest that the event sequences themselves should be considerably enriched on lines suggested in earlier chapters. In particular they should be capable of being structured into *nested* structures of causal or, more generally, contingent sequences of the kind discussed at the end of chapter 2, as when an iterated event of, say, *playing the minute waltz in less than sixty seconds for an hour* causes another, such as *tearing* *ones hair out with irritation*.

Since we have also observed that main clause events may be simultaneous with, as well as consequent upon, the *when* clause event, events must also be permitted to be simultaneous, perhaps using the connective $\cap$ introduced by Peleg (1987) to capture the relation between embedded events like *starting to write "In Disgust of Vulgar Superstition"* and *starting to write*, generalising the above rules accordingly. Partial ordering of events must also be allowed. The inferential possibilities implicit in the notion of the nucleus must be accommodated, in order to capture the fact that one event may cause the preparation of another event to start, thereby embodying a non-immediate causal effect.

Very little of this work has been done, and it may be unwise to speculate in advance of concrete solutions to the many real problems that remain. However the limited fragment outlined above suggests that the dynamic event calculus may be a promising framework in which to pursue this further work and bring together a number of earlier approaches. In this connection, it is perhaps worth remarking that, of the seven putative limitations of the situation calculus and its relatives claimed in the critical review by (Shoham and Goyal, 1988, p.422-424), five (limitation to instantaneous events, difficulty of representing non-immediate causal effects, ditto of concurrent events, ditto of continuous processes, and the frame problem) either have been overcome or have been addressed to some extent in the published work within the situation calculus. Of the remaining two (the qualification problem and the ramification problem) the ramification problem in the narrow sense of *known* causal effects of actions has been addressed above and by Schubert and Reiter. In the broader sense of *unanticipated* contingencies or ramifications, and similarly unanticipated preconditions, or qualifications, these problems have not been overcome in *any* framework, possibly because they do not belong in the logic at all. The next chapter explores these questions further.

The non-computational approaches to temporal anaphora, in contrast, to those just described, have tended to equate all temporal anaphora with structurally bound anaphora. DRT treats temporal referent(s) much like nominal referents, as localised side-effects. This mechanism is used to extend the scope of the temporal referent beyond the scope that surface syntax would most immediately suggest, in much the same way that the scope of nominal referents is extended to capture such varieties of nominal anaphor as "donkey pronouns", and the approach is generalised to modals and conditionals by Roberts (1989), and Kamp and Reyle (1993). They, like Dowty (1986), assume that events invariably move the temporal reference point forward while states do not—cf. (Kamp and Reyle, 1993, p. 528), which in general is not the case. (Indeed in the case of the latter authors, this assumption is somewhat at odds with their adoption elsewhere of a nucleus-like structure over events—cf. p. 558.) However, both note the oversimplification, and their theories remain entirely compatible in principle with the present proposal to bring this question under the control of context and inference, perhaps along lines suggested by Lascarides and Asher (1993b), who incorporate persistence assumptions of the kind discussed above.

Interestingly, (Groenendijk and Stokhof, 1991, p.50) show how the scope-extending mechanism of DRT can be captured straightforwardly in a first-order variety of dynamic logic, dynamic predicate logic (DPL —cf. Dekker (1979)). DPL is in some ways quite distantly related to dynamic logic in the sense that the term is used here. (Indeed, many of the uses to which these authors apply DPL—such as capturing the slightly extended notion of quantifier scope that may be implicit in the interpretation of "donkey anaphora"—can be seen as extensions to the first-order core of the dynamic logic developed here, rather than as having to do with the dynamic modalities themselves.) However, the mechanism that Groendijk and Stokhof propose, which directly models the side-effects implicit in assignment to variables, seems to be generalisable to the DRT treatment of inter-clause temporal anaphora and the Reichbachian reference point, suggesting a way to unify all of the approaches discussed here. Recent work by Muskens (1995) is very much along these lines.

Chapter 5

The Frame Problem Revisited

It is ironic, since the event calculus and its dynamic logic relatives are descended from Kowalski's declarativisation of the STRIPS solution to the frame problem, that much of the opposition to such logics as a representation for knowledge about time and action has come from AI researchers working on Knowledge Representation and the Frame Problem itself. To understand these criticisms, it is important to realise that in the years since the first formulation of the Frame Problem, it has come to have a much more extended sense.

5.1 The Frame Problem Revisited

In the narrow original sense discussed in earlier chapters, the Frame Problem is simply the problem of efficiently representing models for dynamic worlds.

The original situation calculus represented dynamic worlds in terms of a) actions pairing preconditions on the situations in which they can occur with the "relevant" properties of the situation that ensues, and b) a large number of "frame axioms" pairing each fluent with each action and saying which predication involving that fluent hold in that ensuing situation.

Because the way we think about the world (necessarily?) has the property that actions are modular and affect relatively few properties of situations, frame axioms of this kind are inefficent. It is precisely *because* color is not mentioned in the definition of movements that it has a very simple frame axiom. Indeed, we should be able to *derive* the frame axioms (b) from the interesting rules (a), as Reiter shows (1991; 2001).

The STRIPS solution is to represent actions solely in terms of preconditions and consequences, and to build inertia into the model, so that nothing changes from one situation to the next unless a strips rule says it does.

STRIPS did this via sideffects on a database. Kowalski shows how to declarativize this notion of inertia in the logic, in the form of very general frame axioms, of which there is exactly *one per event-type*. Chapter 3 showed that the full Vendlerian taxonomy of event-types and how the Reichenbachian referential notion of tense, together with notions of counterfactuality can be represented in this sort of framework. In particular, progressive and perfective/consequent *states* can be represented in terms of instantaneous initiating and ending *events*.

We saw how the specific situation indices, and hence the *holds* and *result* predicates, could be stripped out of the notation, using the dynamic event calculus modal logic. This reduction in temporal baggage is important because we are soon going to reintroduce the idea of linear time.

Note that the situation/event calculus version of the STRIPS solution involves a *closed world assumption*, which says that we have complete and consistent knowledge about the way the world works. We just leave a lot of the mechanics implicit in the frame axioms. Notice that if we thought about the world in terms of processes that changed most fluents, this might be *less* efficient than just listing all facts explicitly. the fact that we *don't* think of the world that way is a tangible consequence of the way we are grounded in the continuum.

Note also that this system is equally compatible with searching forwards in time for a sequence of actions resulting in a goal state *(Planning)* and searching backwards in time for a sequence of actions that brought a current situation to pass *(Explanation)*. (This is characteristic of all default systems based on closed world assumptions—if you know Tweety is a bird, that birds fly, but Tweety cannot fly and you have a good circumscription of the exceptions to your default rules then you will infer that maybe he is a penguin. If you do *not* have a good circumscription, you are doomed as far as inference goes. The same goes for frame axioms.)

In the wider sense of the frame problem, we saw that the solvable part of the ramification problem could be handled by appropriate choices of representation within the Dynamic version of the Situation/Event Calculus and that this version of the representation naturally supported notions like simultaneity and delayed action. The qualification problem is less easy to handle for any knowledge representation, and requires a rexamination of the place of nonmomotonicity in the system.

5.2 Nonmonotonicity

We should distinguish nonmonotonicity arising from acquisition of new knowledge from nonmonotonicity in the logic itself.

What do you do when you know that Tweety is a bird, so you assumed be could fly, but subsequently hear that he's a penguin, which you know means he cannot? With any default logic, there are two solutions.

1. Redo the entire inference from scratch with the new piece of information. If you have a good theory of the world you will end up this time inferring neither that Tweety can fly, *nor anything that you inferred from him flying*.

2. Do sufficient bookkeeping via strings of indices expressing support relations between inferences that you can do *reason maintainance*. That is, keep track of the fact that Tweety flying was a default inference, and withdraw it *and everything whose inference depended on it* leaving everything else intact.

Both tactics are likely to be expensive. If you have a nearly perfect theory the first might be cheaper. If you have a really terrible theory the second might win (but maybe you should get a better theory). Cheaper alternatives such as probabilities or confidence estimates do not seem to offer a way of compensating for a basically wrong theory (although probablistic approaches may be very helpful in finding a *better* theory—see below).

Both tactics amount to identifying an alternative totally monotonic logic for which the closed world assumption again holds.

5.3 Nixon Diamonds

What do you do when you know that Nixon is a quaker and all quakers are by default pacifists, and that Nixon is a republican and (let us suppose) that no republicans are pacifists? Clearly there is absolutely nothing that you can do as far as inference with this logic goes. You have a bad representation. Either Nixon is not a true republican, or he's not a true quaker, or something.

There have been a number of attempts to resolve logic-internal nonmonotonicity via criteria for choosing one closure of the logic over another. Most of these depend on some idea of choosing simpler chains of reasoning over more complex ones (usually with some invocation of ideas from philosophy of science like Ockham's Razor). However, a moments thought about Nixon diamonds suggests that one will always be able to construct some example in which such content-free techniques will give the answer you don't expect.

This in turn suggests there should be a strict separation between *the logic*, good or bad, and *processes for compensating for bad knowledge representations*.

5.4 Shooting Problems

Schubert (1994), following Sandewall (1994), examines a number of variants of the Temporal Projection Problem or "Yale shooting problem," following Sandewall, including the Stokholm delivery scenario, from which the Auchterarder pizza delivery problem of chapter 3 was adapted.

The original Yale Shooting Problem was set up by Hanks and McDermott as a temporal Nixon Diamond. Axioms say that if you *load* a gun it is *loaded*, if the gun is *loaded* and you *shoot* someone they are ¬*alive*, that you can *wait* and that if you *unload* a gun it is ¬*loaded*. They intended the problem to show that once you let Nixon diamonds in, there is no *theory internal* way of deciding which closure of the diamond to prefer as a simpler explanation.

The way this is set up makes crucial use of classical negation (as opposed to negation-as-failure that we are using), and allows them to show that if you hear that a gun was *loaded* at s_0, we *waited* for a bit, then fired, that a closure in which someone is ¬*alive* is no more minimal than one in which an unmentioned *unloading* event is assumed.

The details don't matter because it is clear that what is wrong here is that H&S have (deliberately) not represented the knowledge. That is their point.

As with the original Nixon diamond, a vast literature ensued attempting to show that you could have criteria of parsimony to choose between closures of the diamond. None of the proposed solutions of this kind appear to work, and it is hard to see how they could.

Schubert shows that by assuming "explanation closure"—the dual of Kowalski's planning-oriented closed world assumption—then the problem goes away, for the obvious reason. That is, you assume that if there was an unloading you would know about it. (Schubert has an elegant justification of this in terms of Gricean assumptions).

Interestingly, nothing in Schubert's proofs seems to depend on the explicit representation of clock time. Parallel proofs must exist in the pure event calculus, and are suggested as an exercise.

5.5 The Way Forward

I have argued at some length that the situation calculus, as modified by Kowalski as the event calculus, and as transformed here into the dynamic event calculus, is a much better basis for a representation of common-sense knowledge about time and causation than has been assumed. The one problem to to which it is indisputably vulnerable appears to be one to which any default representation must by definition be heir, namely the qualification problem.

The only general solution to the qualification problem seems to be to change to a new but still monotonic logic, perhaps by adding a new precondition concerning absence of potatoes from tailpipes or by recomputing circumscription. This could easily get out of hand, defeating the original purpose of default knowledge representation.

The ways of computing such changes to the logic seem quite different in character to the logic of change itself. To build large robust knowledge representations is likely to need machine learning techniques. While such techniques seem likely to essentially be limited to learning finite-state machines and simple classifiers, such techniques may be applied at several levels of a structured representation. Interesting techniques are:

1. Neural network representations. Inducing representations for natural concepts shows every sign of being very hard to do without grounding the system in the sensory manifold in much the way that we and other animals are grounded ourselves.

2. The full hypothetico-deductive method. For rather similar reasons, no-one really knows how to automate this but the idea of "active learning" seems related.

For the moment we are probably stuck with building badly-grounded restricted dynamic logics with nicer proof-search properties, and the theorem provers to go with them. Wallen (1990), Ohlbach (1991), and Stone (1999); Stone and Hardt (1997); Stone (1998) are important in this connection.

Chapter 6

Conclusion

The analysis so far has built solely upon observations from English. Nevertheless, the claim that the semantics outlined above depends directly upon the conceptual representation of action and contingency suggests that this semantics might be universal, despite considerable differences in its syntactic and morphological encoding across languages. Discussion of the evidence for this claim would take us beyond the scope of this essay. However, the available reviews of this extensive literature (e.g. Dahl (1985) and Smith (1991)) seem to lend some support to the following brief observation on this question.

Benjamin Lee Whorf once observed that the auxiliaries and inflections associated with verbs in Amerindian languages appeared to be semantically quite unlike the corresponding categories in English and other European languages. The Amerindian categories seemed to be more concerned with various aspects of the speakers' evidential and consequential relation to events, rather than the strictly temporal relations which Whorf assumed were implicated in the corresponding devices of English. He suggested, controversially, that these differences relected differences in modes of thinking about events and time.

The work described above suggests that such differences across languages are superficial. Ironically, the English tense/aspect system seems to be based on semantic primitives remarkably like those which Whorf ascribed to Hopi. Matters of temporal sequence and temporal locality seem to be quite secondary to matters of perspective and contingency. This observation in turn suggests that the semantics of tense and aspect is profoundly shaped by concerns with goals, actions and consequences, and that temporality in the narrow sense of the term is merely one facet of this system among many.

English and other "Standard Average European" languages are on this account also deeply imbued with an affordance-based view of objects that is similarly explicit in the forms of less familiar languages. For example, we saw that North American Indian languages such Navaho appear to lexicalize nouns as a default affordance, of the kind we have assumed is available universally for purposes of plan construction.

Such properties seem to force the logic that is required to capture natural language semantics to be the particular kind of dynamic event calculus outlined above, whose structure is intimately related to knowledge of action, the structure of episodic memory, and the computational process of inference. Crucially, that calculus builds into the logic itself, via an event-based accessibility relation and an inertial representation of possible situations, structures based on the insights of the AI planning literature. In particular, the emphasis in the Kripke structure that models the logic is on the arcs, in which much of the specification of states is left implicit. Because of these properties, the search space for constructive proofs in this logic has the same structure as the proofs themselves, a property which is fundamental to the treatment of logics as logic programming languages, and which is lacking in many temporal logical approaches to the planning problem. The Dynamic Event Calculus is also one in which the richer ontology of telic and durative events, intervals, and chronological time that is manifested in natural languages is defined in terms of a rich set of primarily causal and contingent relations over timeless instantanteous primitive events.

Further Reading:

The literature on temporality and representation of causal action is vast, and I am painfully aware of having been forced to pass over entirely or to treat rather superficially a great deal of important and relevant work. The following sources are intended to provide a means of entry to a more extensive literature than I have been able to discuss within the confines of these notes.

Hughes and Cresswell 1968 remains an important source for early axiomatic approaches to modal logic, and its early historical development. van Benthem 1983 is a very readable survey of Tense Logic, with particular attention to the effects of different ontological commitments, including those related to the representation of various views of time implicit in modern physics. A number of papers in volume II of Gabbay and Guenthner's 1984 *Handbook of Philosophical Logic* cover recent developments, including those by Bull & Segerberg, Burgess, Thomason, van Benthem, and Åqvist. Harel 1984 in the same volume and Goldblatt 1992 are resources for dynamic logic and related systems, including temporal logic. The latter approach is presented at greater length in Gabbay et al. 1994, which includes some discussion of the linguistic and computational issues raised here. The *Handbook of Logic and Language* edited by van Benthem and ter Meulen 1997 brings together a number of recent survey chapters on related logical issues in other areas of natural language semantics, including an earlier version of some of the material presented here. Kamp and Reyle 1993 discuss tense and aspect within DRT. Groenendijk and Stokhof 1991 discuss the expressibility of DRT in dynamic predicate logic, which Muskens 1995 extends to tense. The early paper by McCarthy and Hayes 1969 remains an excellent review of modal logic from a computational perspective, and is one of the few sources to explicitly relate computational and logical approaches, as is Nilsson's elegant 1980 text. The invaluable collections of readings in artificial intelligence, non-monotonic reasoning, and planning respectively edited by Webber and Nilsson 1981, Ginsberg 1987, and Allen, Hendler, and Tate 1990 are sources which reprint many of the computational papers discussed above, and much other recent work in AI knowledge representation which it has not been possible to survey here. Galton 1987 is a recent collection of essays bringing together logicians and computer scientists on the question of temporal representations. The special issue on tense and aspect of *Journal of Computational Linguistics*, (1988, vol. 14.2) is another source for computational linguistic approaches, including the one developed here. Lyons 1977, Comrie 1976; 1985, Dahl 1985, Binnick 1991, Declerck 1991, and Smith 1991 survey the tense and aspectual systems of English and a considerable number of languages from a linguistic standpoint.

References

Abusch, Dorit. 1997a. "Generalizing Tense Semantics for Future Contexts." In Susan Rothstein, ed., Events and Grammar, 13–33. Dordrecht: Kluwer.

Abusch, Dorit. 1997b. "Sequence of Tense and Temporal *de re*." *Linguistics and Philosophy*, 19, 1–50.

Agre, Phillip, and David Chapman. 1987. "Pengi: An Implementation of a Theory of Activity." In *Proceedings of the Sixth National Conference on Artifi cial Intelligence (AAAI-87)*. Los Altos, CA: Morgan Kaufmann.

Allen, James. 1984. "Towards a General Theory of Action and TimeD." *Artifi cial Intelligence*, 23, 123–154.

Allen, James, and George Ferguson. 1994. "Actions and Events in Interval Temporal Logic." *Journal of Logic and Computation*, 4, 531–579.

Allen, James, and Patrick Hayes. 1989. "Moments and Points in an interval-based temporal logic." *Computational Intelligence*, 5, 225–238.

Allen, James, James Hendler, and Austin Tate, eds. 1990. *Readings in Planning*. Palo Alto CA: Morgan-Kaufmann.

Bacchus, Facchiem, and Froduald Kabanza. 2000. "Using Temporal Logics to Express Search Control Knowledge for Planning." *Artifi cial Intelligence*, 116, 123–191.

Bach, Emmon. 1980. "Tenses and Aspects as Functions on Verb-Phrases." In Christian Rohrer, ed., *Time, Tense, and Quantifi ers*, 19–37. Tübingen: Niemeyer.

Bach, Emmon, Colin Brown, and William Marslen-Wilson. 1986. "Crossed and Nested Dependencies in Dutch and German." *Language and Cognitive Processes*, 1, 249–262.

Barwise, Jon, and John Perry. 1983. *Situations and Attitudes*. Cambridge MA: MIT Press.

Bäuerle, Rainer. 1979. "Tense Logics and Natural Language." *Synthèse*, 40, 225–230.

Bennett, Michael, and Barbara Partee. 1972. "Towards the Logic of Tense and Aspect in English." Published by Indiana University Linguistics Club 1978.

van Benthem, Johan. 1983. *The Logic of Time*. Kluwer: Dordrecht.

van Benthem, Johan. 1991a. "General Dynamics." *Theoretical Linguistics*, 16, 159–172.

van Benthem, Johan. 1991b. *Language in Action*. Amsterdam: North Holland.

van Benthem, Johan. 1995. "Temporal Logic." In Dov Gabbay et al., ed., *Handbook of Logic in Artificial Intelligence and Logic Programming*. Oxford: Clarendon Press.

van Benthem, Johan, and Alice ter Meulen. 1997. *Handbook of Logic and Language*. Amsterdam: Elsevier/North Holland.

Bibel, Wolfgang. 1986. "A Deductive Solution for Plan Generation." *New Generation Computing*, 4, 115–132.

Binnick, Robert. 1991. *Time and the Verb: a Guide to Tense and Aspect*. Oxford: Oxford University Press.

Blackburn, Patrick, Claire Gardent, and Maarten de Rijke. 1993. "Back and Forth through Time and Events." In *Proceedings of the 9th Amsterdam Colloquium*, 161–175. ILLC, University of Amsterdam.

Blass, Andrew. 1992. "A game semantics for linear logic." *Annals of Pure and Applied Logic*, 56, 183–220.

Boyd, Julian, and James Thorne. 1969. "The Semantics of Modal Verbs." *Journal of Linguistics*, 5, 57–74.

Bratman, Michael, David Israel, and Martha Pollack. 1988. "Plans and Resource-Bounded Practical Reasoning." *Computational Intelligence*, 4, 349–355.

Caenepeel, Mimo. 1989. *Aspect, Temporal Ordering, and Perspective in Narrative Fiction*. Ph.D. thesis, University of Edinburgh.

Carlson, Greg. 1977. *Reference to Kinds in English*. Ph.D. thesis, University of Massachusetts, Amherst.

Comrie, Bernard. 1976. *Aspect*. Cambridge: Cambridge University Press.

Comrie, Bernard. 1985. *Tense*. Cambridge: Cambridge University Press.

Cooper, Robin. 1986. "Tense and Discourse Location in Situation Semantics." *Linguistics and Philosophy*, 9, 17–36.

Cresswell, Max. 1974. "Adverbs and Events." *Synthèse*, 28, 455–481.

Cresswell, Max. 1985. *Structured Meanings*. Cambridge MA: MIT Press.

Cresswell, Max. 1988. *Semantical Essays*. Dordrecht: Kluwer.

Cresswell, Max. 1990. *Entities and Indices*. Dordrecht: Kluwer.

Crouch, Richard, and Steve Pulman. 1993. "Time and Modality in a Natural Language Interface." *Artificial Intelligence*, 63, 265–304.

Dahl, Östen. 1985. *Tense and Aspect Systems*. Oxford: Blackwell.

Davidson, Donald. 1967. "The logical form of action sentences." In Nicholas Rescher, ed., *The Logic of Decision and Action*. Pittsburgh: University of Pittsburgh Press.

Declerck, Renaat. 1991. *Tense in English: Its Structure and Use in Discourse*. Routledge.

Dekker, Paul. 1979. *Transsentential Meditations: Ups and Downs in Dynamic Semantics*. Ph.D. thesis, University of Amsterdam. ILLC dissertation Series, 1993, no. 1.

deKleer, Johan. 1984. "Choices without Backtracking." In *Proceedings of the 4th National Conference on Artifi cial Intelligence*, 79–85.

Dowty, David. 1979. *Word Meaning in Montague Grammar*. Dordrecht: Reidel.

Dowty, David. 1982. "Tenses, Time-Adverbs, and Compositional Semantic Theory." *Linguistics and Philosophy*, 5, 23–55.

Dowty, David. 1986. "The Effects of Aspectual Class on the Temporal Structure of Discourse: Semantics or Pragmatics?" *Linguistics and Philosophy*, 9, 37–62.

Doyle, Jon. 1979. "A Truth Maintenance System." *Artifi cial Intelligence*, 12, 231–272.

Eberle, Kurt. 1990. "Eventualities in Natural Language Understanding Systems." In K. Bläsius, U. Hedtstück, and C. Rollinger, eds., *Sorts and Types in Artifi cial Intelligence*. Berlin: Springer-Verlag.

Enç, Murvet. 1981. *Tense Without Scope*. Ph.D. thesis, University of Madison, Wisconsin.

Enç, Murvet. 1987. "Anchoring Conditions for Tense." *Linguistic Inquiry*, 18, 633–657.

Fernando, Timothy. 2004. "Inertia in Temporal Modification." In *Semantics and Linguistic Theory (SALT), 14*. Northwestern University.

Fernando, Timothy. 2005. "Schedules in a Temporal Interpretation of Modals." *Journal of Semantics*, 22, 211–229.

Fikes, Richard, and Nils Nilsson. 1971. "STRIPS: a New Approach to the Application of Theorem Proving to Problem Solving." *Artifi cial Intelligence*, 2, 189–208.

Fine, Kit. 1975. "Critical notice of Lewis' "Counterfactuals"." *Mind*, 84, 451–458.

van Fraassen, Bas. 1971. *Formal Semantics and Logic*. New York: Macmillan.

Frisch, Alan, and R Scherl. 1991. "A General Framework for Modal Deduction." In *Proceedings of the Conference on Knowledge Representation*, 196–207. Los Altos, CA: Morgan Kaufmann.

Gabbay, Dov, and Franz Günthner. 1984. *Handbook of Philosophical Logic, II*. Dordrecht: Reidel.

Gabbay, Ian Hodkinson, Dov, and Mark Reynolds. 1994. *Temporal Logic: Mathematical Foundations and Computational Aspects, I*. Oxford: Clarendon Press.

Gagnon, Michel, and Guy Lapalme. 1995. "From Conceptual Time to Linguistic Time." *Computational Linguistics*, 91–127.

Galton, Anthony. 1984. *The Logic of Aspect*. Oxford: Clarendon Press.

Galton, Anthony, ed. 1987. *Temporal Logics and their Applications*. New York: Academic Press.

Galton, Anthony. 1990. "A Critical Examination of Allen's Theory of Action and Time." *Artifi cial Intelligence*, 42, 159–188.

Garrod, Heathcote. 1954. "Commentary on Keats' sonnet "Written in disgust of vulgar superstition"." In Heathcote Garrod, ed., *Keats' Poetical Works*. Oxford: Clarendon Press.

Geffner, Hector. 1990. "Causal Theories for Nonmonotonic Reasoning." In *Proceedings of the National Conference on Artifi cial Intelligence*, 524–530. Boston: AAAI.

Gelfond, Michael, and Vladimir Lifschitz. 1993. "Representing Actions and Change by Logic Programs." *Journal of Logic Programming*, 17, 301–321.

Gibson, James. 1966. *The Senses Considered as Perceptual Systems*. Boston, MA: Houghton-Mifflin Co.

Gibson, James. 1979. *The Ecological Approach to Visual Perception*. Boston, MA: Houghton-Mifflin Co.

Ginsberg, Matthew. 1986. "Counterfactuals." *Artifi cial Intelligence*, 30, 35–81.

Ginsberg, Matthew, ed. 1987. *Readings in Nonmonotonic Reasoning*. Palo Alto CA: Morgan-Kaufmann.

Girard, Jean-Yves. 1995. "Linear Logic: its Syntax and Semantics." In Jean-Yves Grirard, Yves Lafont, and Laurent Regnier, eds., *Advances in Linear Logic*, vol. 222 of *London Mathematical Society Lecture Notes*, 1–42. Cambridge: Cambridge University Press.

Goldblatt, Robert. 1980. "Diodorean Modality in Minkowski Spacetime." *Studia Logica*, 39, 219–236.

Goldblatt, Robert. 1992. *Logics of Time and Computation*. Chicago: CSLI/Chicago University Press. Second edition.

Goldman, Alvin. 1970. *A Theory of Human Action*. Englewood Cliffs, NJ: Prentice Hall.

Groenendijk, Jeroen, and Martin Stokhof. 1991. "Dynamic Predicate Logic." *Linguistics and Philosophy*, 14, 39–100.

Harel, David. 1984. "Dynamic Logic." In Dov Gabbay and F. Guenthner, eds., *Handbook of Philosophical Logic*, vol. II, 497–604. Dordrecht: Reidel.

Hayes, Patrick. 1971. "A Logic of Actions." *Machine Intelligence*, 6, 495–520.

Hayes, Patrick. 1995. "A Catalog of Temporal Theories." Tech. Rep. UIUC-BI-AI-96-01, University of Illinois.

Heinämäki, Orvokki. 1974. *Semantics of English Temporal Connectives*. Ph.D. thesis, University of Texas, Austin.

Hinrichs, Erhard. 1985. *A Compositional Semantics for Aktionsarten and NP reference in English*. Ph.D. thesis, Ohio State University.

Hinrichs, Erhard. 1986. "Temporal Anaphora in Discourses of English." *Linguistics and Philosophy*, 9, 63–82.

Hintikka, Jaako. 1962. *Knowledge and Belief*. Ithaca NY: Cornell University Press.

Hitzeman, Janet. 1993. *Temporal Adverbials and the Syntax Semantics Interface*. Ph.D. thesis, University of Rochester.

Hitzeman, Janet. 1997. "Semantic Partition and Temporal Adverbials." *Natural Language Semantics*, 5, 87–100.

Hoare, C. Anthony. 1969. "An Axiomatic Basis for Computer Programming." *Communications of the Association for Computing Machinery*, 30, 576–583.

Hornstein, N. 1990. *As Time Goes By: Tense and Universal Grammar*. Cambridge MA: MIT Press.

Hornstein, Norbert. 1977. "Towards a theory of tense." *Linguistic Inquiry*, 8, 521–557.

Horty, John, and Martha Pollack. 2001. "Evaluating New Options in the Context of Existing Plans." *Artifi cial Intelligence*, 127, 199–220.

Hughes, George, and Max Cresswell. 1968. *Introduction to Modal Logic*. Methuen.

Hwang, Chung Hee, and Lenhart Schubert. 1992. "Tense Trees as the Fine Structure of Discourse." In *Proceedings of the 30th Annual Meeting of the Asssociation for Computational Linguistics*, 232–240.

Inoue, Kyoko. 1979. "An Analysis of the English Present Perfect." *Linguistics*, 17, 561–590.

Isard, Stephen. 1974. "What would you have done if . . . ." *Theoretical Linguistics*, 1, 233–255.

Isard, Stephen, and H.Christopher Longuet-Higgins. 1973. "Modal Tictacto." In Radu Bogdan and Ilkka Niiniluoto, eds., *Logic, Language and Probabaility*, 189–195. Dordrecht: Reidel.

Jackendoff, Ray. 1991. "Parts and Boundaries." *Cognition*, 41, 9–46.

Jackson, Peter, and Han Reichgelt. 1987. "A General Proof Method for First-Order Modal Logic." In *Proceedings of the International Joint Conference on AI*, 942–944.

Japaridze, Giorgi. 1997. "A Constructive Game Semantics for the Language of Linear Logic." *Annals of Pure and Applied Logic*, 85, 87–156.

Japaridze, Giorgi. 1998. *The Logic of Resources and Tasks*. Ph.D. thesis, University of Pennsylvania.

Japaridze, Giorgi. 2002. "The Logic of Tasks." *Annals of Pure and Applied Logic*, 117, 261–293.

Kamp, Hans. 1979. "Events, Instants, and Temporal Reference." In Rainer Bäuerle, Urs Egli, and Arnim von Stechow, eds., *Semantics for Different Points of View*. Berlin: Springer Verlag.

Kamp, Hans, and Uwe Reyle. 1993. *From Discourse to Logic*. Dordrecht: Kluwer.

Kamp, Hans, and Christian Rohrer. 1983. "Tense in Texts." In Rainer Bäuerle, C. Schwarze, and Arnim von Stechow, eds., *Meaning, Use, and Interpretation in Language*, 250–269. Berlin: de Gruyter.

Karlin, R. 1988. "Defining the Semantics of Verbal Modifiers in the Domain of Cooking Tasks." In *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, 61–67.

Köhler, Wolfgang. 1925. *The Mentality of Apes*. New York: Harcourt Brace and World.

Kowalski, Robert. 1979. *Logic for Problem Solving*. Amsterdam: North Holland.

Kowalski, Robert, and Fariba Sadri. 1994. "The Situation Calculus and Event Calculus Compared." In *Proceedings of the 1994 International Symposium on Logic Programming*, 539–553.

Kowalski, Robert, and Maurice Sergot. 1986. "A Logic-Based Calculus of Events." *New Generation Computing*, 4, 67–95.

Kratzer, Angelika. 1977. "What *Must* and *Can* Can and Must Mean." *Linguistics and Philosophy*, 1, 337–355.

Kratzer, Angelika. 1991. “The Representation of Focus.” In Arnim von Stechow and Dieter Wunderlich, eds., *Semantics: An International Handbook of Contemporary Research*, 825–834. Berlin: de Gruyter.

Krifka, Manfred. 1989. “Nominal Reference, Temporal Constitution, and Quantification in Event Semantics.” In Renate Bartsch, Johan van Benthem, and Peter van Emde Boas, eds., *Semantics and Contextual Expressions*, 75–115. Dordrecht: Foris.

Krifka, Manfred. 1990. “Four Thousand Ships Passed Through the Lock: Object-Induced Measure Functions on Events.” *Linguistics and Philosophy*, 13, 487–520.

Krifka, Manfred. 1992. “Thematic Relations as Links between Nominal Reference and Temporal Constitution.” In Ivan Sag and Anna Szabolcsi, eds., *Lexical Matters*. Chicago: CSLI/Chicago University Press.

Kripke, Saul. 1965. “Semantical Analysis of Intuitionistic Logic, I.” In John Crossley and Michael Dummett, eds., *Formal Systems and Recursive Functions*. Amsterdam: North Holland.

Kripke, Saul. 1972. “Naming and Necessity.” In Donald Davidson and Gilbert Harman, eds., *Semantics of Natural Language*. Dordrecht: Reidel.

Laird, John, Allen Newell, and Paul Rosenbloom. 1987. “SOAR: an Architecture for General Intelligence.” *Artifi cial Intelligence*, 33, 1–64.

van Lambalgen, Michiel, and Fritz Hamm. 2005. *The Proper Treatment of Events*. Oxford: Blackwell.

Landman, Fred. 1986. *Towards a Theory of Information*. Dordrecht: Foris.

Landman, Fred. 1991. *Structures for Semantics*. Dordrecht: Kluwer.

Landman, Fred. 1992. “The Progressive.” *Natural Language Semantics*, 1, 1–32.

Lansky, Amy. 1986. “A Representation of Parallel Activity Based on Events, Structure and Casusality.” In *Proceedings of the Workshop on Planning and Reasoning about Action*, 50–86. Timberline Lodge, Mount Hood OR.

Lascarides, Alex, and Nicholas Asher. 1993a. “A Semantics and Pragmatics for the Pluperfect.” In *Proceedings of the Annual Meeting of the European Chapter of the Association for Computational Linguistics*, 250–259. Utrecht NL.

Lascarides, Alex, and Nicholas Asher. 1993b. “Temporal Interpretation, Discourse Relations, and Commonsense Entailment.” *Linguistics and Philosophy*, 16, 437–494.

Lehman, D., and S. Shelah. 1983. “Reasoning with Time and Chance.” In *Proceedings of the 10th International Colloquium on Automata, Languages and Programming*, vol. 154 of *Lecture Notes in Computer Science*. Berlin: Springer.

Levesque, Hector. 1988. “Logic and the Complexity of Reasoning.” *Journal of Philosophical Logic*, 17, 355–389.

Lewis, David. 1971. “Completeness and Decidability of Three Logics of Counterfactual Conditionals.” *Theoria*, 37, 74–85.

Lewis, David. 1973. “Causation.” *Journal of Philosophy*, 70, 556–567.

Lewis, David. 1979. “Scorekeeping in a Language Game.” *Journal of Philosophical Logic*, 8, 339–359.

Lin, Fangzhen. 1995. "Embracing Causality in Specifying the Indirect Effect of Actions." In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, 1985–1991. Montreal.

Lin, Fangzhen, and Raymond Reiter. 1995. "How to Progress a Database II: The STRIPS connection." In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, 2001–2007. Montreal.

Link, Godehard. 1983. "The Logical Analysis of Plurals and Mass Terms." In Rainer Bäuerle, C. Schwarze, and Arnim von Stechow, eds., *Meaning, Use, and Interpretation in Language*, 302–323. Berlin: de Gruyter. Reprinted in Link 1998.

Link, Godehard. 1987. "Algebraic Semantics of Event Structures." In Jeroen Groenendijk, Martin Stokhof, and Frank Veltman, eds., *Proceedings of the 6th Amsterdam Colloquium*. Reprinted in Link 1998.

Link, Godehard. 1998. *Algebraic Semantics in Language and Philosophy*. Stanford CA: CSLI Publications.

Lyons, John. 1977. *Semantics, Vol. II*. Cambridge: Cambridge University Press.

McCain, Norman, and Hudson Turner. 1997. "Causal Theories of Action and Change." In *Proceedings of the National Conference on Artificial Intelligence*, 460–465. Providence, RI: AAAI.

McCarthy, John. 1977. "Epistemological Problems of Artificial Intelligence." In *Proceedings of the 5th International Joint Conference on Artificial Intelligence*, 1038–1044.

McCarthy, John, and Patrick Hayes. 1969. "Some Philosophical Problems from the Standpoint of Artificial Intelligence." In Bernard Meltzer and Donald Michie, eds., *Machine Intelligence*, vol. 4, 473–502. Edinburgh: Edinburgh University Press.

McCarty, Thorne. 1994. "Modalities over Actions I: Model Theory." In *Proceedings of the 4th International Conference on Principles of Knowledge Representation and Reasoning*, 437–448. Bonn.

McCawley, James. 1971. "Tense and Time Reference in English." In Charles Fillmore and Terence Langendoen, eds., *Studies in Linguistic Semantics*, 96–113. New York: Holt, Rinehart, and Winston.

McDermott, Drew. 1982. "A Temporal Logic for Reasoning about Processes and Actions." *Cognitive Science*, 6, 101–155.

ter Meulen, Alice. 1984. "Events, Quantities, and Individuals." In Fred Landman and Frank Veltman, eds., *Varieties of Formal Semantics*. Dordrecht: Foris.

ter Meulen, Alice. 1986. "Locating Events." In Jeroen Groenendijk, Dirk de Jonge, and Martin Stokhof, eds., *Foundations of Pragmatics and Lexical Semantics*. Dordrecht: Foris.

ter Meulen, Alice. 1995. *Representing Time in Natural Language: the Dynamic Interpretation of Tense and Aspect*. Cambridge MA: MIT Press.

Miller, George, E. Galanter, and Karl Pribram. 1960. *Plans and the Structure of Behavior*. London: Holt.

Miller, George, and Phillip Johnson-Laird. 1976. *Language and Perception*. Cambridge: Cambridge University Press.

Mittwoch, Anita. 1995. "The English Perfect, Past Perfect, and Future Perfect in a Neo-Reichenbachian Framework." In Pier Marco Bertinetto, Valentina Bianchi, Östen Dahl, and Mario Squartini, eds., *Temporal Reference, Aspect, and Actionality*, 255–267. Torino: Rosenberg and Sellier.

Moens, Marc, and Mark Steedman. 1987. "Temporal Ontology in Natural Language,." In *Proceedings of the 25th Annual Meeting of the Association for Computational Linguistics, Stanford CA, July*, 1–7.

Moens, Marc, and Mark Steedman. 1988. "Temporal Ontology and Temporal Reference." *Computational Linguistics*, 14, 15–28.

Moore, Robert. 1980. *Reasoning about Knowledge and Action*. Ph.D. thesis, MIT, Cambridge MA. Published as TN-191, Menlo Park CA, SRI International.

Moore, Robert. 1985. "A Formal Theory of Knowledge and Action." In Jerry Hobbs and Robert Moore, eds., *Formal Theories of the Commonsense World*, 319–358. Norwood, NJ: Ablex. Reprinted as Ch. 3 of Moore (1995).

Moore, Robert. 1995. *Logic and Representation*, vol. 39 of *CSLI Lecture Notes*. Stanford CA: CSLI/Cambridge University Press.

Morgenstern, Leora. 1988. *Foundations of a Logic of Knowledge, Action, and Communication*. Ph.D. thesis, NYU, Courant Institute of Mathematical Sciences.

Morgenstern, Leora. 1994. "A Proper Ontology for Reasoning About Knowledge and Planning." In *Proceedings of the 1st International Workshop on Temporal Representation and Reasoning*, 69–76. Pensacola FL.

Morgenstern, Leora., and Lynn Stein. 1988. "Why Things Go Wrong: a Formal Theory of Causal Reasoning." In *Proceedings of the 7th National Conference on Artifi cial Intelligence*, 518–523. AAAI.

Mourelatos, Alexander. 1978. "Events, Processes and States." *Linguistics and Philosophy*, 2, 415–434.

Muskens, Reinhard. 1995. "Tense and the Logic of Change." In Urs Egli, P. Pause, C. Schwarze, A. von Stechow, and G. Wienold, eds., *Lexical Knowledge in the Organization of Language*, 147–183. Amsterdam: John Benjamins.

Muskens, Reinhard, Johan van Benthem, and Albert Visser. 1997. "Dynamics." In J. van Benthem and Alice ter Meulen, eds., *Handbook of Logic and Language*, 587–648. Amsterdam: North Holland.

Naumann, Ralf. 1995. *Aspectual Composition and Dynamic Logic*. Ph.D. thesis, University of Düsseldorf. Habilitationsschrift.

Naumann, Ralf. 1997. "A Dynamic Temporal Logic for Aspectual Phenomena in Natural Language." In *Proceedings of the 2nd International Conference on Temporal Logic*, 175–199. Manchester.

Naumann, Ralph. 2001. "Aspects of Changes: A Dynamic Event Semantics." *Journal of Semantics*, 18, 27–81.

Nelken, Rani, and Nissim Francez. 1995. "Splitting the Reference Time in Temporal Anaphora and Quantification in DRT." In *Proceedings of the 7th Conference of the European Chapter of the Association for Computational Linguistics*. Dublin: ACL.

Nerbonne, John. 1984. *German Temporal Semantics: Three-dimensional Tense-logic and a GPSG fragment*. Ph.D. thesis, Ohio State University, Comulmbus OH.

Nerode, Anil, and Duminda Wijesekera. 1990. "Constructive Concurrent Dynamic logic I." Tech. Rep. 90-43, Mathematical Sciences Institute, Cornell University, Ithaca.

Nilsson, Nils. 1980. *Principles of Artifi cial Intelligence*. Palo Alto CA: Tioga.

Ohlbach, Hans-Jürgen. 1991. "Semantics-based Translation Methods for Modal Logics." *Journal of Logic and Computation*, 1, 691–746.

Ortiz, Charles. 1994. "Causal Pathways of Rational Action." In *Proceedings of the Twelfth National Conference on Artifi cial Intelligence*, 1061–1066.

Ortiz, Charles. 1999a. "A Commonsense Language for Reasoning about Causation and Rational Action." *Artifi cial Intelligence*, 111, 73–130.

Ortiz, Charles. 1999b. "Explanatory Update Theory: Applications of Counterfactual Reasoning to Causation." *Artifi cial Intelligence*, 108, 125–178.

Palmer, Martha, Rebecca Passonneau, Carl Weir, and Tim Finin. 1993. "The KERNEL Text Understanding System." *Artifi cial Intelligence*, 63, 17–68.

Parsons, Terence. 1990. *Events in the Semantics of English*. Cambridge MA: MIT Press.

Partee, Barbara. 1973. "Some Structural Analogies between Tenses and Pronouns in English." *Journal of Philosophy*, 70, 601–609.

Partee, Barbara. 1984. "Nominal and Temporal Anaphora." *Linguistics and Philosophy*, 7, 243–286.

Pednault, Edward. 1989. "ADL: Exploring the Middle Ground between STRIPS and the Situation Calculus." In *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning*, 324–332. Palo Alto CA: Morgan Kaufmann.

Peleg, David. 1987. "Concurrent Dynamic Logic." *Journal of the Association for Computing Machinery*, 34, 450–479.

Petri, Carl Adam. 1962. "Kommunikation mit Automaten." Schriften des IMM 2, Institut für Instrumentelle Mathematik, Bonn.

Pratt, Vaughan. 1979. "Process Logic." In *Proceedings of the 6th Annual ACM Conference on Principles of Programming Languages*, 93–100.

Pustejovsky, James. 1991. "The Syntax of Event Structure." *Cognition*, 41, 7–82.

Quine, Willard van Ormond. 1960. *Word and Object*. Cambridge MA: MIT Press.

Reichenbach, Hans. 1947. *Elements of Symbolic Logic*. Berkeley, CA: University of California Press.

Reiter, Raymond. 1978. "On Closed World Databases." In Hervé Gallaire and Jack Minker, eds., *Logic and Databases*, 119–140. New York: Plenum.

Reiter, Raymond. 1991. "The Frame Problem in the Situation Calculus: a Simple Solution (Sometimes) and a Completeness Result for Goal Regression." In Vladimir Lifshitz, ed., *AI and Mathematical Theory of Computation: Papers in Honour of John McCarthy*, 359–380. New York: Academic Press.

Reiter, Raymond. 1993. "Proving Properties of States in the Situation Calculus." *Artificial Intelligence*, 64, 337–351.

Reiter, Raymond. 2001. *Knowledge in Action*. Cambridge MA: MIT Press.

Rescher, Nicholas. 1964. *Hypothetical Reasoning*. Amsterdam: North Holland.

Richards, Barry. 1982. "Tense, Aspect, and Time Adverbials." *Linguistics and Philosophy*, 5, 59–107.

Ritchie, Graeme. 1979. "Temporal Clauses in English." *Theoretical Linguistics*, 6, 87–115.

Robert, Kowalski, and Fariba Sadri. 1997. "Reconciling the Event Calculus with the Situation Calculus." *Journal of Logic Programming*, 31, 39–58.

Roberts, Craige. 1989. "Modal Subordination and Pronominal Anaphora in Discourse." *Linguistics and Philosophy*, 12, 683–721.

Rosenschein, Stanley. 1981. "Plan Synthesis: a Logical Perspective." In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, 331–337. Vancouver.

Sandewall, Eric. 1994. *Features and Fluents*. Oxford: Oxford University Press.

Schank, Roger. 1975. "The Structure of Episodes in Memory." In Daniel Bobrow and Allan Collins, eds., *Representation and Understanding*, 237–272. New York: Academic Press.

Schein, Barry. 1993. *Plurals and Events*. Cambridge, MA: MIT Press.

Schein, Barry. 2003. "Events and the Semantic Content of Thematic Relations." In Gerhard Preyer and Georg Peter, eds., *Logical Form, Language and Ontology: On Contemporary Developments in the Philosophy of Language and Linguistics*. Oxford: Oxford University Press.

Scherl, Richard, and Hector Levesque. 1993. "The Frame Problem and Knowledge-Producing Actions." In *Proceedings of the 11th National Conference on Artificial Intelligence*, 689–695. Washington: AAAI.

Schubert, Lenhart. 1990. "Monotonic Solution of the Frame Problem in the Situation Calculus: an Efficient Method for Worlds with Fully Specified Actions." In Henry Kyburg, Ronald Loui, and Gregory Carlson, eds., *Knowledge Representation and Defeasible Reasoning*, 23–67. Dortrecht: Kluwer.

Schubert, Lenhart. 1994. "Explanation closure, action closure, and the Sandewall test suite for reasoning about change." *Journal of Logic and Computation*, 4, 679–700.

Shanahan, Murray. 1997. *Solving the Frame Problem*. Cambridge, MA: MIT Press.

Shanahan, Murray. 1997b. "Event Calculus Planning Revisited." In *Proceedings of the Workshop on Robots, Softbots, Immobots: Theories of Action, Planning, and Control*, 96–104. Held in conjunction with the 14th National Conference on Artificial Intelligence, Providence, RI, July 1997.

Shoham, Yoav. 1988. *Reasoning about Change*. Cambridge MA: MIT Press.

Shoham, Yoav., and Nita Goyal. 1988. "Temporal reasoning in AI." In Howard Shrobe, ed., *Exploring Artificial Intelligence*, 419–438. San Mateo, CA: Morgan Kaufmann.

Smith, Carlota. 1991. *The Parameter of Aspect*. Dordrecht: Reidel.

Smullyan, Raymond. 1973. "A Generalization of Intuitionistic and Modal Logics." In Hugues Leblanc, ed., *Truth, Syntax, and Modality*, 274–293. Amsterdam: North Holland.

Song, Fei., and Robin Cohen. 1988. "The Interpretation of Temporal Relations in Narrative." In *Proceedings of the 7th National Conference of the American Association for Artifi cial Intelligence*, 745–750.

Spejewski, Beverly. 1994. *Temporal Subordination in Discourse*. Ph.D. thesis, University of Rochester.

Stalnaker, Robert. 1968. "A Theory of Conditionals." In Nicholas Rescher, ed., *Studies in Logical Theory*. Oxford: Blackwell.

Stalnaker, Robert. 1984. *Inquiry*. Cambridge MA: MIT Press.

Steedman, Mark. 1977. "Verbs, Time and Modality." *Cognitive Science*, 1, 216–234.

Steedman, Mark. 1982. "Reference to past time." In Robert Jarvella and Wolfgang Klein, eds., *Speech, Place, and Action*, 125–157. New York: Wiley.

Steedman, Mark. 1995. "Dynamic Semantics for Tense and Aspect." In *Proceedings of the International Joint Conference on Artifi cial Intelligence, Montreal, Quebec*, 1292–1298. San Francisco, CA: Morgan Kaufmann.

Steedman, Mark. 1997. "Temporality." In Johan van Benthem and Alice ter Meulen, eds., *Handbook of Logic and Language*, 895–938. Amsterdam: North Holland/Elsevier.

Steedman, Mark. 2002a. "Formalizing Affordance." In *Proceedings of the 24th Annual Meeting of the Cognitive Science Society, Fairfax VA, August*, 834–839. Mahwah NJ: Lawrence Erlbaum.

Steedman, Mark. 2002b. "Plans, Affordances, and Combinatory Grammar." *Linguistics and Philosophy*, 25, 723–753.

Steedman, Mark. 2004. "Towards a Compositional Semantics for English Intonation." URL `http://www.cogsci.ed.ac.uk/~steedman/papers.html`. Ms. University of Edinburgh.

Stein, Lynn. 1991. *Resolving Ambiguity in Non-monotonic Reasoning*. Ph.D. thesis, Brown University, Providence RI.

Stone, Matthew. 1994. "The Reference Argument of Epistemic "Must"." In *Proceedings of the International Workshop on Computational Semantics*, 181–190. Tilburg University.

Stone, Matthew. 1998. *Modality in Dialogue: Planning Pragmatics and Computation*. Ph.D. thesis, University of Pennsylvania.

Stone, Matthew. 1999. "Representing Scope in Intuitionistic Deduction." *Theoretical Computer Science*, 129–188.

Stone, Matthew, and Daniel Hardt. 1997. "Dynamic Discourse Referents for Tense and Modals." In Harry Bunt, ed., *Proceedings of the International Workshop on Computational Semantics*, 287–299. Dordrecht: Kluwer.

de Swart, Henriëtte. 1998. "Aspect shift and coercion." *Natural Language and Linguistic Theory*, 16, 347–385.

Taylor, Barry. 1977. "Tense and Continuity." *Linguistics and Philosophy*, 1, 199–220.

Thomason, Richmond. 1970. "Indeterminist Time and Truth-Value Gaps." *Theoria*, 36, 246–281.

Thomason, Richmond. 1981. "Deontic Logic as Founded in Tense Logic." In Risto Hilpinen, ed., *New Studies in Deontic Logic*. Dordrecht: Reidel.

Thomason, Richmond. 1991. "Logicism, AI, and Common Sense: John McCarthy's Program in Philosophical Perspective." In Vladimir Lipschitz, ed., *AI and Mathematical Theory of Computation: Papers in Honour of John McCarthy*, 449–466. New York: Academic Press.

Thomason, Richmond. 1997. "Nonmonotonicity in Linguistics." In Johan van Benthem and Alice ter Meulen, eds., *Handbook of Logic and Language*, 777–831. Amsterdam: North Holland/Elsevier.

Thomason, Richmond. 1999. "Ability, Action and Context." In *Proceedings of the Workshop on Temporality and Discourse Context: Dynamic and Modal Approaches*. URL `http://www.cs.tcd.ie/Tim.Fernando/postTemp.html`.

Thomason, Richmond, and Anil Gupta. 1980. "A Theory of Conditionals in the Context of Branching Time." *Philosophical Review*, 88, 65–90.

Tichy, Pavel. 1985. "Do We Need Interval Semantics?" *Linguistics and Philosophy*, 8, 263–282.

Turner, Raymond. 1981. "Counterfactuals without Possible Worlds." *Journal of Philosophical Logic*, 10, 453–493.

Veloso, Manuela, Jaime Carbonell, Alicia Pérez, Daniel Borrajo, Eugene Fink, and Jim Blythe. 1995. "Integrating Planning and Learning: The PRODIGY Architecture." *Journal of Experimental and Theoretical Artificial Intelligence*, 7, 81–120.

Veltman, Frank. 1984. "Data Semantics." In Theo M. V. Janssen Jeroen Groenendijk and Martin Stokhof, eds., *Truth, Interpretation, and Information*, 43–63. Dordrecht: Foris.

Vendler, Zeno. 1967. *Linguistics in Philosophy*. Ithaca, NY: Cornell University Press.

Verkuyl, Henk. 1972. *On the Compositional Nature of the Aspects*, vol. 15 of *Foundations of Language Supplementary Series*. Dordrecht: Reidel.

Verkuyl, Henk. 1989. "Aspectual Classes and Aspectual Composition." *Linguistics and Philosophy*, 12, 39–94.

Vlach, Frank. 1981. "The Semantics of the Progressive." In Philip Tedeschi and Annie Zaenen, eds., *Tense and Aspect*, vol. 14 of *Syntax and Semantics*. New York: Academic Press.

Wallen, Lincoln. 1990. *Automated Proof Search in Non-Classical Logics: Efficient Matrix Proof Methods for Modal and Intuitionistic Logics*. Cambridge MA: MIT Press.

Webber, Bonnie Lynn. 1978. "On Deriving Aspectual Sense: a Reply to Steedman." *Cognitive Science*, 2, 385–390.

Webber, Bonnie Lynn. 1983. "Logic and Natural Language." *IEEE Computer*, 43–46. Special Issue on Knowledge Representation, October 1983.

Webber, Bonnie Lynn. 1988. "Tense as Discourse Anaphor." *Computational Linguistics*, 14, 61–73.

Webber, Bonnie Lynn, and Nils Nilsson. 1981. *Readings in Artificial Intelligence*. Palo Alto CA: Tioga.

White, Michael. 1993. "Delimitedness and Trajectory of Motion Events." In *Proceedings of the 6th Conference of the European Chapter of the Association for Computational Linguistics*, 412–421. Utrecht.

White, Michael. 1994. *A Computational Approach to Aspectual Composition*. Ph.D. thesis, University of Pennsylvania.

Wilensky, Robert. 1983. *Planning and Understanding*. Reading MA: Addison-Wesley.

Winograd, Terry. 1972. *Understanding Natural Language*. Edinburgh: Edinburgh University Press.

von Wright, Georg. 1964. *Norm and Action: a Logical Enquiry*. Routledge.

von Wright, Georg. 1967. "The Logic of Action—a Sketch." In Nicholas Rescher, ed., *The Logic of Decision and Action*. Pittsburgh: University of Pittsburgh Press.

Young, Robert, and William Sr. Morgan. 1987. *The Navaho Language*. Albuquerque, NM: University of New Mexico Press.

Zwarts, Joost, and Henk Verkuyl. 1994. "An Algebra of Conceptual Structure: an Investigation into Jackendoff's Conceptual Semantics." *Linguistics and Philosophy*, 17, 1–28.