

PART I

Tense, Aspect, and Event Structure



Introduction to Part I

I INTRODUCTION

The science of computational linguistics requires a variety of tools, both theoretical and practical. In order to process temporal information in natural language, we need core theories from linguistics, computational approaches to interpreting tense, aspect, and discourse structure, as well as temporal reasoning capabilities. The papers in this part cover the core linguistic theories dealing with tense and aspect, as well as computational approaches to interpreting them. To assist the reader in navigating through the various theories, we will first provide some relevant linguistic background, explaining the variety of terminology used in these papers. We then summarize and situate each of the papers in terms of this background.¹

In order to survive in the world, our ancestors at some point evolved the ability to reason in relative terms like ‘near’ and ‘far’ about how far away a source of food was. Likewise, they also evolved the ability to reason about past events and planned or hypothetical events, and thus to reason based on how far away an event was in time. As a result of evolution, by a process which is still largely mysterious, all these kinds of reasoning got built into distinctions found in the languages of the world. In particular, **tense** is a specific mechanism built into language for locating information in time.

Verbs and other expressions in natural languages convey meaning about the temporal properties of events, such as whether events last, change, or complete. **Lexical aspect** or *Aktionsarten* (German for ‘kinds of action’), distinguishes between different subclasses of events based on these temporal properties. Languages also provide a variety of resources to distinguish between external versus internal viewpoints on an event. In particular, **grammatical aspect** expresses whether an event is viewed as finished, completed, or ongoing. Finally, the notion of **event structure** is a representation of events as complex objects with specific components. Taken together, aspect and event structure provide what, in computational terms, is called an **event ontology**: a theory of the different subclasses (related by what are sometimes called, in the AI literature, **isa** links), components (related by what are sometimes called **part-of** links), and temporal properties of events. The event ontologies discussed here are motivated by the representational distinctions found in natural language, rather than those motivated by nonlinguistic reasoning tasks. This is appropriate, since the ontologies are used for processing natural language.

Sections 4 and 5 below make use of logical concepts and notation, as do subsequent parts of this book. The reader who lacks an introductory background in logic may benefit from a text such as Suppes (1999).

Tense can be defined as ‘the grammaticalized expression of location in time’ (Comrie 1986: 9). This grammaticalized expression involves marking, via change of form, of particular syntactic elements, e.g. the verb and auxiliaries. For example, in *John ran a marathon*, the past tense morpheme represented as *-ed* (producing the inflected verb form *ran*) is used to indicate that the event occurred at a time earlier than the speech time. In *John will run a marathon*, the modal auxiliary *will* is used to locate the event as occurring at a future time, i.e. latter than the speech time. While tense is mainly marked on the verb and auxiliaries associated with the verb group, in some languages, like the North American Indian language Nootka (Comrie 1986), tense is expressed on the noun phrase.

Tense isn’t the only mechanism for expressing location in time. In languages such as Mandarin Chinese, which lacks tense morphemes, aspectual markers (see Section 3.2 below) can be used to express location in time, though sometimes even these may be absent (Lin 2003). There are also nongrammaticalized expressions of location in time given by temporal adverbials, e.g. *tomorrow*, *yesterday*, *two hours later*, etc. In the case of *tomorrow* or *yesterday*, the temporal location is with respect to the speech time. Temporal locations can also of course be expressed relative to a coordinate system given by a calendar, e.g. 1991 (*AD*), or a cyclically occurring event, e.g. *morning*, *spring*, or an arbitrary event, e.g. *the day after he married her*.

The few languages that lack tense altogether are not able distinguish past from present or future. However, they all have a **realis/irrealis** distinction. In Burmese, for example (Comrie 1986), events that are ongoing or that were observed in the past are expressed by sentence-final realis particles *-te*, *-tha*, *-ta*, and *-hta*. In other cases, i.e. for unreal or hypothetical events (including future events, present events, and hypothetical past events), the sentence-final irrealis particles *-me*, *-ma*, and *-hma* are used.

We now turn to the meanings of individual tenses, as found in English. In English, the **present tense** usually locates events as occurring roughly at the speech time. However, the use of the present does not necessarily mean the event coincides with speech point. In fact, except for performatives like *I warn you . . .*, the event doesn’t usually coincide with the speech point. The typical use of the English present is the reporting use (as in sports commentaries), and in stative sentences (see below). The present tense can also be used to locate events as occurring in the past, as in *Like, then he tells me . . .* It can also be used to locate events in the future, as in *I leave next week*, as also to express a habitual use indicating a generalization over a set of times, as in *John loves sushi*.

The **past tense** usually refers to a time prior to speech time. Thus, *John slapped Bill* means that there is a time prior to the speech time when the slapping event occurred. The past tense can also involve **definiteness**, i.e. the speaker has a particular time in mind, as in the example of Partee (1973) (*Oops*,) *I didn’t turn off the stove* (uttered halfway down the turnpike). Note that this reading is stronger than the indefinite reading (there was some unspecified time when the stove wasn’t turned off), and weaker than the reading where there is no earlier time when I turned off the stove.

The **future tense** usually refers to a time after the speech time though, like the other tenses, it can also be used to refer to other times, e.g. *He’ll be home by now*. Our concept of the future is not symmetric with the past, since the future involves branching. In most Germanic languages, the expression of the future tense always seems to involve the non-tense feature of modality, e.g. *I will/may go tomorrow*. See Comrie (1986) for extensive discussions of tense across languages.

3 ASPECT

3.1 Lexical Aspect

3.1.1 Vendler

VENDLER's classic paper (Chapter 1) groups verbs into various subclasses based on their temporal properties. He notes that verbs which describe **activities** like running, working, etc., express actions that 'consist of successive phases following each other in time.' As a result, it is natural to express events by means of a 'continuous tense', i.e. a verb in the progressive form (*John is running*). Vendler characterizes verbs that describe activities as **processes**. By contrast, **states** do not involve successive phases, as a result, they sound odd in the progressive form, e.g. **John is knowing* (the oddness indicated here with an asterisk). Vendler also observes that 'while running or pushing a cart has no set terminal point, running a mile and drawing a circle do have a "climax".' He points out that a man who stops running did run, while a man who stops running a mile didn't run a mile. He argues that running for half an hour involves running for every subperiod within that half hour, whereas having run a mile in four minutes precludes having run a mile in any subperiod. Thus processes are distinguished from a further class of events that culminate, called **accomplishments**. Processes allow adverbials with 'for' but sound odd with 'in', as in *pushing a cart for/*in half an hour*. Accomplishments have the opposite behavior, e.g. *draw a circle in*/for twenty seconds*. Vendler then goes on to distinguish the class of **achievements**, namely events like reaching a hilltop or winning a race that can be predicated for single moments of time. Since achievements don't extend over time, they can't in general co-occur with 'for' adverbials.

The distinctions Vendler draws and subsequently illustrates with a variety of natural language verb usage examples, are based on solid arguments. His work has stimulated research that has been applied across many languages; see (Binnick 2002) for a comprehensive bibliography on lexical aspect. Over the years, a number of refinements have been made to Vendler's scheme, both in terms of the lexical and syntactic categories involved, their linguistic tests, and their semantic formalization. These are discussed next.

3.1.2 Refinements

Overall, we can characterize a general class of **eventualities** (also called **situations**), that are either **events** or **states** (also called **statives**), with events being subdivided into **processes** (also called **activities**), **accomplishments**, and **achievements**. We will now describe the linguistic tests for each. It is worth bearing in mind, though, that linguistic tests can vary in reliability. While they provide a well-established form of linguistic argumentation for drawing distinctions between categories, they rely on all-or-nothing judgments of ungrammaticality or awkwardness of sentences; a more graded approach to such judgments may be more realistic. Finally, the tests do have exceptions, but usually such exceptions can be explained (e.g. see the discussion of the *for*-adverbial test for achievements below).

Statives are expressed by verbs like *have, love, believe, know, realize, sit, and stand*, etc., but also by adjectives with a copula, as in *is clever, is happy*, and nouns like *killing, accident*, etc. For any state p (like John's being hungry) that holds over a period t , p must also hold for every subinterval of t (Dowty, Chapter 15). This **subinterval property** is a characteristic of states. Statives either denote a situation or entry into the situation

(**ingressive** or **inceptive** readings, respectively). Thus, *John sits* is ambiguous: it might denote a state or entry into a state; *John knows* describes a state, whereas *John realizes* describes entry into a state; hence *John is (gradually) realizing what's happening* is acceptable, but **John is knowing what's happening* is odd. Some of the standard linguistic tests for stative readings are as follows:

1. Sentences describing statives (in noninceptive readings) are, as Vendler points out, odd when they occur in the progressive (**John is knowing Bill*).
2. Statives are odd in imperatives (**Know the answer*).
3. Statives can't be used in pseudo-clefts with *do* (**What John did was know the answer*).

Activities (associated with verbs like *walk*, *run*, etc.) differ from other eventualities in that if an activity *p* (like John's walking) occurs in period *t*, a part of the activity (also an activity) must occur for most subintervals of *t*. Note that here the subinterval property must be constrained so as to allow for gaps. For example, in contrast to Vendler's argument about running, I usually pause while running. Also, only parts of the activity down to a particular **grain size** count as the same activity; clearly, lifting a leg is not running. Thus, *X Vs for an hour*, where *V* is an activity, entails *X Vs for all/most times in that hour*; also *X is Ving* entails that *X has Ved*, since the latter sentence is true of a subactivity of *V*. Some of the standard tests are:

1. Activities take temporal adverbials with *for* (e.g. *John ran for an hour*).
2. Activities do not take temporal adverbial phrases with *in* (e.g. **John ran in an hour*).

Accomplishments (associated with verbs like *build*, *cook*, *book*, etc.) are eventualities which can **culminate**. Unlike activities, *X Vs for an hour*, where *V* is an accomplishment, does not entail *X Vs for all times in that hour*; likewise *X is Ving* does not entail that *X has Ved*. Thus, *John is cooking dinner* does not entail *John has cooked dinner*. Standard tests are:

1. Accomplishments take temporal adverbial phrases with *in* (e.g. *John booked a flight in an hour*).
2. Accomplishment verbs with *stop* indicate that the event did not succeed, e.g. *John stopped building a house* (compare with activities, where they indicate the event succeeded e.g. *John stopped running*).

Achievements (associated with verbs like *win*, *blink*, *find*, *reach*, etc.) are instantaneous (or short-duration) accomplishments, namely events that finish and that occur in a very short time period. The standard tests are:

1. Achievements cannot be modified by temporal *for*-adverbials. Thus, **X Vs for an hour*, where *V* is an achievement verb. For example, **John dies for an hour*, **John wins for an hour*, **John reached New York for an hour*. In cases where these sentences are acceptable, they tend to have readings which refer to the duration of the result state (e.g. of being dead or a winner) rather than the duration of the achievement event itself.
2. Achievement verbs do not take *stop* (e.g. **John stopped reaching New York*).

Accomplishment and achievements, which are events that can culminate, are sometimes called **telic** eventualities. Achievements are instantaneous, and events in general involve change from the state that held before the event began. Finally, there are also instantaneous activities, called **semelfactives**, like *knock* or *cough*, which are

TABLE I Aspectual Features

| | <i>Telic</i> | <i>Dynamic</i> | <i>Durative</i> | <i>Example</i> |
|-----------------------|--------------|----------------|-----------------|------------------------------|
| Stative | – | – | + | <i>Know, have</i> |
| Activity | – | + | + | <i>walk, paint</i> |
| Accomplishment | + | + | + | <i>build, destroy</i> |
| Achievement | + | + | – | <i>notice, win</i> |

instantaneous, atelic, and dynamic. However, they are usually not broken out as a separate class. Based on this, the feature analysis shown in Table I can be constructed to capture these distinctions, based on Smith (1991):

So far, these aspectual classes have been defined by linguistic criteria, but not formalized in terms of a semantic theory. Such a formalization is implicit in the work of Dorr and Olsen (Chapter 7), and is carried out explicitly in the paper by Dowty in Part III.

3.1.3 Aspectual Composition

While these eventuality subclasses characterize individual lexical items, they also can be used to characterize phrases and sentences. In fact, as Dowty points out in his paper in part III, there is a lack of consistency in the literature in terms of what linguistic element the classification applies to. Different approaches have considered different kinds of text units, including lexical items (e.g. *hungry*), verb phrases (e.g. *eating a peach*), sentences, and semantic entities like events. Despite these differences, however, it is clear that expressions of one class can be transformed into one of another class by combining with another expression. This type of transformation is called **aspectual composition**. For example, an activity can be changed into an accomplishment by adding an adverbial phrase expressing temporal or spatial extent, e.g. *I walked* (activity) and *I walked to the station/a mile/home* (accomplishment). Likewise, an accomplishment like *I built my house* can be changed into an activity by combining with a *for*-adverbial, viz., *I built my house for an hour* (activity). A system of rules for carrying out such aspectual composition is called an **aspect calculus**; examples are found in Dowty (1979), see also Dowty (Chapter 15), Jackendoff (1991), de Swart (1998) and Pustejovsky (Chapter 2), among others.

3.2 Grammatical Aspect

While tense allows the speaker to relate the time of an eventuality to a deictic center or some other reference point, grammatical aspect allows the speaker to represent the structure of an eventuality. Here there is a distinction between **perfective aspect**, where an entire eventuality is presented without its internal temporal structure, e.g. *John built a house*, and **imperfective aspect**, where the speaker represents internal phases of the eventuality, e.g. *John is building a house*. Perfective aspect can express termination or completion of an eventuality, while imperfective aspect can express the ongoing nature of an activity. It is important to realize that many of the traditional tenses, e.g. Spanish imperfective, as in *Juan leía cuando entre* (*John was reading when I entered*), may combine both tense and aspect, e.g. past and imperfective. The same is true of the ‘complex tenses’ in English, such as present progressive, present perfect, etc.

Grammatical aspect is expressed in systematic ways across languages, depending on the lexical aspect of the eventuality. The following account, derived from cross-linguistic arguments by Smith (1991) and work on Chinese by (Lin 2003), summarizes some of this systematicity.

In English and French, **perfective aspect** is signaled by verbal tense and aspect morphemes. Termination is expressed in activities, completion is expressed in accomplishments and achievements, and statives can either express termination (e.g. French *passé composé* tense morpheme) or not (English, e.g. *I have lived in Paris*). In Mandarin Chinese, which lacks tense markers but which does have the semantic notion of tense (Lin 2003), the perfective is signaled by morphemes *-le* and *-guo*, usually indicating termination for activities, accomplishments and achievements; completion is indicated by a separate resultative morpheme *-wan*. In Russian, the perfective doesn't apply to statives, but is signaled by prefixes *po-* (short duration) and *pro-* (unexpected interval) in activities.

The **imperfective aspect** is signaled in English by the progressive morpheme *-ing*. It occurs in activities, accomplishments, and achievements. In French, as in Russian, it is signaled by tense morphemes (e.g. the French *imparfait*). In Mandarin, it is signaled by the progressive morpheme *-zai* and resultative morpheme *-zhe*. The particle *-le* can also have an imperfective use with atelic predicates (Lin 2003).

4 EVENT STRUCTURE

As mentioned earlier, from an event ontology standpoint, the component structure of events is of considerable importance, in addition to their various subclasses. In order to represent the internal structure of events, one must first understand how meanings of event-denoting expressions are to be represented. In a logical framework, there are two fundamental ways of representing meaning: (i) the meaning of an accomplishment verb, say *close*, is the predicate *CLOSE*, so that a sentence of the form *X closes Y* is mapped to *CLOSE(X, Y)* (ii) the meaning of *close* is decomposed into the meaning of primitive elements, e.g. *CAUSE TO BECOME NOT OPEN*, so that a sentence like *X closes Y* can be represented as *CAUSE(X, BECOME(NOT(OPEN(Y))))*.

As Dowty (1979) shows, by decomposing the meaning of any accomplishment verb into a description of a causative event and a result state, we can explain a number of apparent scope ambiguities with adverbs, as well as some data from derivational morphology. For example, *John closed the door again* is ambiguous between an 'external' reading, where John performed the action of closing the door at least once before, and an 'internal' reading, where John merely brought it about that the door is again closed. Postulating a decomposed meaning, the external reading may be explained by *again* having (wide) scope over the entire expression, i.e. *AGAIN(CAUSE(X, BECOME(NOT(OPEN(Y)))))*, whereas in the internal reading, *again* has narrow scope over *CLOSED*, i.e. *CAUSE(X, BECOME(AGAIN(NOT(OPEN(Y)))))*.

Despite Dowty's plausible arguments in favor of his decompositional scheme, decompositional theories are prone to problems, the most crucial of which is the proliferation of primitives. In practical situations, the semantic primitives are often inadequate for the task at hand, and new primitives have to be added. Approaches such as Jackendoff (1990) are hindered by the inability to distinguish differences in meaning among similar words, for example, *run/walk/trot* are all mapped to a *MOVE* primitive. In dealing with this, approaches often end up proliferating 'manner'

primitives that help distinguish the different senses. (In addition, some decompositional accounts, such as Schank (1975) are further limited by the lack of a compositional semantics.)

Instead of decomposing meanings in this (somewhat ungainly) way, PUSTEJOVSKY (Chapter 2) suggests a more modular approach. In his approach, rather than the decomposition of event meanings into primitives, the substructure of events, together with meaning-forming rules for ‘event composition’ give rise to the relationships expressed by decomposed meanings. In particular, he argues that accomplishment events like *close* are made up of two aspectually salient parts, a preparatory process with the meaning $ACT(X, Y) \& NOT(CLOSED(Y))$ and a result state with meaning $CLOSED(Y)$. Note that primitives like *CAUSE* and *BECOME* are not explicitly part of the representation. The resulting state overlaps with the event, and can of course extend beyond it. Activities are made up of (an indeterminate number of) subevents each of which corresponds to the activity, e.g. *X pushed Y* would be made up of subevents *e* each represented as $ACT(X, Y)$ and $MOVE(Y)$. States have no such decomposition, so that *Y is closed* is simply represented as $CLOSED(Y)$. His treatment of achievements is a little less uniform, as it depends on whether it involves an agentive action or not. A non-agentive sentence like *X died* is represented in terms of a preparatory process with the meaning $NOT(DEAD(X))$ and a result state with meaning $DEAD(X)$, while an agentive achievement like *win* is represented just as accomplishments are.

A crucial point about event composition is that it is part and parcel of a compositional semantics. Thus, in *Mary ran to the store*, *to the store* has a meaning such that it combines with an activity like *ran* to yield a phrase *ran to the store* that is an accomplishment. The latter event has a structure with a preparatory process given by the activity $RAN(X)$ and the result state given by $AT(X, STORE)$. Pustejovsky goes on to apply his account to resultative constructions (like *Mary hammered the metal flat*, where the result state is directly expressed in the construction), and the syntactic behaviors of psychological verbs like *fear* and *frighten*, as well as problems involving the scopes of adverbial modifiers, and the mapping from argument structure to syntactic structure. While the theory has some explanatory advantages, it does not directly relate event structure to temporal structure. For example, the point of culmination of an accomplishment, at which the result is generated, isn’t explicitly identified.

The notion of event structure calls out for representational detail. Let us assume that the accomplishment of building a cabin involves a number of subevents as part of its preparatory process, including the accomplishment of hammering in a nail. What sort of precise structure do these events and processes have, given the nesting of subevents? One way of answering this question is by means of an **event algebra**, described by BACH in this part.

Bach’s work draws inspiration from the work of Link (1983) on a lattice-theoretic account of the semantics of mass, count, and plural nouns. In Link’s scheme, sets of individuals (as in the case of plurals like *John and Mary*) are individuated by using a join operation that creates super-individuals. Members of the set are linked by an ‘individual part-of’ relation to the set super-individual. Thus John is an individual part of the plural individual John and Mary, as also of John and the chair, etc. There is a similar structure having to do with the stuff that John (or the chair) is made up of: there is a join operation defined over bits of stuff, with the similar result that the stuff of John is part of the stuff of John and Mary. There is also a mapping between the two structures: the mapping from an individual to the stuff she is made out of can be extended to a mapping from joins of individuals to the joins of the stuff they are made out of.

Bach's goal is to have a similar mapping between telic events and the process stuff they are made out of. His event subclasses are the Vendlerian ones, except that he distinguishes, following Carlson (1977), between two subclasses of achievements, 'happenings' like *recognize* and 'culminations' like *die*. The algebra itself is very concise. (However, please note that there is a misprint in his paper in Section 4, line (2): it should say that the atomic events A_e are a subset of E_e .) Bach shows that the algebra can make precise a variety of intuitions. For example, the process of the event of Sally pounding in a nail is a process-part-of the process of the event of Sally building a cabin. The event of Mary stumbling is an individual part-of the plural event of Mary stumbling and Mary twisting her ankle, and the process of Mary stumbling is a process-part-of the process of Mary stumbling and twisting her ankle. Moreover, events which are distinct, such as John pouring poison into the water (deliberately to kill fish) and John poisoning the populace (which could be unintentional) could be made up of the same processes. The paper then goes on to speculate about other interesting parallels between the domain of nominals and the domain of events.

Further work on event structure applied to event and aspectual semantics is found in the lattice-theoretic work of Hinrichs (1985) and Krifka (1989).

5 FORMAL THEORIES OF TENSE AND ASPECT

5.1 Tense as Anaphor: Reichenbach

We have already seen that times of events can be located with respect to a deictic center, so in a sense they are like pronouns. This anaphoric view of tense derives from the classic work of HANS REICHENBACH, first published in 1947 (Chapter 4). Tensed utterances, according to Reichenbach, introduce references to three 'time points': the speech time S , the event time E , and the **reference time** R . Reference times, for Reichenbach, can coincide with the times of expressions like *Tuesday*, *the next week*, *now*, etc., but they aren't necessarily made explicit in text. Now, consider sentence (1). Here the event of leaving is located before the speech time, and also before the time of five o'clock, which coincides with the reference time, which in turn offers a vantage point from which the event of leaving is seen. This would be represented as $E < R < S$. Here $E < R$ means the event time **precedes** the reference time. Likewise, $E = R$ means the event time is **simultaneous** with the reference time.

- (1) By five o'clock, John had already left.

Likewise, in (2), the vantage point, which isn't expressed by a temporal adverbial, is some time in the future, so we have E and S before R , with the relation between E and S left unspecified.

- (2) John will have caught the train.

Reichenbach used the relations between S , E , and R to represent the meaning of tenses in English and a few other languages. Reichenbach's examples are mainly drawn from English, but in keeping with other well-educated scholars of his time, he also discussed examples in German, French, Latin, and Greek.

Reichenbach's constraint-based analysis of the English tense system is shown in Table 2. Tense (present, past, future) is determined by the relation between R and S . Aspect (simple, anterior, or posterior) is determined by relation between E and R .

TABLE 2 Comparison of Reichenbach and Prior

| <i>Relation</i> | <i>Reichenbach's Tense Name</i> | <i>Prior</i> | <i>English Tense Name</i> | <i>Example</i> |
|-----------------|---------------------------------|--------------|---------------------------|--|
| E<R<S | Anterior past | PP ϕ | Past perfect | <i>I had slept</i> |
| E = R<S | Simple past | P ϕ | Simple past | <i>I slept</i> |
| R<E<S | Posterior past | PF ϕ | | <i>[I expected that]</i> |
| R<S = E | | | | <i>I would sleep</i> |
| R<S<E | | | | |
| E<S = R | Anterior present | P ϕ | Present perfect | <i>I have slept</i> |
| S = R = E | Simple present | ϕ | Simple present | <i>I sleep</i> |
| S = R<E | Posterior present | F ϕ | Simple future | <i>I will sleep French: Je vais dormir</i> |
| S<E<R | Anterior future | FP ϕ | Future perfect | <i>I will have slept</i> |
| S = E<R | | | | |
| E<S<R | | | | |
| S<R = E | Simple Future | F ϕ | Simple future | <i>I will sleep French: Je dormirai</i> |
| S<R<E | Posterior future | FF ϕ | | <i>I shall be going to sleep</i> |

Only seven out of thirteen relations are realized in English; there are six different forms, with the simple future being ambiguous. However, as Comrie (1986) points out, the posterior past can be given an unambiguous representation $E>R<S$, and similarly the anterior future can be given an unambiguous representation as $E<R>S$. The progressive forms (not shown) are treated by Reichenbach as having extended intervals in time instead of time points, but are otherwise similar to the simple tenses.

Reichenbach's work has had an enduring influence both on subsequent work in formal semantics and computational linguistics. In the Introduction to Part III, we discuss the use of Reichenbachian notions in a formal representation of discourse meaning, namely Discourse Representation Theory (DRT) (Kamp and Reyle 1993). The notion of tense as anaphor is also explicated further in Webber's paper in Part III.

5.2 Tense as Operator: Prior

When discussing the semantics of the past tense, it was pointed out that *John slapped Bill* means that there is a time prior to the speech time when the slapping event occurred. We might restate this as follows: *John slapped Bill* is true at the speech time if and only if there is some time t before the speech time when *John slaps Bill* is true. This is, in a nutshell, the analysis ARTHUR N. PRIOR offers of the past tense.

Prior published most of his work on tense logic in the 1950s and 1960s (Prior 1957, 1967, and 1968). His 1968 paper *Tense Logic and the Logic of Earlier and Later* (Chapter 5 in this volume) has the twin virtues of being self-contained and not requiring the reader to have a technical background in philosophy (this cannot be said of his other papers). The first section of his paper discusses the minimal tense logic K_t , developed first by E. J. Lemmon. K_t uses the operators P (for past) and F (for future), defined as below. Here, ϕ and ψ are propositions in propositional logic j 'not' is a logical negation.

1. $P\phi$ is true at time t iff there is a t' prior to t such that ϕ is true at time t' .
2. $F\phi$ is true at time t iff there is a t' after t such that ϕ is true at time t' .

AQ: Please check logic j or logic i.

Two additional operators H (for ‘has always been’, defined as not P not ϕ) and G (‘is always going to be’, defined as not F not ϕ) are also specified.

The logical system K_t consists of the following four axioms:

- (a) $\phi \rightarrow H F \phi$: What is, has always been going to be;
- (b) $\phi \rightarrow G P \phi$: What is, will always have been;
- (c) $H(\phi \rightarrow \psi) \rightarrow (H\phi \rightarrow H\psi)$: Whatever always follows from what always has been, always has been;
- (d) $G(\phi \rightarrow \psi) \rightarrow (G\phi \rightarrow G\psi)$: Whatever always follows from what always will be, always will be.

The rest of the chapter enhances this logic. First, the minimal propositional tense logic above is related to a minimal calculus defined over instants, that uses the operators T and U : here Tap means ‘it is the case at instant a that p ’ and Uab means ‘instant a is earlier than instant b ’. A logic based on the latter calculus forms the first of four logical systems of increasing power developed in the paper. For example, the second logical system relies on allowing the variable p in Tap to stand for complex expressions like Tbp or Uab , resulting in expressions like $TaTbp$ and $TcUab$.

Unfortunately, the paper by Prior reprinted in Part I uses a somewhat opaque prefix notation. This notation may be explained as follows, from Prior (1957, chapter 2):

- Apq : Either p or q
- Cpq : If p then q
- Epq : Iff p then q
- Kpq : Both p and q
- Np : Not p
- Σp : For some p
- Πp : For all p

Thus, for example, $CFFpFp$ is the same as $FFp \rightarrow Fp$ and the rule (d) of the System K_t above is the same as $CGC\phi\psi CG\phi G\psi$.

Our illustration of Priorean logic applied to specific English tenses is based in part on de Swart and Verkuyl (1999), and is shown in Table 2. As can be seen, Prior’s approach allows free iteration of tense operators to get more complex tenses. However, the free iteration is too powerful: it generates many expressions which can’t possibly be natural language (NL) tenses, e.g. $PPPP\phi$ **It was the case that it was the case that John had slept*. The progressive isn’t distinguished; also, the past perfect has a treatment which is analogous to Reichenbach’s. Unlike Reichenbach, however, there is no distinction between present perfect and simple past. Nor is there any operator for present tense. Despite these shortcomings, Prior offers a set of formal logics for representing tense meanings, and it has been extended in various ways, for example, by Kamp (1971) and others. For more on logics for reasoning about time and events, see Part II.

5.3 Imperfective Aspect: Semantics of the Progressive

The **progressive**, which in English expresses imperfective aspect, has attracted the attention of semanticists on account of its peculiar semantic behavior. In the case of activities, the use of the progressive implies culmination: *John was running* implies *John ran*. In the case of telic eventualities, however, this isn’t so: *John was running home* doesn’t imply *John ran home*. Since the progressive expresses imperfective aspect, this

peculiarity or puzzle is referred to as the **imperfective paradox** (Dowty 1979). Note that earlier approaches such as Reichenbach and Prior gave the progressive a similar semantics to the simple tenses.

Dowty addresses this paradox by considering future branches of an event that seem ‘normal’. Accordingly, he introduces the concept of an **inertia world**: i.e. a world which is like the current world up to *Now* and is a ‘normal’ continuation of it thereafter, i.e. a world where nothing ‘abnormal’ happens. (Such worlds are of course a philosophical abstraction!) He postulates the following meaning for the progressive morpheme (e.g. *-ing*):

Progressive V is true in the current world at a time interval I if and only if there is some superinterval I' \supset I extending beyond I into the future such that V is true at I' in every corresponding inertia world.

Consider an event *V* true at *I'* in every inertia world. If *V* is an activity like *John was running*, since it has the subinterval property, *V* is true at *I* in every inertia world. Since the inertia world is like the current world up to *Now*, *V* is true in the current world as well. In contrast, if *V* is a telic event like *John was running home*, since the subinterval property doesn't hold for it, no such inference can be made about *V* being true in the current world.

Inertia worlds, while restricted to normalcy, can still result in noncompletion of an event, e.g. *Mary was crossing the street, when a truck hit her* (Vlach 1981). As an alternative, Parsons (1990) proposes a nonmodal theory, where the progressive changes a telic event to be nontelic, i.e. instead of telic *V* culminating at time *t*, it merely holds at time *t*. This approach, Parsons admits, commits to the existence of possibly incomplete objects, e.g. *John was building a house*, on Parson's account, implies that the house exists even if the building event doesn't culminate for some reason (e.g. after some initial activity, the building event might have been aborted due to cost overruns). For more details on the progressive, see Kuhn and Portner (2002), Landman (1992), and White (1994).

5.4 Perfective Aspect: Semantics of the Perfect

The perfect, which is an expression of perfective aspect, poses several semantic puzzles. It has an existential use, locating the event somewhere in the indefinite past. There is also a universal use involving the predicate holding over an entire period, as in *John has lived in Boston*. Whether this period extends up to now or not depends in part on the context.

The present perfect, in particular, also involves **recency**, cf. *Mary has run into trouble*. As Kuhn and Portner (2002) point out, that recency is part of the meaning is supported by the view that the present perfect can always take temporal modifiers that pick out intervals overlapping the present, but never those that pick out intervals entirely preceding the present; viz. *Mary has bought a dress since Saturday*, but **Mary has bought a dress last week*.

The use of the present perfect also involves **felicity conditions**. As Chomsky pointed out, *Einstein has visited Princeton* is infelicitous today, but *Princeton has been visited by Einstein* is ok today. As Kuhn and Portner point out, *My mother has changed my diapers many times* is infelicitous when spoken by an adult. For more on the formal semantics of the perfect, see Klein (1994), Kuhn and Portner (2002), and Portner (2003).

6 COMPUTATIONAL APPROACHES

6.1 Introduction

In the above sections, we have introduced some of the theoretical tools that linguistics and logic provide for analyzing temporal information in natural language. These tools supply, at the very least:

- (1) a theory, expressed in rules of some form, as to how particular tense features and temporal adverbials contribute to temporally locating events;
- (2) an event ontology, namely a description of event subclasses and components, and their temporal properties. In linguistic terms, the ontology falls out of a theory of aspect, including an account of lexical aspect, an aspect calculus, and an account of the semantics of perfective and imperfective aspects, along with a theory of event structure.

The computational approaches in this section demonstrate how some of these tools can be applied in the design of systems to process temporal information in NL. These approaches go well beyond merely applying these tools; they further refine the theories, integrate multiple theories, and test some of the analyses on actual data from an application.

6.2 Computing Aspect: Transition Networks

MOENS AND STEEDMAN (Chapter 6) offer an event ontology with event subclasses that are distinguished similarly to those in Table 1. Telicity is similar to its sense in Table 1, and associated here with a result or consequent state. States are viewed as extending indefinitely in time, while events have beginnings and ends. Event subclasses are also distinguished based on whether they are atomic or extended, analogous to the Durativity feature in Table 1. Thus, there are atomic activities (corresponding to semelfactives) as well as atomic accomplishments (i.e. achievements). In addition to these event subclasses, the event ontology of Moens and Steedman provides a tripartite model of event structure. They represent events that culminate with an event structure consisting of a preparatory process, a culmination, and a result state. Note that the breaking out of the culmination distinguishes it from Pustejovskian version of event structure. Any of these three components can in turn be compound; the preparation for a climb of Everest can include various events, the reaching of the top can be broken down further, and the result state can include further events that are in the same sequence of contingently related events as the culmination. Moens and Steedman motivate this account of event structure by the different readings one gets with *when*-adverbials adjoined to a main clause.

Moens and Steedman go on to develop an aspect calculus, represented as a remarkably compact transition network. Thus, for example, to get an interpretation of *Harry was hiccupping*, the semelfactive (i.e. a point process) *Harry hiccupped* is changed into a process of iteration of the point, and that process of iteration is then in turn changed by the progressive into a progressive state. Temporal adverbials can trigger specific transitions, e.g. *Harry hiccupped for three hours* would be analyzed as a process of iteration of the point event of hiccupping.

Their paper goes on to show how the transition network is able to express a wide variety of aspectual distinctions. It accommodates multiple interpretations for

a sentence by allowing for multiple paths through the network. The network can admit ambiguities that seem spurious, for example, *Harry is speaking* could mean a process changed into a progressive state, or a process that is changed into a point and iterated before being changed into the progressive state. The authors show in some detail how particular readings are preferred, and state upfront that the appropriateness of any particular transition is ‘conditional on support from (world) knowledge and context’. Their account also doesn’t allow transitions out of states, so that a progressive can’t be changed into an activity (excluding, for example *He was winning for an hour at the tables when his luck turned*). Their treatment of the progressive is a nonmodal account, and as such it commits to the existence of incomplete objects. Further, some of their assumptions are incorrect, e.g. their claim that the perfect applies only to culminations—what about *John has lived in Boston*? Also, their paper stops short of offering rules for a compositional semantics tied to the specific syntactic structure of sentences. (For an alternative, biologically motivated approach to aspectual composition, see Chang et al. (1998).)

6.3 Computing Aspect: Decomposed Event Ontologies

The work of DORR AND OLSEN (Chapter 7) uses an aspectual classification similar to Table 1. However, they argue, based on Olsen (1994), that while aspectual composition can change a negative feature to positive, the converse doesn’t occur. As a result, they get rid of the negative feature values. (This step may be misguided: an accomplishment can change into an activity, as in *Ulrich was winning the race*, and an activity can change into a state, as in a stative treatment of the progressive, as we saw in the Moens and Steedman work.) The result, from a computational point of view, is that features just get added, never deleted, resulting in aspectual composition being a monotonic process. The paper goes on to represent verb meanings in terms of ontological primitives based on the decompositional theory of Jackendoff (1990). Telicity, Dynamicity, and Durativity can then be mapped to the presence of particular primitives in the verb meanings. For example, Jackendoff represents states with one of two primitives; the absence of these is a test for Dynamicity.

Having specified lexical aspect based on the meaning representations of lexical entries (this requires some modification to their pre-existing lexical entries), Dorr and Olsen show how the aspectual class of composed elements can be derived by the same algorithm which determines lexical aspect, namely, by testing for particular elements in the semantic representation. The paper goes on to describe how these aspectual features are exploited to construct translations in Machine Translation applications. For example, their system requires that sentential translations must match in telicity, even if the verb and its translation do not, cf. *marched across the field*, which has the atelic verb *march*, versus the Spanish translation *cruzo el campo marchando* (*crossed the field by marching*), which has the telic *cruzo*. Their lexicon is also available for free use by other researchers (Dorr 2003). This work has also been exploited in natural language generation (Dorr and Gaasterlaand 2002).

Overall, their approach shows how lexical aspect and aspectual composition can both be derived from a particular decomposed representation for lexical meanings. In other words, the primitives of the event ontology determine aspectual classes of lexical and composed expressions. In their framework, aspectual composition is a side-effect of semantic composition. The downside of their approach is of course that this advantage

is won by committing to a decompositional theory that is prone to the problem of proliferation of semantic primitives. Most importantly, however, their paper shows how this kind of representation can be made practical in terms of reusable lexical resources.

6.4 Computing Temporal Locations of Actual Events

The final computational paper in this part is that of PASSONNEAU. Her goal is to process references to eventualities, locating them temporally, as part and parcel of the PUNDIT information extraction system. Passonneau's event ontology distinguishes states, activities, and accomplishments (which she calls 'transition events'). As with Vendler, semelfactives aren't distinguished. Achievements aren't distinguished from accomplishments, since she views them as having short durations rather than being instantaneous; also, whether an achievement is viewed as a point or an interval is considered an issue of granularity. As a result, she doesn't have a Durativity feature. She introduces a feature for Boundedness: states are inherently unbounded, whereas transition events are bounded; activities may be bounded or not depending on sentential context. Like Dorr and Olsen, she uses a decompositional approach to representing aspectual meaning; however, unlike them, she uses the semantic framework of Dowty (1979). Passonneau also independently develops a model of event structure for transition events that is extremely similar to the one proposed by Moens and Steedman.

As in most information extraction research, Passonneau focuses on **actual** events, namely, realis events that are specific, that have occurred earlier or that are occurring at the time of the text. Realis events are characterized as being those not in modal, negated, intensional, or frequency adverbial contexts. Her system works in three steps.

In Step I, each utterance is analyzed in terms of the aspectual class, the tense (past/present) and grammatical aspect (perfect or progressive), the main predicate and its arguments, and (finally) the time over which the event occurs or (in the case of states) holds—the latter is explained below in the account of Step II. Any present tense nonstative event that isn't in the progressive or perfect is treated as nonactual, with all other events, excluding those on a stoplist, being treated as actual.

In Step II, the temporal structure of the eventuality is determined. If it is a state, since it is inherently unbounded, the event producing the state occurs in a subinterval of the state; for example, for the progressive, stative sentence *Metallic particles are clogging the strainer*, the time of the ongoing clogging event is assumed to be an arbitrary moment included within the unbounded interval of the state of being clogged. (Note that unlike Moens and Steedman's approach, there is no notion of iteration here.) In contrast, in the case of a transition event, as in the sentence *The pump failed*, the moment of the event's culmination starts the period of the result state.

In Step III, Reichenbach's tense representation is used to order the event time with respect to the reference time and situation time. In this process, temporal adverbials are interpreted. Thus, given *The compressor failed before the pump seized*, the reference time for each clause is computed as being arbitrary moments included in the failing and seizing events respectively. These reference times are then ordered based on the semantics for *before*.

Passonneau's focus is on the contributions of tense and lexical aspect. Her work is characterized by many interesting details, including the handling of subordinated events. While the theoretical tools she uses (event ontology and tense analysis) are very general, the types of texts her system is used for allow her not to be directly concerned with

aspectual composition. Also, checks for realis events in terms of the presence of modal and other contexts aren't discussed in her paper. Her work, nevertheless, demonstrates how these theoretical tools can be used within a practical system to locate events in time. We shall see that such work forms a natural precursor to subsequent work in temporal information extraction, including work discussed in Part IV. We turn now to a system that uses tense and aspect information to locate events in time, but one that is not aimed at the information extraction problem.

Ion Androutsopoulos's work (Androutsopoulos 1996, 2002; Androutsopoulos et al. 1998) addresses problems in natural language interfaces to databases. His NLITDB system allows users to pose temporal questions in natural language to an airport database. NLITDB maps English queries to a temporal extension of the SQL database language, via an intermediate semantic representation. The latter is influenced by the work of Crouch and Pulman (1993), and uses temporal operators, including the Priorean Past operator, along with Reichenbachian indices. (The source code, which relies on Head-Driven Phrase Structure Grammar (HPSG), is freely available.) Although we cannot include a reprint in this volume for reasons of length, we provide a summary of his approach. Readers are referred to his book (Androutsopoulos 2002) for more details.

In NLITDB, verbs in natural language questions are classified into (at least one of the) four aspectual classes based on manual analysis using linguistic tests on question sentences from an application domain test suite. For example, verbs that occur in single-clause questions with nonfuture meanings are treated as states (e.g. *Does any tank contain oil?*). A verb like *land*, which can be used habitually, e.g. *Which flight lands on runway 2?*, or nonhabitually, e.g. *Does flight BA737 land on runway 2 this afternoon?*, will be treated as ambiguous, with a habitual state entry and an accomplishment entry. Activities are distinguished using the imperfective paradox: *Were any flights taxiing?* implies that they taxied, while *Were any flights taxiing to gate 2?* does not imply that they taxied. So, *taxi* will be given an activity verb sense, one that doesn't expect a destination argument, and also an accomplishment verb sense, one that expects a destination argument.

NLITDB has extensive coverage of temporal adverbials. The semantic relation between a question event and the adverbial it combines with is inferred by a variety of inference rules. In the case of a state combining with a 'point' adverbial, e.g. *Which flight was queueing for runway 2 at 5:00 pm?*, the state is coerced to an achievement and is viewed as holding at the time specified by the adverbial. When an activity combines with a point adverbial, it can mean that the activity holds at that time, but an inceptive reading that the activity starts at that time is also possible, as in *Which flight queued for runway 2 at 5:00 pm?* An accomplishment may indicate inception or termination; for example, *Which flight taxied to gate 4 at 5:00 pm?* can mean the taxiing starts or ends at 5 pm. In the case of 'interval' adverbials, the event holds for at least part of the adverbial's duration or is included within the duration of the adverbial. The system also has a fairly extensive treatment of *while*, *before*, and *after* clauses.

7 CONCLUSION

This overview has introduced the theoretical tools for analyzing how natural language expresses the location of events in time (via tense), whether they are ongoing, completed, or terminated (via grammatical aspect), and the ontological distinctions among events

that language makes (in terms of aspectual subclasses of events, as well as event structure). We have also explored computational approaches to aspectual composition based on particular event ontologies. Finally, this Introduction has presented information extraction as a means of temporally locating extracted events. It is worth noting that the computational approaches rely heavily on these theoretical tools. In Part III, we will show how the analysis of sentences in a discourse context is used to locate and order events in time, examining formal approaches to discourse analysis. Finally, In Part IV, we will explore current corpus annotation schemes for representing events and their temporal locations, that can in turn give rise to advanced temporal information extraction systems.

REFERENCES

- Androutsopoulos, I. (1996). 'A Principled Framework for Constructing Natural Language Interfaces to Temporal Databases'. Ph.D. dissertation, University of Edinburgh.
- (2002). *Exploring Time, Tense and Aspect in Natural Language Database Interfaces*. Amsterdam: John Benjamins.
- Ritchie, G. D., and Thanisch, P. (1998). 'Time, tense and aspect in natural language database interfaces', *Natural Language Engineering* 4(3): 229–76. (See <http://www.aueb.gr/users/ion/nlitdb_book/> for freely downloadable software.)
- Binnick, R. (2002). 'Project on the Bibliography of Tense, Verbal Aspect, Aktionsart, and Related Areas'. <<http://www.scar.utoronto.ca/~binnick/TENSE/Aktions.html>>.
- Carlson, G. (1977). 'Reference to Kinds in English'. Ph.D. dissertation, University of Massachusetts, Amherst.
- Chang, N., Gildea, D., and Narayanan, S. (1998). 'A Dynamic Model of Aspectual Composition', in M. A. Gernsbacher and S. J. Derry (eds.), *Proceedings of the 20th Annual Meeting of the Cognitive Science Society*, Madison, Wisconsin. Mahwah, NJ: Lawrence Erlbaum Associates, 226–31.
- Comrie, B. (1986). *Tense*. Cambridge: Cambridge University Press.
- Crouch, R. S. and Pulman, S. G. (1993). 'Time and modality in a natural language interface to a planning system', *Artificial Intelligence*, 63: 265–304.
- Dorr, B. (2003). 'Lexical Resources' <www.umiacs.umd.edu/~bonnie/LCS_Database_Documentation.html>.
- and Gaasterland, T. (2002). 'Constraints on the Generation of Tense, Aspect, and Connecting Words from Temporal Expressions'. Technical Report CS-TR-4391, UMIACS-TR-2002-71, LAMP-TR-091, University of Maryland, College Park, MD.
- Dowty, D. (1979). *Word Meaning and Montague Grammar*. Boston: D. Reidel.
- Hinrichs, E. A. (1985). 'Compositional Semantics for Aktionsarten and NP Reference in English'. Ph.D. dissertation, Ohio State University.
- Jackendoff, R. (1990). *Semantic Structures*. Cambridge, MA: MIT Press.
- (1991). 'Parts and Boundaries', *Cognition* 21:9–45.
- Kamp, H. (1971). 'Formal Properties of "now"', *Theoria* 27: 237–73.
- and Reyle, U. (1993). *From Discourse to Logic: Introduction to Model-theoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*, 2 vols. Dordrecht, Holland: Kluwer.
- Klein, W. (1994). *Time in Language*. London: Routledge.
- Krifka, M. (1989). 'Nominal Reference, Temporal Constitution, and Quantification in Event Semantics', in R. Bartsch, J. van Benthem, and P. van Emde Boas (eds.), *Semantics and Contextual Expressions*. Dordrecht, Holland: Kluwer, 75–115.
- Kuhn, S. and Portner, P. (2002). 'Tense and Time', in D. Gabbay and F. Guenther (eds.), *The Handbook of Philosophical Logic*, vol. 7 (2nd edn.). Dordrecht, Holland: Kluwer.

- Landman, F. (1992). 'The Progressive', *Natural Language Semantics* 1: 1–32.
- Lin, J-W. (2003). 'Temporal Reference in Mandarin Chinese', *Journal of East Asian Linguistics* 12: 259–311.
- Link, G. (1983). 'The Logical Analysis of Plurals and Mass Terms', in R. Bauerle, C. Schwarze, and A. von Stechow (eds.), *Meaning, Use and Interpretation of Language*. Berlin: Walter de Gruyter, 302–23.
- Olsen, M. B. (1994). 'The Semantics and Pragmatics of Lexical and Grammatical Aspect', Ph.D. dissertation, Northwestern University, Evanston/Chicago, IL.
- Parsons, T. (1990). *Events in the Semantics of English*. Cambridge, MA: MIT Press.
- Partee, B. (1973). 'Some structural analogies between tense and pronouns in English', *Journal of Philosophy*, 18: 601–10.
- Portner, P. (2003). 'The (Temporal) Semantics and (Modal) Pragmatics of the Perfect', *Linguistics and Philosophy* 26: 459–510.
- Prior, A. N. (1957). *Time and Modality*. Oxford: Oxford University Press.
- (1967). *Past, Present, and Future*. Oxford: Oxford University Press.
- Schank, R. C. (1975). *Conceptual Information Processing*. Amsterdam: North-Holland.
- Smith, C. (1991). *The Parameter of Aspect*. Dordrecht, Holland: Kluwer.
- Suppes, P. (1999). *Introduction to Logic*. New York: Dover.
- Swart, H. de (1998). 'Aspect shift and coercion', *Natural Language and Linguistic Theory* 16: 347–85.
- and Verkuyl, H. (1999). 'Tense and Aspect in Sentence and Discourse'. Ms.
- Vlach, F. (1981). 'The Semantics of the Progressive', in P. Tedeschi and A. Zaenen (eds.), *Syntax and Semantics 14: Tense and Aspect*, New York: Academic Press, 271–92.
- White, M. (1994). 'A Computational Approach to Aspectual Composition', Ph.D. dissertation, University of Pennsylvania.

