UNIVERSITY OF ALBERTA

A Logical Approach to Narrative Understanding

By

Chung Hee Hwang

A Thesis

Submitted to the Faculty of Graduate Studies and

Research in Partial Fulfillment of the Requirements for

The Degree of Doctor of Philosophy

Department of Computing Science

Edmonton, Alberta

Fall 1992

ISBN 0-315-77117-8

UNIVERSITY OF ALBERTA

RELEASE FORM

Name of Author: **Chung Hee Hwang**

Title of Thesis: **A Logical Approach to Narrative Understanding**

Degree: **Doctor of Philosophy**

Year This Degree Granted: **1992**

Chung Hee Hwang

1100 S. Goodman Street

Rochester, New York 14620

U. S. A.

Date October 8, 1992

UNIVERSITY OF ALBERTA

FACULTY OF GRADUATE STUDIES AND RESEARCH

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research for acceptance, a thesis entitled **A Logical Approach to Narrative Understanding** submitted by **Chung Hee Hwang** in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

L.K. Schubert

Lenhart K. Schubert, Co-supervisor

F. Jeffry Pelletier

F. Jeffry Pelletier, Co-supervisor

R. Goebel

Randy Goebel

Peter van Beek

Peter van Beek

Bernard Linsky

Bernard Linsky

Bonnie L. Webber, External Examiner

Date Sept 11/92

To Mother

Abstract

It is argued that a "theory bottleneck" encountered in the 70's and early 80's in attempts to build comprehensive NLU systems led to a fragmentation of NLU research, which still persists. NLU is an organic phenomenon, and enough has been learned about the vexing problems of the 80's to try to integrate these insights and build more comprehensive theories and extensible implementations. On that premise, a new comprehensive framework for narrative understanding has been developed. Its centerpiece is a new situational logic called *Episodic Logic* (EL), a highly expressive knowledge and semantic representation well-adapted to the interpretive and inferential needs of general NLU.

EL is Montague-inspired and influenced by situation semantics. It provides an easily computed first order logical form for English. It allows propositional attitudes, unreliable generalizations, and other non-standard constructs, including ones involving events, actions, facts, kinds and donkey sentences. It incorporates a DRT-like treatment of indefinites, and makes systematic use of *episodic* variables in the representation of episodic sentences, using them to capture temporal and causal relationships. The rules of inference in EL include probabilistic versions of deduction rules resembling forward and backward chaining rules in expert systems.

Also developed is a uniform, compositional approach to interpretation in which a parse tree leads directly (in rule-to-rule fashion) to a preliminary, *indexical* logical form, and this indexical LF is deindexed with respect to the current *context* (a well-defined structure). The initial translation is obtained using a GPSG-style grammar; the latter transformation is accomplished by a new recursive deindexing mechanism. Deindexing simultaneously transforms the LF and the context: context-dependent constituents of the LF, including tense, aspect and temporal adverbials, are replaced by explicit relations among quantified *episodes*, bringing the context information into the LF, thus removing context dependency; and new structural components and episode tokens are added to the context. The relevant context structures are called *tense trees*. The mechanism allows reference episodes to be correctly identified even for embedded clauses.

Finally, a pilot implementation is able to make many (though not all) of the inferences described in this thesis, and has been successfully used in several domains.

Acknowledgements

I am grateful to Len Schubert for his help and guidance throughout my thesis research. When I first plunged into NLU fresh from computer graphics, he gladly took me as a student and started to teach me virtually everything — from English grammar to Montague grammar, from how to form an idea to how to stick out one's neck, and from how to pronounce "wolf" to how to shop for a used car. But more importantly, he was a great source of ideas; and when things got tough with Episodic Logic, it was he who maintained course and would not let me give up.

Episodic Logic, as reported in this thesis, is indeed the result of our joint research — the result of years of discussion, skirmishes, and technical trial-and-error. It was a great experience to work with him, and without his continuous input and rigorous checking, Episodic Logic would not exist today. I sometimes felt frustrated in trying to keep up with Len's boundless resourcefulness and diligence. It is no accident that the writing in my thesis often slips into the first person plural. In fact, many topics remain that we jointly worked on, but which didn't make it into the thesis. I hope we can come back to them. Any remaining errors and shortcomings of the thesis (and I am sure there are many) are my own.

This research was done away from Alberta, and I thank Jeff Pelletier for making it possible. He gladly assumed the responsibility of co-supervisor, when Len and I moved to Rochester, and helped me to see the end. I especially appreciate his encouragement and cheerful email messages, received while I was writing the thesis. Also, his pertinent comments and questions helped improve the thesis. Special thanks also go to my external examiner, Bonnie Webber. Her careful reading caught many mistakes and lacunae in the thesis draft, both technical and stylistic in nature, and I am grateful she detected them before they reached print. She also gave me invaluable advice about tone in writing, which I will keep in mind. I would also like to thank Randy Goebel, Peter van Beek and Bernie Linsky for reading the thesis and asking many thoughtful questions, Jonathan Schaeffer for chairing the exam, Stan and Lisa Cabay for their support throughout my old Alberta days and during my last stay for the exam, Tony Marsland for cajoling me into finishing my thesis, and Paul Sorenson for extending his support to me in various forms as Department Chair, even though we had never met.

This gigantic project would not have seen the light of day without the tremendous efforts of Stephanie Schaeffer, undertaken with financial support from the Boeing Company. Stephanie transformed a tiny inference program into a huge EPILOG system, never giving up on the ever-changing nature of Episodic Logic. I also greatly appreciate the help of Phil Harrison at Boeing. His thoughtful questions, examples, and suggestions led us to change our theory on many occasions. I especially benefited from his proofreading while I was writing the second half of the thesis. Many thanks also to Edith Drummond and Debi Owens for patiently and expertly handling the many administrative problems and chores that go with managing a "remote" Ph.D. student's program and financial support.

I am grateful to James Allen for inviting me to Rochester, extending me financial support, and providing helpful and important comments while Episodic Logic was being developed. Rochester is indeed a great place to work. I especially enjoyed the company of David Traum, Massimo Poesio, and Marc Light (in the order of proximity of their offices to mine). I learned a lot from discussions and debate with them, and from their reading of various chapters of my thesis. I also thank my graduate student friends near and far who had to answer my never-ending questions about how a native English speaker would interpret or react to this or that: Hannah Blau, Virginia de Sa, Janet Hitzeman, Alice Kyburg, Catherine Chronaki, George Ferguson, Nat Martin, Jack Veenstra, Lambert Wixson, Ray Rimey, Peter Heeman, Leo Hartman, Mike Swain and Josh Tenenberg, to name a few. I would also like to thank Jill Forster for taking care of the various needs of a "special" student, and other staff members at Rochester for their moral support. Stephanie (once again) and Rebecca and Marina have my gratitude for sustaining me with friendship and glimpses of other worlds than the academic one, on numerous occasions when discouragement and frustration got the better of me. And I heartily thank my other friends: Choi, Yoon and Shin.

Finally, my dearest love and thanks to Brian, Paul, Tim and E Young, for their love and support and prayer. And to Mom, for everything.

This research was supported in part by a Killam Memorial Scholarship, an Andrew Stewart Graduate Prize, the Boeing Co. in Seattle (under Purchase Contract W 288104), NSERC Operating Grant A8818, ONR/DARPA research contract no. N00014-82-K-0193, NSF Research Grant IRI-9013160, and graduate assistantships from the Department of Computing Science, University of Alberta, and the Department of Computer Science, University of Rochester.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Background: Need for an Organic Approach to NLU

During the heady days of the 70's and early 80's, each year appeared — at least to a casual observer — to bring a more complete picture of the essential structures and processes underlying language understanding and a more impressive set of working systems demonstrating the power of the new ideas and mechanisms in small, but subtle, "human interest" domains. For example, CYRUS [Kolodner, 1981] and BORIS [Dyer, 1983; Lehnert *et al.*, 1983] showed a remarkable degree of understanding in complicated human domains such as divorce stories, diplomatic visits, and newspaper reports, as opposed to blocks worlds or airline reservations. The main obstacle to building truly comprehensive understanding systems then appeared to be the "knowledge bottleneck," i.e., the problem of how to impart staggeringly large amounts of well-structured knowledge to them. Consequently, efforts were launched to develop effective methods for autonomous knowledge acquisition, or alternatively, to cook up the requisite knowledge through coordinated, massively parallel human effort.

Although these assaults on the "knowledge bottleneck" are interesting, they have not led to any breakthroughs in NLU. The truth is that those earlier ambitious systems proved to be very complex and hard to extend in scope beyond the initial domains they were designed to handle. Indeed, the efforts of the 70's and early 80's ran not only into a knowledge bottleneck, but also into a multifaceted "theory bottleneck." Even syntax and parsing, regarded by many in the 80's as virtually solved problems, confronted researchers attempting to build comprehensive grammars and parsers with numerous difficult problems (e.g., correct formulation of feature propagation principles, constraints on unbounded dependencies, apparent non-context-freeness, rampant ambiguity, need for error tolerance, not to mention specifics like the syntax of coordination, auxiliaries, verb subcategorization, etc.). And these problems were relatively well-explored compared to problems in computing logical form, deriving unambiguous, formally interpretable and usable semantic representations, and performing well-founded inferences based on these semantic representations in conjunction with background knowledge.

Gradually, the AI/CL community came to appreciate the existence and magnitude of the theory bottleneck. This began with a growing awareness that the fragility and inextensibility of existing NLU systems was attributable not just to their meager knowledge but also to the myriad rough-and-ready assumptions and domain-specific hacks employed to achieve demonstrable output; the appreciation of the theoretical difficulties grew further as more and more insights from theoretical linguistics, psycholinguistics and philosophical logic were assimilated into AI, which made clear the subtlety of problems in syntax, semantics, and pragmatics yet to be overcome.

As a result, the field as a whole has shifted toward specialization and fragmentation, in a striking contrast with its character in those earlier days. To some extent, this fragmentation represents an appropriate response to the variety and subtlety of remaining problems. But with most NLU researchers addressing quite narrow issues only, and most work being done within a specialized theoretical framework — a particular type of grammar, a particular class of parsers, a particular style of semantic representation, a particular theory of discourse structure, etc. — what has been missing is a global perspective, a comprehensive framework and an effort to connect various aspects of language understanding. It is true that attempts at comprehensive theorizing and system-design have not been completely abandoned. For example, Hobbs *et al.* [1986; 1988; 1990] and Charniak *et al.* [Charniak, 1988] [Charniak and Goldman, 1988, 1989a, b] [Charniak and Shimony, 1990] have kept their sights on full understanding, and in particular, have made very intriguing proposals for full integration of all types of disambiguation based on abduction. But the overwhelming majority of researchers in recent years have shown a curious reserve toward full understanding, most strikingly in their approach to the core problem of semantic representation.

So, if the theoretical bottleneck encountered in the 80's was real (as it undeniably was), is the current emphasis on highly focused, fragmented studies to be applauded? In some ways, surely yes. Divide and conquer is a tried and true research strategy, and there are many intriguing problems in all facets of formal and computational linguistics left to conquer. In taking receipt of the insights of linguists, philosophers, logicians, and psychologists, computational linguists have also taken receipt of their long-standing problems, which promise to provide many more decades — or centuries — of food for thought.

At the same time, there has been something of a loss of nerve within the AI/CL community. Language understanding (more generally, linguistic communication) is an *organic* phenomenon in the sense that each facet is strongly dependent on the others. If surface form determines logical form, and logical form determines the ultimate meaning representation, and the ultimate meaning representation determines further conversational (and other) behavior — and if all of these transductions are mediated by inferential use of world knowledge and by a shifting context of salient features of the discourse situation — then surely there is a point where study of isolated features of this organic whole becomes less profitable than an attempt to see it in its entirety. The need for integration and a global perspective is even more pressing in computational linguistics than in other sciences, since most facets of our problem do not even have a clearly discernible shape *independently* of their relation to other facets. Even syntactic structure, the most accessible aspect of language, is moot, and logical form, semantic representation, knowledge representation, context, and inferential operations are utterly hypothetical, and tightly interlocked. Thus, work done on one issue while simplifying or ignoring the rest is almost certain to go off in quite different directions than work which attempts to keep in mind all constraints and desiderata at once.

Luckily, compared to the 70's and early 80's, present prospects for principled, integrated NLU are greatly improved. Considerable strides have been taken in our understanding of all aspects of language processing. Current versions of GPSG, HPSG, categorial grammars, and other grammatical formalisms now account for a wide range of syntactic phenomena as well as shedding light on semantic type structure. LR-like parsers and techniques for exploiting statistical correlations are bringing NL-parsing closer to practicality. Theories of intention, speech acts and discourse structure are coming to grips with language as goal-directed interaction. And most importantly from our perspective, new logical frameworks such as DRT, situation semantics, and type and property theories have been developed to address various long-standing semantic conundrums, such as the semantics of attitudes, anaphora, kinds, substances and collections, properties, propositions, events, and tense and aspect. Thus, the time seems ripe for working on NLU theories or systems that aspire to be complete and comprehensive, with respect to all the major syntactic, semantic, and pragmatic phenomena encountered in NL texts.

This thesis describes an attempt that takes up that call — an effort to "put it all together," exploiting the insights from various subfields, and building more comprehensive theories and extensible implementations. Such an integration is certainly not a trivial task. What is needed here is a *creative synthesis*. Simply "shopping" for the right building blocks, among all those currently being offered by various authors and schools of thought, will not work. The blocks just don't fit together — at least not in any obvious way. So, many adaptations, modifications, and augmentations of extant ideas will be required. Moreover, a comprehensive design cannot be expected to be flawless from the outset. Rather, there must be gradual "debugging" and theoretical deepening of the entire account, both in the theoretical and computational details and the overall abstract architecture. As an important "reality check" on this process of gradual refinement, development of the framework should go hand-in-hand with implementation, in the best tradition of AI. The research reported in this thesis — an attempt to provide a comprehensive framework for general NLU — has been undertaken with that in mind. The battle is still going on, and it will be many years till the full goal is met. In that sense, this thesis may be considered as a progress report. But I believe the endeavor has been successful, with significant results in several aspects of the theoretical framework and encouraging results in preliminary implementations. Before getting into the discussion of what the research has specifically accomplished though, I will first discuss what I view as the core problems in building general NLU systems.

1.2 Desiderata for General NLU Systems

Before actually attempting to "put it all together," we need to develop a general view of the stages or facets of the understanding process. These stages or facets should be intuitively natural, simply interrelated, in principle mathematically analyzable, and computationally implementable. At the most general level, the task can be viewed in terms of the schema shown in Figure 1.1. Utterances in conjunction with suitable context structures need to be

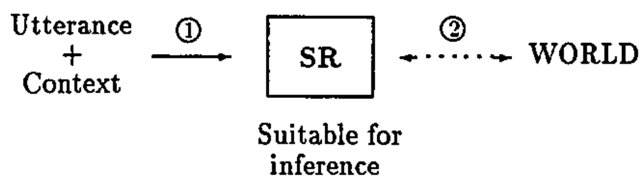

Figure 1.1: Schema for Narrative Understanding Process

mapped to a semantic representation SR (the process indicated as ①), and this SR should be *interpretable* relative to the world (the relation indicated as ②), with truth conditions that are in accord with intuitions about the original utterances.

Both during the computation of utterance meaning, and upon its completion, a great deal of "spontaneous" (input-driven) inferencing is presumed to occur, working out plausible interpretations and consequences based on the discourse interpreted so far, meaning postulates and world knowledge. This includes computing unique referents for referring expressions, predictions, and explanations which ultimately give a causally coherent elaboration of what has been said. Therefore, an essential requirement is that the representation support such inferences and the knowledge behind them. It should do so in a way that is both intuitively transparent and analyzable in terms of a formal notion of interpretation. The interpretability of SR will give us some assurance that inferences based on the SR will be those intuitively warranted.

1.2.1 Representation as the "Core" of the Problem

This general picture points to the centrality of the issue representation. The ease of mapping from syntax to a semantic representation, "deindexing" (amalgamating the context information into the representation of an utterance so that the resulting representation becomes context-independent), and performing inferences all depend on the representation used.

A basic methodological assumption of this thesis is that these multiple demands on the representation are best met by using a highly expressive logic closely related to NL itself. The possibility of handling tense, causes, facts, modifiers, propositions, beliefs, etc., simply and directly depends on the expressiveness of the representation. To see the importance of this issue, let us consider the following sentence from the dialog which was processed by the initial TRAINS-implementation [Allen and Schubert, 1991]:

"*We have to make orange juice.*"

As Allen and Schubert point out, this simple sentence exemplifies the following interesting semantic phenomena:

- It expresses an obligation (and puts that obligation, in part, on the hearer);
- It expresses joint agency;
- It is tensed (thereby expressing that the joint obligation is upon the hearer and the speaker at the time of speech);
- It involves a mass noun phrase, 'orange juice', which on most modern accounts denotes an abstract *kind*;
- It involves a verb of creation, 'make', which in combination with its kind-denoting object implies the coming-into-being of some *quantity* of orange juice (this coming-into-being cannot be easily expressed by existential quantification); and
- It involves an infinitive whose interpretation is arguably a reified property.

This illustrates what subtle meanings even seemingly simple sentences may have. Note that this is a typical situation we run into with real spoken and written language. Without an expressively rich semantic representation, a comprehensive approach to natural language understanding will scarcely get off the ground.

These remarks apply as much to *knowledge* representation, KR, as to *semantic* representation, SR. In fact, given the tight link between interpretation and inference in the understanding process, the simplest assumption is that the two are one and the same.[1]

[1]Here I am only concerned with knowledge that can be verbally expressed. Knowledge that is given in

1.2.2 Desiderata for a SR/KR

I now summarize the desiderata for general NLU systems, especially with respect to representation and inference.

(1) It should have an expressive SR/KR. The representation should be able to express anything that is verbally expressible, including, for instance, complex quantification ("most people with two or more cars"), logical compounding ("If he fails, he is either lazy or a fool"), complex concepts ("the type of person who never forgets a slight"), modification ("a nearly invisible pale brown birthmark"), temporal relations ("He had seen her twice the previous week"), intension ("He is looking for a unicorn"), opaque contexts ("He wants to marry a blonde"), generics and habituals ("She makes a delicious pie"), and so on.

(2) The system needs to be able to reason about agents' beliefs, goals, plans, etc., both at input time and question-answering time. For this reasoning to be theoretically analyzable, the SR/KR should be formally interpretable. The representation should admit not only straightforward intuitive interpretation but also a formal denotational semantics. Words and phrases intuitively correspond to things in the world, so should their SR-translations. Thus, SR/KR needs to have a model theory. The more knowledge is added to such a system, the greater the uncertainty becomes as to the coherence of that knowledge, and what conclusions it supports. Soundness and completeness may be too much to ask as some kind of nonmonotonic reasoning—e.g., probabilistic reasoning—is a necessity in NLU; but as long as we have an interpretable SR/KR, we will also be able to interpret the current state of the system.

(3) The SR/KR should be easily computable. As will be seen, in many systems the mapping ① in Figure 1.1 is often very informal or limited to a very restricted set of constructions. Representation languages should permit systematic computation of meaning representations from English input; ideally, the logical form (LF) of an English syntactic constituent should be obtainable as a simple function of the LFs of its immediate subconstituents. The strongest form of this desideratum is that the mapping from surface form to LF should implement a *compositional* semantics in Montague's sense; i.e., the semantic value of the LF of a constituent should depend only on the semantic values of the LFs of the immediate subconstituents.[2]

diagrams, images, gesture, etc., is beyond the scope of this thesis.

[2]Note that computing LFs as a function of the LFs of immediate subconstituents does *not* automatically guarantee compositionality in this sense. After all, such a computational rule might "examine" the syntactic details of the input LFs and return an output LF dependent on those details, whether or not they make any *semantic* difference. For instance, one can easily make up semantic rule which returns radically different resultant LFs depending on what variable names are used in the input LFs, or depending on the ordering of disjuncts in a disjunction — even when these syntactic properties are not regarded as making any difference

Non-compositional approaches based on simple sentence types run the risk of being ungeneralizable.

(4) The SR/KR should be direct, transparent, and conceptually modular. By being direct and transparent, it allows all types of linguistic, world and inference knowledge to be represented in an explicit, analyzable form, thereby assuring only intuitively truthful (or at least plausible) assumptions are made. Also, it should be easy to add more lexical knowledge and more world knowledge, expanding the knowledge base indefinitely, to make the system understand discourse from any domains. Finally, it should have cleanly separated knowledge and conceptually modular inference control structure. That is, it should partition knowledge (both lexical and world knowledge) in a manageable way and strictly separate knowledge representation and inference control structure.

These are not unrealistic desiderata, and in fact pretty close to the proposal adopted at the *Workshop on the Evaluation of NLP Systems* [Palmer and Finin, 1990].[3] Obviously, it may not be easy to come up with a system or a SR/KR that meets all the requirements. Progressive refinement of an unrestricted representation seems to us to hold the greatest promise for breaking through the "theory bottleneck."

to semantic values. Equally importantly, compositionality precludes injection of "world knowledge" in the computation of an LF. For instance, suppose we had a semantic rule corresponding to phrase structure S $\leftarrow$ NP VP, stating in part that if NP denotes an inanimate object and VP denotes an action (as in "The car drove away"), then the LF of the sentence S is $(\exists x)\,\text{agent}(x) \wedge \text{VP}'(x, \text{NP}')$ (e.g., some agent drove the car away), where the "primes" indicate LFs of the corresponding constituents. This rule is noncompositional, assuming that NP$'$ could be something like "Johns-Buick" (an individual constant), and that the semantic value of an individual constant is just an individual. For in that case, we need to access stored knowledge about "Johns-Buick" to decide on the animacy or inanimacy of the subject — a property that is simply not part of its semantic value.

[3]The following is their proposal on the "dimensions under which a KR&R (knowledge representation and reasoning) system might be evaluated."

- Theory: Is there an underlying theory which gives meaning to the KR&R system? What is known about the expressiveness of the language and the computational complexity of its reasoning?
- Languages: How does the KR&R system function as a practical language for expressing knowledge? How easy or difficult is it to define certain concepts or relations or to specify computations?
- Systems: KR&R systems are more than just an implementation of an underlying theory. They require good development environments: knowledge acquisition tools, debugging tools, interface technology, integration aids, etc. How extensive and good is this environment?
- Basic models: A KR&R system often comes with some basic, domain-independent modules, such as temporal reasoning, spatial reasoning, naive physics, etc. Are such models available and, if they are, how extensive and detailed are they?

1.2.3 Informal Semantics? First-order Logic? Or?

At this point, a brief review of existing approaches to SR/KR is in order. There seem to be roughly two prevalent approaches: the informal approach and the first-order logic, FOL, approach. That is, most researchers/system designers either limit themselves to (putative) first-order translatable sentences, or deal with a much wider range but provide no formal denotational semantics.

The advantage of the informal approach is that the practitioner can quickly accommodate a rich variety of concepts and ideas in the representation — for instance, beliefs, actions, goals, habitual behavior, etc. — without being detained very much by such issues as whether the various types of symbols are being used in a coherent and consistent way, and whether or not the proposed inference methods have a rational basis in some sort of consequence relation.

CYRUS [Kolodner, 1981], BORIS [Dyer, 1983; Lehnert *et al.*, 1983], and SWALE [Schank and Leake, 1989] are examples of the informal approach. For instance, CYRUS takes conceptual representations of episodes as input (cf., Schank's conceptual dependency (CD) representation). The memory is organized with conceptual categories for events called E-MOPs. Since everything relies on memory search in CYRUS, updating and retrieval of information in and from the memory do the job of inferencing. So, the "predictive power" of a feature depends on the context in which it is found in the memory. This severely limits its inferential power, if any.

Systems like TELI [Ballard and Stumberger, 1986], PUNDIT [Dahl *et al.*, 1987] and PAULINE [Hovy, 1990] also take a relatively informal approach. For instance, the nominal 'contamination' is represented in PUNDIT as follows:

$$state(S, contaminatedP(instrument(metal1), location(filter1)), (period(S))).$$

This appears to be a "case-frame" style of representation, quite possibly reducible to FOL. However, one should note that merely because a system employs a case-frame like syntax, and certain case-frame formalisms are FOL-reducible, it does not follow that the system's particular case-frame representation is formally interpretable.

Of course, such semantically informal work is not necessarily bad. It may well be insightful and persuasive with respect to the issues it addresses. Nevertheless, the avoidance of the issue of formal interpretability leaves the overall framework excessively ill-defined. As long as we remain unclear about what sorts of things in the world our symbols can stand for, or how a putative knowledge base can conform with (or deviate from) how things actually are, we risk having the system lapse into total inconsistency and incoherence for all but trivial knowledge bases. Truth, or, "being in accord with how things actually are," is a crucial notion in inference and rational behavior, and serves as a consistency check on informal semantics. Recall that it has also illuminated reasons for intuitive consequences (or their absence) in many kinds of locutions, e.g., those involving intensions such as *seek a unicorn* or *resemble a leprechaun*, reified entities such as *to dance* (a kind of action), attitudes such as *believe*, tense, *when*-sentences, etc.

Among those who do take a more formal approach, most have restricted themselves to FOL, especially when their goal has been the actual implementation of a NLU system. The FOL approach consists of limiting the fragments of language considered to those which appear to be expressible in FOL, at least in rough-and-ready fashion. This has the advantage that FOL is well-understood syntactically and semantically, e.g., one can use standard proof techniques of FOL. But it also has the disadvantage that very little real language is easily expressible in it, as was illustrated by the semantic subtleties of the meaning of the sentence "We have to make orange juice." While with practice one can become quite good at inventing *ad hoc* FOL-approximations to English sentences, it is quite implausible that any reasonably simple, systematic transduction from syntax to semantics would deliver such FOL approximations. So the FOL approach factors out most of language — at least for any *algorithmic* (as opposed to *ad hoc*) mapping from syntax to semantic representation.[4]

Systems like UC [Wilensky *et al.*, 1988], FRAIL [Charniak, 1988], ABSITY [Hirst, 1988], and KT [Dahlgren *et al.*, 1989], use a kind of FOL language for their SR, while JANUS [Ayuso, 1989] uses an intensional logic indexed by time and world indices.

Let us look at ABSITY as a fairly representative example. It is described as a Montague-inspired compositional semantic interpreter, in which the Montague semantic objects (functions and truth conditions) are replaced with elements of the frame language FRAIL. Since it is a purely extensional first-order formalism, it cannot handle intensional contexts. Among the phenomena not handled, Hirst lists NP modifiers, habituals, certain kinds of predication, complex quantifiers, inherent vagueness, time and space, moral and contingent obligation, negation, conjunction, etc. So, for instance, the following kind of sentences are reported as beyond the scope of ABSITY.

Ross, whose balloon had now deflated completely, began to cry. (NP modifier)

Naida resembles a pika. (intension)

Ross sleeps on the floor. (habitual)

All but five of the students whose fathers like cheese gave three peaches to many of the tourists. (complex determiners/quantification)

[4]A reasonable reply might be that one may be able to push nonstandard semantic entities from the metalanguage into the object language, e.g., introducing possible-world arguments into predicates and quantifying over them (e.g., [Gawron *et al.*, 1982; Rosenschein and Shieber, 1982]). However, it is unclear how to do this for a logic that quantifies over intensional objects — objects such as functions from possible worlds (or situations) to sets. More to the point, if one *has* a formally interpretable nonstandard logic adequate for NL semantics, and knows how to map into it, what is the point of further mapping this into (undoubtedly much more cumbersome) FOL translations? Apart from short-term expediency (e.g., use of an available FOL theorem prover), there seems to be no good motivation for doing so.

Ross ought to swim home tomorrow. (contingent obligation)
Ross in a bad mood should be avoided. (problematic NP modifier)

None of these examples are particularly outlandish, and further underscore the need to break away from the restrictiveness of FOL, and settle for nothing less than NL-like expressiveness, without retreating to informal representations.

A few systems do already use significantly extended versions of FOL as a representation language. The semantic representation language used in the *SRI Core Language Engine* [Alshawi and van Eijck, 1989; Alshawi, 1990] can express, among other things, event and state variables, (indexical) tense operators, generalized quantifiers, collectives and measure terms, natural kinds, and comparatives and superlatives. TACITUS [Hobbs *et al.*, 1986, 1987] allows for event variables, quantification over predicates, sets, scales, time, spaces and dimension, material, causal connection, force, systems and functionality, etc. (It is also claimed to handle normatives, where a norm is a pattern which is established either by conventional stipulation or by statistical regularity.) But where they go beyond FOL, these systems do not have a formally defined semantics. Also, they still fall short of comprehensive expressiveness; for instance, the representation for both the *Core Language Engine* and TACITUS lacks means to express nominalization, intensional verbs, and generic sentences. As well, the process of mapping syntax to semantics in these systems appears to remain rather *ad hoc* — perhaps necessarily so, since the representation languages have not been defined to make this mapping as direct and simple as possible.

1.2.4 Modularity and Interfacing

We just saw that the majority of the NLU systems are either not capable of representing attitudes, beliefs, actions, plans and goals, or temporal or causal relations or do not provide a formal semantics for their representation. In addition to expressively weak or informal SR/KR, however, NLU systems have suffered from lack of transparency and modularity. CYRUS and BORIS are cases in point.

In these (and many other) systems, a great deal of knowledge about language and about the world are buried in procedures (e.g., procedures which seek semantically appropriate fillers for frame slots) in a way that makes it very hard to determine what linguistic and factual assumptions have been made. Also, in the case of BORIS, it runs as a single module, in which all inferencing, instantiation, and memory searching are invoked as side-effects of a single parsing process. Attached to each lexical item are one or more knowledge structures and associated "demons," which procedurally encode expectations and other information, and the parser produces high-level as well as low-level memory structures. All memory searches, episodic instantiations, and inferencing occur on a word-by-word basis. While of course integration of all the processes that contribute to understanding is desirable at run time, lack of *interface* transparency and proceduralization of semantic knowledge seem to be major obstacles to extending such a system. Only with a modular, explicit, analyzable conceptual architecture and representation, will we be able to expand lexical and world knowledge indefinitely, and attain nontrivial, domain independent understanding.

1.3 Episodic Logic: A Comprehensive Framework for General NLU

The motivation of this thesis was to develop a comprehensive framework for a general NLU system. Hence, the major concerns have been (1) to develop a SR/KR that meets the desiderata discussed in previous section, (2) to develop a translation that maps surface English to such a SR/KR, (3) to develop inference techniques that can be applied to this SR/KR, and (4) to implement these to test the theory.

This is obviously an extremely ambitious program, clearly one that cannot be completed within the framework of one dissertation. But the overriding need for a *comprehensive* framework has never been lost sight of in this work. This has meant that we could not *from the outset* demand absolute hygiene and mathematical perfection from a representational logic which attempts to encompass all the semantic phenomena of NL which have occupied linguists and philosophers for decades, and for which no agreed-upon unified framework exists. Rather, we have taken the risk of initially compromising theoretical rigor and thoroughness in favor of expressive, interpretive, and inferential adequacy. We then progressively refined the framework, gradually firming up the foundations. EL initially had little more than a tentative syntax, ontology, and type structure, and we are only now gaining a better understanding of semantic entailment and finding soundness proofs for some of the inference methods. But, in the meantime, the expressive completeness of EL allowed us to experiment freely with the syntax/logical form interface, tense and aspect deindexing, and inferences based on simple stories.

In taking this somewhat pragmatic view of formal semantics and proof theory, and relying on progressive refinement, we seem to fall onto a curiously unpopulated middle ground between the informal approach and the FOL approach discussed earlier. We realize that this leaves us in a somewhat vulnerable position from the perspective of each of the "subdisciplines" which our work necessarily intersects: grammar, semantic representation, knowledge representation, interpretation, discourse pragmatics, and inference (both at understanding time and at "question answering time"). For instance, we may not meet the most exacting standards of the specialties in some of these areas, or justify each of our theoretical postulates and design decisions with the same exhaustiveness as some work on a much more limited aspect of language understanding. However, what is important about our theory is that it is integrated representation (allowing for the full semantic richness of language), transduction from the former to the latter, deindexing, inference and question-answering.

The centerpiece of the thesis work has turned out to be the semantic representation/knowledge representation called *Episodic Logic* (EL), a highly expressive knowledge representation well-adapted to the interpretive and inferential needs of general NLU. As emphasized earlier, it is the choice of representation which determines how easily we can derive content from surface from, how fully we can capture the semantic nuances of NL text, and how readily we can perform needed inferences. I will now summarize the features of EL and its theoretical and practical role in language understanding, thereby providing an overview of the accomplishments of this thesis.

1.3.1 Meeting the Interlocking Needs of LF-computation, Deindexing, and Inference

Episodic logic is a first order logic with many extensions designed specifically for general NLU, though its power and generality make it suitable for many AI applications. It was Montague-inspired and influenced by situation semantics [Barwise and Perry, 1983; Barwise, 1989]. That is, it is based on a Montague-style coupling between syntactic form and logical form, while incorporating from situation semantics the idea that sentences describe situations (events, states, episodes, circumstances, eventualities, etc.). Most importantly, it meets interlocking needs of LF-computation, deindexing and inference, is formally interpretable and easily derived from surface utterances, yet allows efficient inference. I now briefly describe how EL treats these key aspects, accomplishing its role in a comprehensive, modular approach to NLU.

Expressive and Direct SR/KR

EL serves simultaneously as SR and KR, i.e., it is capable of representing both the explicit content of discourse and the linguistic and world knowledge needed to understand them in a uniform and transparent manner. There are several intermediate forms of the representation, beginning with an "unscoped" form very close to surface structure, continuing with a scoped but still indexical form, and terminating in EL proper, a nonindexical (context-independent) form suitable for storage in permanent memory.

In the process of mapping from an indexical to a nonindexical form, a DRT-like treatment of indefinites is employed, and explicit episode variables are introduced and tense, aspect, and time adverbials interpreted as relations over these variables. The final representation can still be described as "natural-language-like," however. It allows the representation of restricted quantifiers, propositional attitudes, predicate modifiers, nominalized predicates, and perhaps most importantly, unreliable generalizations. Such generalizations have recently received much attention in the non-monotonic reasoning literature and elsewhere (e.g., linguistic semantics).

LF Computation and Deindexing

EL is not only close to surface form but also allows for the relationship between surface form and logical form to be specified in a modular, transparent way. As the grammatical representation, a variant of Generalized Phrase Structure Grammar (GPSG) [Gazdar *et al.*, 1985] has been chosen. GPSG is a particularly perspicuous grammatical formalism which is expressively adequate for almost all English grammatical phenomena, and is relatively easy to use by a parser and logical-form generator. A GPSG-like grammar has been developed for a fragment of English that allows logical forms for EL to be easily computed from surface English. Also developed were a set of "deindexing rules," that make use of utterance contexts and convert indexical logical forms (LFs) into nonindexical episodic logical forms (ELFs). The deindexing rules make use of *tense trees*, a type of context structure that can be viewed as the "fine structure" of discourse. These rules are capable of correctly identifying reference episodes and analyzing the interaction between tense, aspect and temporal adverbials.

Sound and Efficient Inference

The deindexing stage is followed by inference stages which discharge "context-charged" relations (ambiguous relations whose specific meaning depends on the nature — such as aspectual class — of its arguments and on "what makes sense" in the current discourse situation) and more generally do *input-driven* plausible inference based on the discourse interpreted so far, meaning postulates and world knowledge. Very general inference rules, *rule instantiation* and *goal chaining*, in particular, have been developed that allow for deductive and probabilistic inferences, both of which are crucial in narrative understanding and commonsense reasoning. The rules are natural, delivering intuitively warranted conclusions (often combining multiple steps of more standard methods into a single step); the nonprobabilistic versions have been proved sound under certain (not very restrictive) assumptions. The inference rules provide methods of making deductive and probabilistic inferences in both input-driven and goal-driven modes.

1.3.2 Experience with Implementation

What has been accomplished so far seems to vindicate the methodology of working toward a comprehensive framework for NLU. The deindexing algorithm has been successfully implemented and used in the TRAINS domain [Hwang, 1992], producing deindexed episodic logical form. The logic has also been successfully implemented in the EPILOG system [Schaeffer *et al.*, 1991], a hybrid inference system combining efficient storage and access mechanisms, forward and backward chaining, and multiple "specialists" for taxonomies, temporal reasoning, sets, strings, etc.[5] This system has been applied to several domains, proving EL's practicality. It makes quite complex inferences, e.g., with utterances from the TRAINS domain [Allen and Schubert, 1991]; it understands small excerpts from the *Little Red Riding Hood* story [Schubert and Hwang, 1990a]; and it reasons with telex reports for aircraft mechanical problems in the ARMS application, a message processing application for the Boeing Commercial Airplane Reliability and Maintainability Project [Namioka *et al.*, 1991, 1992].

These experiments show that inferencing is straight-forward, despite the richness of the logic, or — I might argue — because of it, and that the knowledge it is based on is uncontrived. It corresponds quite directly to English sentences, and each individual piece of knowledge arguably is formulated at a maximally general level, rather than being particularized to the needs of a specific story.

1.4 Organization of the Thesis

This thesis consists of two parts. Part I presents Episodic Logic, EL, through Chapters 2–5, and Part II discusses the transduction from English to Episodic Logic, with particular emphasis on English tense and aspect, through Chapters 6–9.

In Part I, the emphasis is in EL as a tool for semantic and knowledge representation and inference for NLU, rather than as a logic *per se*. Chapter 2 provides motivation and a preview of EL; it also includes a review of some related previous work. Chapters 3 and 4 present the logical syntax and semantics of EL, respectively. Chapter 5 shows the rules of inference in EL, concluding Part I.

The main purpose of Part II is to show that linguistic input could be mapped into episodic logical representation in a principled and transparent fashion. Chapter 6 sketches the derivation of preliminary, indexical logical form from English surface structure with a GPSG-like grammar. Chapters 7–9 discuss how to deindex indexical logical form. The emphasis here is in *temporal* deindexing. Chapter 7 motivates the deindexing algorithm to be developed by reviewing relevant previous work. Chapter 8 describes the tense tree component of the context structure and the deindexing algorithm that computes non-indexical episodic logical form from indexical logical form using tense trees. Chapter 9 provides some "advanced" deindexing rules and discusses possible ways of extending the deindexing mechanism.

Next, Chapter 10 discusses how EL fits in the general framework for NLU, and reports

[5] Much acknowledgement is due for the rules of inference in EL and their implementation in EPILOG: Len Schubert formulated the inference rules; Stephanie Schaeffer wrote most of the implementation. She also found out that the original rules as presented in [Schubert and Hwang, 1990a] were not adequate, which caused them to be modified. As well, the Boeing Co. not only funded the implementation, but used it in their ARMS message processing application.

the experience with computer implementation, illustrating it with an extended example based on a small fragment of *Little Red Riding Hood*. Chapter 11 assesses the progress made and work still to be done.

Part I

EPISODIC LOGIC

Chapter 2

Situations and Episodic Logic: Motivation and Preview

Episodic Logic, EL, is a first order, situational logic with many extensions designed specifically for natural language processing. The adjective "episodic" is intended to suggest that in narrative texts the focus is on transient types of situations rather than on "eternal" ones.[1] Situations can be used, among other things, as causal antecedents and consequents and as anaphoric referents, which play important roles in AI applications including natural language processing. In this chapter, I first discuss why situations, episodes in particular, need to be taken as individuals in the logical ontology, and review how situations are represented in some of the existing formalisms (including the ones that influenced EL). Then I provide a preview of EL, introducing the characterizing relation between utterances and described episodes, the permissive ontology and the DRT-like parameter mechanism of EL that form the basis its expressiveness, and semantic preliminaries of EL.

2.1 Situations as Discourse Entities

2.1.1 Situations We Live With

We live with situations. We participate in them, with or without realizing it; we perceive them; we talk about them; we laugh at them; we imagine them; we dream of them; we get excited or frustrated at the thought of them. Some situations—like the one in which my computer came back from a crash five minutes ago—are event-like, taking place at a

[1]As will be seen, the word 'situation' (or, interchangeably, 'episode') is used as a generic term that covers states, events, eventualities, worlds (situations that are maximal in terms of time, space, and informational content), etc.

particular time at a particular place, and bring changes into the world. Others — like the one of the economy being in recession — are more enduring, taking place over a prolonged period. Yet other situations — e.g., the one of 7 being a prime number — may be limited in neither time nor space.[2] Some situations — like the one in which my roommate is looking out of the window — are real; others — like the one in which I talk with my great-grandfather who passed away before I was born — are not real, at least in this world I live in, no matter how real it was in my dream last night. Yet some situations may not be physically possible — like the one in which water flows from low places to higher places. Nevertheless, we can imagine and discuss such situations. And sometimes we question whether some kind of situations could exist at all. Finally, some situations — like the one in which I hit the return key (and nothing else) — are "small"; others — like the one in which every citizen in Chicago wears a yellow ribbon — are "large." Some could be *very* large, like Easter Sunday in the northern hemisphere in 1992, all of human history, the world, or reality. All these situations are related with each other in complicated ways. They overlap or meet with each other in time and space in a complicated way. They may have been caused by some other situations, and in turn cause other situations. Some consist of many, often infinitely many, situations, and yet they themselves may be a part of still larger situations.

Narratives are "about" such situations, describing what is going on in the world — real or imaginary — when and where, caused by what situations, and causing what situations. In fact, each utterance of the narrator is also a situation in which the narrator "describes" some situations to the (potential) readers. Hence, a narrative may be regarded as an ordered set of utterances by which the narrator describes a set of situations that compose a larger situation often called an episode, and similarly so for a discourse or a dialog.

One of the main concerns of this thesis is what situations each utterance — or sentence in a narrative — describes, and how they can be represented. Ironically, although situations are at the heart of the meaning of every sentence, they had been widely ignored in traditional logics until recently. Recently, however, many — especially among those working in NLU — have started to incorporate events, facts, situations, eventualities, etc., into their formal ontology. Hobbs [1985b], for example, argued for treating events as individuals, especially in representing sentences with *time* and *place* adverbials such as "John ran on Monday" and "John ran in San Francisco." Treating events, situations and facts as individuals is indeed essential in representing statements involving cause and effect or objects of perception or propositional attitudes. For example, consider the following:

(2.1) a. Mary turned on the radio
       b. *Then*, she picked up the newspaper

[2]In EL, the (factual) proposition that 7 is a prime number is supported by (sufficiently "informed") situations anywhere, anytime.

(2.2) a. Bill kissed Mary
b. John saw *it*
c. *It* made him furious

(2.3) a. John is in love with Mary
b. Everybody knows *that*

In (2.1), the situation in which Mary picked up the newspaper was *right after* the one in which Mary turned on the radio. This temporal relation is easily represented once event variables are allowed in the logical form (although one may use time variables instead in this case). In (2.2 b), *it* refers to the situation, or event, in which Bill kissed Mary, and (arguably) so does *it* in (2.2 c). Event variables seem essential to represent these sentences. Especially the causal relation in (2.2 c) is very hard to capture without using event variables, if not impossible. In (2.3 b), *that* refers to the fact or proposition that John is in love with Mary. Thus, it would be convenient if facts (or propositions) were considered to be individuals.[3]

Webber [1987b] and Schuster [1988] also emphasize the need to treat events (and actions) and event types as individuals in the domain of discourse. Their main argument is that events and event types must be individuals as they are frequently pronominalized. For instance, consider the following example from [Schuster, 1988]:

(2.4) a. John was shot in broad daylight in Philadelphia
b. *It* happened at 10am
c. *It* never happened before

It in (2.4 b) refers to the event described by (2.4 a), while *it* in (2.4 c) refers to an event type found in (2.4 a), namely, the one in which someone was shot in broad daylight in Philadelphia.

In fact, since Barwise and Perry [1983], there has been great interest in situations as individuals both in AI/CL and philosophy, and there seems to be an emerging consensus that situations, events, episodes, facts, etc., need to be included in the ontology as legitimate individuals. This is also one of the underlying positions in EL. In EL, situations, also called *episodes*, form the basis of its model structure. Before getting into EL, I will briefly discuss some of the well-known approaches that have taken events and situations seriously.

[3]This fact or proposition seems closely related to the situation of John being in love with Mary, but I leave the exact connection open for now.

2.1.2 Situations in the Logical Form: A Review of Previous Work

The idea of having event variables as part of the logical form is now well accepted; so, the question is how to incorporate them in the logical form. There have apparently been two schools of thought. One derives from Davidson [1967], who considered an event variable as an extra argument of predicates. This approach has been particularly popular among NLU researchers. The other is one that regards sentences as describing events. This includes Reichenbach's work [1947] and the situation semanticists' approach [Barwise and Perry, 1983]. Here, events are *described by* sentences or formulas, rather than being predicate arguments within those formulas. I will briefly discuss a couple of event representations in both approaches.

Davidsonian Event Variables

Davidson [1967] was a pioneer in advocating the importance of treating events as individuals. He put forward an analysis of action sentences, the basic idea of which is that verbs of action should be construed as containing a place, for singular event terms or variables, that they do not appear to, as in the following.

(2.5) a. Shem kicked Shaun
b. $(\exists e)(\text{Kicked}(\text{Shem}, \text{Shaun}, e))$

One can read (2.5 b) as "There is an event e such that e is a kicking of Shaun by Shem." Davidson's motivation for treating events as legitimate individuals in the ontology was that an adequate theory must give an account of adverbial modification [Davidson, 1970]. Consider, for example,

(2.6) Sebastian strolled through the streets of Bologna at 2 am.

Now to get entailments

a. Sebastian strolled through the streets of Bologna
b. Sebastian strolled at 2 am
c. Sebastian strolled

from (2.6), it needs to be analyzed as

> *There exists an x such that Sebastian strolled x, x took place in the streets of Bologna, and x was going on at 2am.*

Here are slightly more complicated examples of Davidson's and their logical forms obtained with the same kind of analysis.

(2.7) a. I flew my spaceship to the Evening Star
b. $(\exists e)$(Flew (I, My-Spaceship, e) & To (Evening-Star, e))

(2.8) a. Jack fell down at 3pm
b. $(\exists e)$(Fell-down (Jack, e) $\wedge$ $t(e) = 3$)

(2.9) a. Earwicker slept before Shem kicked Shaun
b. $(\exists e1)$(Slept (Earwicker, $e1$) $\wedge$
$(\exists e2)$(Kicked (Shem, Shaun, $e2$) $\wedge$ before ($e1$, $e2$)))

Notice the *before* relation in (2.9 b) that relates two events in terms of time. One can read (2.9 b) as "There exists two events, e_1 and e_2, such that e_1 is Earwicker's sleeping and e_2 is Shem's kicking Shaun, and e_1 was before e_2."

Incidentally, though Davidson was interested in representing temporal relations between events, he concerned himself only with explicit time adverbials or temporal conjunction, completely neglecting tense. (Notice the past form of predicates *kicked, flew, fell,* etc., in the above logical forms.) Also, there was no concern with *how* to get logical form translations from English sentences. Later, Harman [1972], who considered the Davidsonian approach to be advantageous as it minimizes axioms and is compatible with English syntax, separated tense from predicates as shown below:

(2.10) a. John walked in the street
b. $(\exists e)$(walk (John, e) $\wedge$ past(e) $\wedge$ in(e, Street))

(though he failed to specify how to interpret *past*). He also gave a sketch for a possible derivation of sentence (2.10 a), from logical form (2.10 b), using Chomsky-like transformational rules [1972, p. 307], but not vice versa.

Although Davidson was a pioneer in emphasizing the importance of treating events as individuals and had a great influence on NLU (as will be seen shortly), his method of attaching event variables only to atomic predicates is insufficiently general. For instance, the method cannot handle sentences with quantifiers or negation such as "Everyone departed, and this left Mary all alone" or "Mary did not eat for three days. As a result, she was famished." However, such complex situations are frequently encountered in stories and dialogs, often in cause and effect relations. Davidson was apparently unaware of this shortcoming. He even emphasized that his proposal is unique in that in his analysis, sentences like "Shem kicked Shaun" nowhere appear inside his analytic formulas. What he means is that since predicates directly take event variables as arguments in his logical form, there is no need to introduce a modal operator that will embed sentences and event variables as arguments. His concern is that substitution of equals for equals or logical equivalence substitutions would not be allowed if a formula were embedded in a modal context, yet such substitution is clearly truth-preserving in sentences like (2.9). This seems to be an unnecessary worry, however. What Davidson failed to notice was that when a truth-preserving substitution is made in an English sentence, the event or situation implicitly described by the sentence need not (and intuitively often does not) stay the same. Thus while "Shem kicked Shaun" might describe event E, and "Shem kicked Shaun and either kicked Bill or didn't kick Bill" is a logical consequence; this consequence may well describe a different event E', one which (unlike E) involves Bill, for instance, even if only trivially. This is quite different from making a similar substitution in a modal sentence like "Mary *suspects* that Shem kicked Shaun," where the modal operator continues to relate the modified proposition to the *same* individual (viz., Mary).

Hobbs' Nominalization Operator ' ′ '

Under the slogan that the logical form of English sentences should be close to English as well as syntactically simple, Hobbs [1985b] proposes a first-order, nonintensional logical notation. The main idea behind his notation is that it is better to expand one's ontology to allow more kinds of entities than complicating the logical notation, the logical form of sentences, or the semantic translation process. Since he recognized the need for events as arguments of causal and temporal relations and objects of propositional attitudes, he included events as individuals in his ontology.

In his treatment of events, he basically follows Davidson, allowing each predication to have an extra event argument.[4] That is, corresponding to any predicate that can be expressed in natural language, one can say there is an event, or state, or condition, or situation, or "eventuality," in the *world* that it refers to. He uses a "nominalization" operator, "′", to introduce such a reified event or condition into a predicate. He also retains the option of not specifying that extra argument when it is not needed and uses the following axiom schema to relate the two sets of predicates systematically.

$$(\forall x_1, \ldots, x_n)\ p(x_1, \ldots, x_n) \equiv (\exists e)\ Exist(e) \wedge p'(e, x_1, \ldots, x_n)$$

That is, p is true of $x_1, \ldots, x_n$ if and only if there is a condition e of p being true of $x_1, \ldots, x_n$ and e exists in the real world. Thus, corresponding to every n-ary predicate p, there will be an $n+1$-ary predicate p' whose first argument can be thought of as the "condition" that holds when p is true of the subsequent arguments. Notice, however, he uses predicate *Exist* to indicate that its argument entity exists in the actual universe (as opposed to in the Platonic universe of possible individuals that is the domain of quantification). Presumably, this is because he allows *possible* entities in the ontology, including possible events. Nevertheless, it is not clear if "′" can be extensionally interpreted. For example, "to read minds" and "to defy gravity" probably have the same extension, namely,

[4] In contrast to Davidson, he allows not only action predicates but *all* predications to have event arguments.

the empty one, but the corresponding "primed" predicates with "possible events" as additional arguments are presumably distinct. Thus, "′" is not interpretable as a function on extensions. But then the primed predicates must either be directly given extensions under any interpretation I, or "′" must be intensionally interpreted. The former option reduces exactly to the Davidsonian one (i.e., the "′" can only be applied to atomic predicates), while the latter leaves the semantics wide open.

Let us now consider some examples.

(2.11) a. John runs
b. $run(John)$
c. $Exist(E) \wedge run'(E, John)$

(2.12) a. John wants to fly
b. $Exist(E_1) \wedge want'(E_1, John, E_2) \wedge fly'(E_2, John)$

(2.11 a) can be interchangeably represented as in (2.11 b) and (2.11 c). Hobbs suggests reading (2.11 c) as "The condition E of John's running exists in the actual universe," or " 'John runs' is true," or simply "John runs." Note that (2.12 a) may be represented only as in (2.12 b), as the other option of $want(John, fly(John))$ is no longer ordinary first-order logic. Notice that in (2.12 b), $John$ and E_1 are actual, but not E_2. This is because predicate *want* is opaque in its second argument although it is transparent in its first argument.

As (2.12 b) shows, in Hobbs' logical form a natural language sentence is reduced to a *conjunction of atomic predications* in which all variables are existentially quantified with the widest possible scope. Predicates are identical or nearly identical to natural language morphemes. He allows no functions, functionals,[5] nested quantifiers, disjunctions, negations, or modals or intensional operators [1985, p. 62]. Hobbs believes such a restriction is not a problem as his nominalization operator, "′", provides a way of expressing these operators in terms of event predications. For instance, "John is almost a man" could be represented as

$$almost(E) \wedge man'(E, J)$$

instead of $almost(man)(J)$. Note that he treats the adverbial *almost* as a property of events or conditions, where $almost(E)$ presumably means something like "condition E almost exists."

Hobbs treats manner adverbials (and other nonlocative, nontemporal adverbials) as properties of events. The following is another example of his, illustrating his adverbial

[5]Hobbs calls predicate modifiers functionals, i.e., those operators that map predicates into predicates, such as *almost*. In EL, the term "predicate modifier" is used for such operators.

treatment [p. 62].

(2.13) a. A boy wanted to build a boat quickly

b. $(\exists e_1, e_2, e_3, x, y)\ Past(e_1)\ \wedge\ boy(x)\ \wedge\ want'(e_1, x, e_2)\ \wedge\ quick'(e_2, e_3)\ \wedge$
$boat(y)\ \wedge\ build'(e_3, x, y)$

The formula says that there is a boy x, and a past event or state e_1 of the boy's wanting some event or state e_2, where e_2 is the quickness of event e_3 of the boy's building a boat y.[6] Although Hobbs' analysis of manner adverbials seems quite plausible in this example, such an analysis may not work in all cases. For instance, in

(2.14) a. John sold the boat to Mary *reluctantly*

b. $(\exists e_1, e_2, x)\ Past(e_1)\ \wedge\ boat(x)\ \wedge\ sell'(e_1, J, M, x)\ \wedge\ reluctant'(e_2, e_1)$

analyzing *reluctant* as a property of the event of John's selling the boat to Mary does not seem intuitively correct. After all, the same event may be viewed as Mary's buying the boat from John, and Mary might have been very willing! In that case, the same event will be asserted to be both reluctant and willing, giving rise to a contradiction. (Also, "reluctant events" does not sound quite right. Similarly, 'skillful action', but *'skillful event'; 'intentional action', but *'intentional event'; etc.) It seems the problem lies in Hobbs' conflating events and actions. Though some adverbials, especially spatiotemporal ones, may be considered as properties of events, many adverbials seem to be about actions. (In our view, actions are event-agent pairs, i.e., events with well-defined agents.)

Another difficulty with Hobbs' method is that it severely limits the range of sentences that can be represented in the logical form. For example, consider the following sentences.

(2.15) Mary feels uncomfortable when John or Jack is around

(2.16) John did not elude the tackle and score a touchdown, disappointing the fans

(2.17) When everyone congratulated Mary, giving her a present, she felt quite overwhelmed

(2.15) suggests the disjunctive situation of either John's being around or Jack's being around makes Mary uncomfortable. In (2.16), two events — John's eluding the tackle and scoring a touchdown — are explicitly described and temporally related, but it is their *non*-occurrence, rather than their occurrence, which is asserted, and furthermore that

[6] Another, perhaps better, way of reading it would be "e_2 is the eventuality that e_3 is quick." (Incidentally, Hobbs seems to have forgotten $Exist(e_1)$ in (2.13 b).)

non-occurrence is itself a relevant situation here, since it is given as cause of the fans' disappointment. Similarly, in (2.17), one gathers that there is a congratulation episode accompanied by a present-giving episode for each guest (in our reading of the sentence), yet the *when*-clause appears to focus on the larger episode subsuming these individual ones, and to provide it as reason for Mary's feeling overwhelmed. Note that such compound situations — that is, a disjunctive situation in (2.15), a non-occurrence situation in (2.16), and a situation which consists of other smaller ones as parts in (2.17) — cannot be represented in Hobbs' method which associates event variables with atomic predications only. In their later work [Hobbs *et al.*, 1986], one can see that Hobbs and his collaborators have disjunctions in their axioms about commonsense metaphysics. So it is not clear whether they do now allow disjunctive events, and other types of compound events, as well.

Hobbs' work probably was the first in NLU which took events seriously in semantic and knowledge representation. However, unless his "′" is intensionally interpreted, the approach essentially reduces to Davidsonian's, and shares the same problems. Also, the rules of translation are not shown beyond the cursory remark that semantic translation for the logical form is "naively compositional." Yet, Hobbs' move toward "ontological promiscuity instead of complicating the logical notation and the logical form of sentences" was an important one, and has greatly influenced EL.

Schuster's Event/Action Anaphora

I will briefly discuss one more Davidsonian approach as proposed in [Schuster, 1988]. According to Schuster, in general each sentence gives rise to an event,[7] and that an action is a "part" of that event. This seems to be based on her observation that actions are characterized by the predicate of the sentence (VP) and that events are characterized by the whole sentence (S).[8] Thus, in her view, events correspond to the conjunction of action predicates with other predicates describing, for instance, time and place, along with the agent performing the action.

In her semantic representation, Schuster adopts a Davidsonian approach which allows predicates to have an extra event argument. First, actions are represented in terms of lambda predicates as in the following example:

(2.18) a. kiss Mary
b. $\lambda x.[kiss(x, Mary, e)]$

[7]As we will see later, what she really proposes is that each sentence gives rise to a *set* of events. For example, "John kissed Mary" is translated in a way that says "There is a set of events, each of which is of type 'John kissed Mary before now'."

[8]In EL, VPs are in 1-1 correspondence with *kinds* of actions or attributes with undefined agents.

I.e., the action corresponding to kissing event e is identified with the property of being the agent of that kissing event e. (The free variable e gets bound by λ at the level of the complete sentence representation.) Sentences characterize event descriptions which are represented as in the following example:

(2.19) a. John kissed Mary
b. E: $\lambda e.[kiss(John, Mary, e) \ \wedge \ past(e)]$

According to Schuster, what (2.19 b) means is that E is an entity describable as 'the event in which John kissed Mary'.[9]

Schuster's primary concern is not so much with details of logical forms in general, but rather with accounting for the multiple event-related entities which seem to become available for pronominal reference as soon as any event has been described. For instance, (2.19 a) might be followed by any of "*It* happened last night," "*It* never happened to Mary before," "Bill had done *it* the night before," and "*It* happens all the time." Here the referents for *it* can seemingly be particular event of John kissing Mary, *someone* kissing Mary, the action of kissing Mary (or *someone*), and *someone* kissing *someone*, respectively. Accordingly, Schuster abstracts the following 'generalized' events from (2.19 b).

$$E1: \lambda e.[\exists x\, [kiss(x, Mary, e)] \ \wedge past(e)]$$
$$E2: \lambda e.[\exists x\, \exists y\, [kiss(x, y, e)] \ \wedge past(e)]$$
$$E3: \lambda e.[\exists x\, \exists y\, [kiss(x, y, e)]]$$

Here, E1 is the set of events in which someone kissed Mary, E2 is the set of events in which someone kissed someone, and E3 is the set of events in which someone kisses someone (without any temporal location specified). Thus, every member of E1 is a member of E2, and every member of E2 is a member of E3. In other words, there is a partial ordering, E1 $\preceq$ E2 $\preceq$ E3.

There is a slight technical problem here though, since if E1–E3 are sets, then it seems that E in (2.19 b) is also a set, rather than a specific event; viz., it is the set of events in which John kissed Mary. This seems undesirable since the chief purpose of E is to serve as anaphoric antecedent (as in "John kissed Mary. *It* made her angry"). The problem seems to lie in Schuster's trying to use the same kind of representation for both specific events described by sentences and the types of events she calls generalized events. What is needed here is something like a quantificational reading of "E :" (there is a unique event E such that ...), but a nominalization reading of "E1 :" – "E3 :" (the *type* of event such that ...).

[9]One could think of ':' as a predicate or operator meaning "of type" so that 'η: π' means the event η is of type π, assuming she is using a higher-order logic here.

In her subsequent work [Schuster, 1992], she uses Allen's temporal logic [Allen, 1984] to describe both specific and general events. For instance, a specific event in which John kissed Mary (or an action of John's kissing Mary) is represented as

(2.20) $OCCURS(kissed(john, mary),\ t1)$.

On the other hand, general events of type "Someone kissed Mary" and "Someone kisses someone at some time" are represented as

(2.21) $OCCURS(kissed(Agent, mary), t1)$, and

(2.22) $OCCURS(kissed(Agent, Object), Time)$

respectively. Here, *Agent*, *Object* and *Time* are *constant-to-variable* "generalization operators" [p. 84]. Thus, the earlier existential generalizations are recast as variable introduction (without binding). It seems, though, that something like a nominalization operator would still be needed to bind these variables, but it is hard to evaluate the logical aspect of her formalism as it is not developed very far. Somewhat in distinction from her earlier position, she considers *actions* to be special *kinds* of *events* that have agents (rather than *parts* of events), i.e., they are events whose *Agent* slot is filled with a specific constant. Thus, (2.20) represents both an event and an action. This is a dubious position, given the earlier observations about *willing/reluctant *events*. However, Schuster's main contribution lies not in her particular logical proposals, but in her enumeration of the multiple types of entities potentially derivable from a particular event, and available as discourse referents.

Reichenbach's Prescient View on Events and Utterances

It is interesting that before Davidson, Reichenbach [1947] had proposed an analysis of sentences and events that is much closer to the modern situationistic view [Barwise and Perry, 1983]. That is, like situation semanticists, Reichenbach viewed a sentence or a proposition as describing a situation, and introduced function $[\ ]^*$ which, when applied to a sentence, gives a function that takes an event as one of its arguments. This approach is also closely related to that in EL.

Reichenbach believed there are two classes of arguments: the "physical objects" that are arguments of functions and those arguments that determine "space-time locations." For example, the sentence 'John met Jeanne in Hollywood on Tuesday at 8 p.m.' has as arguments the two individuals 'John' and 'Jeanne', the space indication 'Hollywood', and the time indication 'Tuesday at 8 p.m.'; therefore it can be represented as a four-place "thing" function as in (2.23 b).

(2.23) a. John met Jeanne in Hollywood on Tuesday at 8 p.m.
b. meet(John, Jeanne, Hollywood, Tuesday-8pm)

Then he points out that sentence (2.23 a) denotes a fact, and introduces a representation which makes that fact explicit. He uses '$[\ ]^*$' as a fact function such that $[\phi]^*(v)$ means "ϕ describes fact v." That is, (2.23 b) is equivalent to either of the following.[10]

(2.23) c. $(\exists v)$ [meet(John, Jeanne, Hollywood, Tuesday-8pm)]$^*(v)$
d. $(\exists v)$ [meet(John, Jeanne)]$^*(v$, Hollywood, Tuesday-8pm)

In general, a thing function $r(a_1, \ldots, a_n, s, t)$ can be transformed into equivalent fact functions

$$(\exists v)\,[r(a_1, \ldots, a_n, s, t)]^*(v)\,,$$
$$(\exists v)\,[r(a_1, \ldots, a_n, s)]^*(v, t)\,,$$
$$(\exists v)\,[r(a_1, \ldots, a_n)]^*(v, s, t)\,,$$

etc., where v denotes a fact (event, state, episode, situation, and so on).

To see the similarity of Reichenbach's proposal to situation semantic approaches [Barwise and Perry, 1983], compare Reichenbach's formula

$$(\exists v)[r(a_1, \ldots, a_n)]^*(v, s, t)\,,$$

and situation semanticists' representation for a *soa* (state of affairs)

$$v = \langle\langle s, t\rangle, r, a_1, \ldots, a_n\rangle,$$

which reads "a *soa* v is a pair that consists of a location, $\langle s, t\rangle$, and a situation-type, $r, a_1, \ldots, a_n$," or more intuitively, "a *soa* v is a location–situation-type pair that says at location $\langle s, t\rangle$, $a_1, \ldots, a_n$ stand in the relation r."

Unfortunately, Reichenbach equates events and facts [1947, p. 222]. As many critics, e.g., Vendler [1967], have pointed out, this is wrong.[11] Events and facts are different things and cannot be used interchangeably. Most importantly, events take place over a certain

[10]Tense is ignored in (2.23). Once tense is incorporated, it will probably look like
$(\exists v)$ [meet(John, Jeanne, Hollywood, Tuesday-8pm)]$^*(v) \wedge$ past(v),
$(\exists v)$ [meet(John, Jeanne)]$^*(v$, Hollywood, Tuesday-8pm) $\wedge$ past(v),
etc.

[11]Vendler argues that facts and events are not of the same category and thus cannot arbitrarily assume the status of causal antecedents or consequences. In his terminology, 'events' could be any of the following kinds: process, action, condition, state, situation, state of affairs, etc. Slightly differently put, any words that can replace or can be ascribed to perfect nominalizations are events, and similarly, those corresponding to imperfect nominalizations are facts. (Gerunds, e.g., *Mary's beautiful singing of the song*, are taken to be perfect nominalizations, and participle-like clauses, e.g., *Mary's having sung the song beautifully*, and

time interval, and may cause and be caused by other events. In contrast, facts do not happen; they are there at all times; their content is true at all times; they may be the result of some events (or the result of the fact that some events took place); they may enable some events (to take place) and they may be the reason for agents behaving in a certain way (e.g., upon reflecting upon them), but do not themselves "cause" events. Suppose John hit a home run last night, bringing a victory to his team, is an event. Then, the event in which John hit a home run *took place* last night. The event *caused* his team's victory. (Notice that we cannot substitute *facts* here: *The *fact* took place last night; *The *fact* caused his team's victory.) But "that John hit a home run," i.e., *that the event took place*, is a fact. You could have said that last night to your sister, then what you said was a true statement. You could say it tomorrow to your neighbor, and it would still be a true statement. Once an event takes place, the fact that the event has taken place is true at all times, and prior to that, it is true that it *will* take place.

It is true that people often speak loosely of facts, or even agents, as causes. For instance, instead of "The event in which John kissed Mary caused Bill to be angry," one might say "The fact that John kissed Mary caused Bill to be angry," or even "John caused Bill to be angry." Strictly speaking, however, facts and agents cannot be causes. It is sufficient to note that their temporal extents (i.e., time spans) are too large for them to be causes, easily extending beyond the caused event. Facts are abstractions (like propositions) and as such provide explanations, rather than causes. However, they are so closely related to events (e.g., it may be a fact *that an event occurred* or *will occur*) that it is not surprising they occur in causal sentences. To talk of facts as causes just seems to be a way of talking about the "events behind the facts," in somewhat metonymic fashion.[12]

Two important advantages of Reichenbach's analysis over Davidson's are, first, that Reichenbach's fact function can be applied to any kind of sentence, whether it is describing an event, a process, a state, or a fact, while Davidson restricts the event predication to action sentences only, and second, that Reichenbach's fact function, $[\,]^*$, allows quantifiers

---

That-clauses, e.g., *that Mary sang the song beautifully*, are considered imperfect nominalizations.) He then postulates the following three possible cause-effect relations between events and facts:

(A) an *event* is the effect of an *event*
(B) a *fact* is the result of a *fact*
(C) a *fact* is the cause of a *event*

Though much of his argument is convincing, not all of it. For instance, according to him, *Lorelei's beautiful singing* would be a perfect nominal, hence an event, and cannot be a causal antecedent; on the other hand, *that Lorelei sang beautifully* would be an imperfect nominal, hence a fact, and can be a causal antecedent. But "That Lorelei sang beautifully was the cause of the boatmen's death" does sound odd, whereas "Lorelei's beautiful singing was the cause of the boatmen's death" doesn't. At any rate, allowing only facts as causes seems too restrictive.

It seems not all three cause-effect relations Davidson proposes are needed if there is a 1-1 map from events to facts, such that the inverse image of a fact is the event whose having occurred is that fact, and the image of an event is the fact that that event occurred. This is in fact the view taken in EL.

[12]Also note that we loosely say, e.g., "Caffeine causes insomnia," meaning "Caffeine intake causes insomnia."

and other compound sentences within it. Thus, one can easily represent events or situations that involve quantifiers or negations. The following are some examples of my own to illustrate this.

(2.24) a. Every girl danced at the party yesterday
b. $(\exists v)\,[(\forall x)\ \mathrm{girl}(x) \rightarrow \mathrm{dance}(x)]^*(v, \mathrm{Party}, \mathrm{Yesterday})$

(2.25) a. John did not sleep yesterday
b. $(\exists v)\,[\neg\, \mathrm{sleep}(\mathrm{John})]^*(v, \mathrm{Yesterday})$

(2.26) a. John did not meet Jeanne in Hollywood
b. $(\exists v)\,[\neg\, \mathrm{meet}(\mathrm{John}, \mathrm{Jeanne}, \mathrm{Hollywood})]^*(v)$, or
c. $(\exists v)\,[\neg\, \mathrm{meet}(\mathrm{John}, \mathrm{Jeanne})]^*(v, \mathrm{Hollywood})$

(2.27) a. John did not meet Jeanne in Hollywood, but in Boston
b. $(\exists v)\,[\neg\, \mathrm{meet}(\mathrm{John}, \mathrm{Jeanne}, \mathrm{Hollywood}) \wedge \mathrm{meet}(\mathrm{John}, \mathrm{Jeanne}, \mathrm{Boston})]^*(v)$

Notice, however, that the $[\ ]^*$ operator is used rather liberally. That is, a formula may be split rather arbitrarily around it, leading to variable polyadicity. As is well known, variable polyadicity complicates semantics, meaning postulates, and inference. And most importantly, Reichenbach provided no semantics for his operator (which would have been well beyond the "state of the art" at the time!).

Finally, although Reichenbach was much interested in the logical form representation of English sentences, he did not attempt to derive logical forms from English sentences systematically. Indeed this problem remained largely untouched until [Montague, 1970].

Barwise and Perry's Situation Semantics

Since its introduction by Barwise and Perry [1983], situation semantics has increasingly influenced research in computational linguistics and artificial intelligence (cf., [Cooper, 1987; Nakashima *et al.*, 1988]).

Barwise and Perry took as their starting point that, contrary to Frege, sentences refer to (or provide the type of) *situations*, rather than truth values. They saw this move as crucial to a satisfactory analysis of perception sentences. At least at first sight, this view is an even more radical departure from classical semantics than possible-worlds semantics *a la* Kripke or Montague. After all, the latter still regards sentences as referring to truth values, though potentially to different ones in different worlds or at different times. Since a particular world contains everything that could possibly affect the truth value of some sentence, naturally $\Phi \vee \neg\Phi$ is true in every world, and $\Psi \wedge (\Phi \vee \neg\Phi)$ is true in exactly the same worlds as Ψ. Such a notion of sentence denotation is inadequate for both perception sentences and attitude sentences. For perception sentences, we need to avoid adjunction for irrelevant truths such as

John is aware of Molly's being asleep;
**Therefore, John is aware of Molly's being asleep and Jackie's barking or not barking.*

For attitude sentences, we want to avoid substitution of logically equivalent sentences; e.g.,

John believes Molly is asleep;
**Therefore, John believes that Molly is asleep and that all maps are 4-colorable.*

The idea that sentences are about situations gets around these problems since situations only *partially* determine the facts of the world. Thus, $\Psi \wedge (\Phi \vee \neg\Phi)$ may well be about quite a different situation than Ψ alone. The latter may well be silent on the truth or falsity of $\Psi \wedge (\Phi \vee \neg\Phi)$, if Φ falls outside its information content.

Besides arguing that sentences are about situations, situation semanticists also emphasize the *relational* aspect of sentences, since they are after all uttered *in* situations. For instance, they take the meaning of a simple declarative sentence as a relation between utterances and described situations as follows.

$u[\![\text{I AM SITTING}]\!]e$ iff there is a location l and an individual a such that
in u: at l: speaks, a; yes
in e: at l: sits, a; yes

where u is the utterance situation, and e is the described situation. Such situations are built out of three kinds of primitives: individuals, properties and relations, and (spatiotemporal) locations. Having firmly committed to a realist view, however, they do not admit *possible* individuals or *possible* situations. A problem with situation semantics is that the attempt to construct situations out of tuples of individuals, avoiding set-theoretic notions, leads to difficulties with negation and quantifier semantics. For instance, negation was avoided altogether in [Barwise and Perry, 1983], and led to very intricate proposals elsewhere (e.g., [Kratzer, 1989]). This has motivated a new line of theoretical work based on operations (especially replacement) on structures (cf., [Aczel, 1990]). Whatever the metaphysical merits of this approach may be, this seems unnecessary as far as linguistic semantic problems raised by perception and attitude sentences are concerned. As long as we adhere to the *partiality* of the information in situations, we have a handle on those problems. Apart from that, a more or less classical set-theoretic approach seems entirely feasible, and avoids the difficulties with negation and quantifiers (and other operators, such as abstraction).

2.2 Situations and their Characterizations in EL

EL is a first-order, intensional logic, which is probably the most expressive yet brought to bear on the problem of narrative understanding. A distinctive feature of the logic, responsible for its name, is the inclusion of *episodic variables*. These are really *situation* variables in general, but we tend to revert to the "episodic" terminology when discussing logical forms for narrative sentences. Like situation semantics, EL takes sentences as describing situations—events, states of affairs, etc. For example, each of the following sentences describes some present or past (as determined by the verb tense) situation.

(2.28) a. Some lawyers are not honest
b. John was a lawyer

(2.29) a. Gold is metal
b. Whales are mammals

(2.30) a. $2 + 2 = 4$
b. World War I preceded World War II

(2.31) a. Everyone looked at Mary
b. She blushed

(2.32) a. John kicked Pluto
b. It caused him to yelp

To understand the approach to sentence meaning taken in EL (much as in situation semantics), one needs to distinguish two things involved in the meaning of each sentence: the thing referred to or introduced by it (a situation) and the thing used to describe or characterize it (a sentence, or semantically, a sentence intension, or situation type). There are some striking differences in the kinds of situation descriptions or characterizations in the above examples. The descriptions in (2.28)–(2.29) ascribe "enduring" properties to certain entities (lawyers, gold, John, etc.). This is particularly so for (2.29 a) and (2.29 b). It is just barely conceivable (especially if one neglects modern physics and biology) that gold might at some earlier time have been nonmetallic, or whales non-mammalian. But one is inclined to see these properties as holding of their subjects at all times at all places. The examples in (2.30) clearly involve "eternal" or "universal" properties, in that they are true of their subjects whenever, wherever. By contrast, the examples in (2.31)–(2.32) involve *episodic* or *telic* properties, holding of their subjects only now and then, here and there. But despite these clear differences in the kinds of description or characterizations they employ, the sentences uniformly introduce some situation of the specified type. To indicate the relation between these sentences (or their intensions) and the situations they describe, the connective '$**$' is used in EL. '$**$' is an episodic operator, i.e., a modal operator, that connects a formula with the situation it describes. Intuitively, for Φ a formula and η an episodic term, $[\Phi ** \eta]$ means "Φ characterizes (or, completely describes) η." The following are rough representations of the above examples. For the time being, I will use simplified representations, avoiding the details of logical form. Note, however, that in EL, a restricted quantification of form ($\mathsf{Q}\alpha$: $\Phi\Psi$) is used, i.e., ($\exists\alpha$: $\Phi\Psi$) is equivalent to ($\exists\alpha$) $\Phi \wedge \Psi$.

(2.28′) a. ($\exists\, e_1$: [e_1 AT ABOUT Now_1]
[[SOME LAWYERS ARE NOT HONEST] ** e_1])
b. ($\exists\, e_2$: [e_2 BEFORE Now_2]
[[JOHN IS A LAWYER] ** e_2])

(2.29′) a. ($\exists\, e_3$: [e_3 AT ABOUT Now_3]
[[GOLD IS METAL] ** e_3])
b. ($\exists\, e_4$: [e_4 AT ABOUT Now_4]
[[WHALES ARE MAMMALS] ** e_4])

(2.30′) a. ($\exists\, e_5$: [e_5 AT ABOUT Now_5]
[[$2 + 2 = 4$] ** e_5])
b. ($\exists\, e_6$: [e_6 AT ABOUT Now_6]
[[WWI PRECEDES WWII] ** e_6])

(2.31′) a. ($\exists\, e_7$: [e_7 BEFORE Now_7]
[[EVERYONE LOOKS AT MARY] ** e_7])
b. ($\exists\, e_8$: [e_8 BEFORE Now_8]
[[MARY BLUSHES] ** e_8])

(2.32′) a. ($\exists\, e_9$: [e_9 BEFORE Now_9]
[[JOHN KICKS PLUTO] ** e_9])
b. ($\exists\, e_{10}$: [e_{10} BEFORE Now_{10}]
[[[e_9 IS CAUSE OF e_{10}] $\wedge$ [PLUTO YELPS]] ** e_{10}])

Thus (2.28′a) says that e_1 is an episode characterized (or completely described) by "Some lawyers are not honest," and similarly for e_2 in (2.28′b). Note the reduction of the present tense to a relation placing episodes e_1 at about the same time as Now_1, i.e., same time as the utterance of the sentence, in (2.28′a). Similarly, the past tense is reduced to a relation placing episodes e_2 prior to Now_2 in (2.28′b). This reduction is obtained from an initial translation involving indexical operators *pres* and *past* as discussed in Chapter 8. In (2.29′), e_3 and e_4 are situations characterized by gold being metal and whales being mammals, respectively, and so on.

The question now arises whether, and how, the "aspectual" distinctions pointed out among characterizations is reflected in the situations they characterize. The answer is that situations are indelibly "stamped" with the properties of their characterizing descriptions. This is because they support the truth of very little else besides those characterizing descriptions. Consequently, given that there *is* a situation, which supports a certain characterization, it is often possible to infer other situations also supporting it. For instance, in (2.29′a) and (2.29′b), we can infer from the current existence of situations in which gold is a metal and whales are mammals that there are such situations at (virtually) *all times*—in fact, that every situation which decides those issues at all, decides them affirmatively. This is even clearer in the case of (2.30′a) and (2.30′b), since what these say about their arguments is surely true anytime, anywhere, if true at the situations (e_5 and e_6) they introduce.

In (2.31′), e_7 and e_8 are episodes characterized by "Everyone looks at Mary," and "Mary blushes." Notice that (2.31′a) would be further expanded to show *individual* episodes of "person x looking at Mary," for each individual x in the domain, occurring *during* the overall episode e_7. Similarly, e_9 and e_{10} in (2.32′) are episodes characterized by "John kicks Pluto" and "It (John's kicking Pluto) causes Pluto to yelp," respectively. Notice that e_9 is the causal antecedent of e_{10}. Given the episodic or telic aspect of the characterizations in these cases, no generalizations to other times and places follow.

As indicated above, a "characterizing" description of an episode is maximal, or complete, in the sense that it provides *all* the facts that are supported by the episode, except possibly for certain ones entailed by those given. For instance, if $[\Phi \wedge \Psi]$ characterizes η, then η also supports Φ, Ψ, $\neg[\neg\Phi \vee \neg\Psi]$, $[\Phi \vee \Psi]$, and other trivial consequences. Instead of saying that characterizing descriptions are informationally maximal, we could also say that the episodes so characterized are *minimal* with respect to the characterizing description, in the part-of ordering among situations. In other words, no proper *part* of such an episode supports the same description. (This is the view taken in the formal semantics.) Since the *part-of* relation for situations will be construed in such a way that if s' is part of s, then s is spatiotemporally at least as inclusive as s', it follows that a situation s characterized by Φ cannot have a spatiotemporally smaller part that also supports Φ. As will be indicated shortly, the notion of a characterization (or complete description) is important for getting *causal* statements right.

Now, there is a more fundamental episodic operator '$*$', where $[\Phi * \eta]$ means "Φ is true in (or, partially describes) η." '$*$' is essentially an object-language embedding of the semantic notion of *truth in an episode or situation*. Note that $[\Phi ** \eta]$ implies $[\Phi * \eta]$. For example, [[Mary blushes] $**$ e_8] in (2.31′b) entails [[Mary blushes] $*$ e_8]. Also, [[Mary blushes] $*$ e_8] entails the truth of [Mary blushes] in episode e_8. However, in contrast with [[Mary blushes] $**$ e_8], the blushing does not have to extend over the entire episode e_8. (I.e., e_8 need not be spatiotemporally minimal with respect to this description.) And, again in contrast with [[Mary blushes] $**$ e_8], [[Mary blushes] $*$ e_8] does not entail that [Mary blush] constitutes the entire factual content of e_8. Thus, '$*$' is different from '$**$' in at least two respects.

The following meaning postulate relates '$**$' to '$*$'. Φ and Ψ are schema variables over formulas, and η is a schema variable over terms.

$$[\Phi ** \eta] \leftrightarrow [[\Phi * \eta] \wedge (\neg(\exists e\colon [e \text{ proper-subep-of } \eta][\Phi * e]))].$$

Here, proper-subep-of can be read as 'proper-part-of'. What the above means is roughly: if Φ characterizes η, then Φ is true in η, and furthermore Φ cannot be true in any part of η other than η itself.

Whereas the operator '**' is introduced by English sentences as above, '*' is typically introduced by meaning postulates. For instance, (2.32′a) implies that e_9 is a part (in an informational sense) of some bigger episode, say, e_{11}, concurrent with e_9, such that

[[JOHN NEAR PLUTO] * e_{11}],
[[JOHN MOVE HIS LEG TOWARD PLUTO] * e_{11}],
[[JOHN TOUCH PLUTO] * e_{11}],

etc. From a situation semantics perspective, we have roughly the following correspondences:

$[\Phi * s]$	<- - ->	$s \models \Phi$
		(s supports Φ),
$[\Phi ** s]$	<- - ->	s is a minimal situation such that $s \models \Phi$
		(s supports only Φ, and supports it as a whole).

However, note that '*' differs from '$\models$' in that it relates sentence intensions (partial mappings from situations to truth values), rather than 'infons',[13] to situations.

We believe that the notion of a complete description (characterization) of a situation using '**' is crucial for interpreting anaphoric reference to situations and for representing causal relationships among situations. In particular, it appears that an event anaphor like the one in (2.32 b) refers to a *minimal* event described in a prior sentence like (2.32 a). In other words, it would be incorrect to interpret the referent of 'it' in (2.32 b) as simply some event *partially* described by (2.32 a). Rather, it is an event *completely* described by (2.32 a), i.e., a (spatiotemporally as well as informationally) *minimal* event supporting (2.32 a). To see this, imagine that (2.32) is a part of the testimony of John's spiteful neighbor at John's trial for animal abuse. Suppose that the facts as the neighbor knows them are as follows. John is an animal trainer and has been training the big, shaggy Labrador, Pluto, to respond to loud singing by yelping pitifully. On the occasion in question, John did kick Pluto, but only playfully and lightly, to get the dog's attention. At the time, John was singing at the top of his voice, and as soon as Pluto was attending to John, he yelped at the singing as usual. Consider, then, the situation

s: John sang loudly and lightly kicked Pluto to get his attention.

[13] Roughly speaking, *infons*—items of information—are situation types (see [Barwise, 1989; Devlin, 1991]). In EL, situation types are individuals (more specifically, *kinds*), rather than sentence intensions.

The proposition under consideration is whether an "expanded" situation like s can properly be the referent of the anaphor in (2.32b). If it can, the neighbor cannot be accused of lying or even distorting the truth. After all, he may then be correctly attributing Pluto's yelping to John's gentle kick combined with loud singing. This strikes us as incorrect. (2.32b) is simply untrue under the assumed circumstances. It attributes Pluto's yelping to John's kick (at least as primary cause), whereas in fact it was primarily a response to John's singing. A slightly different way of putting the argument is to note that if one witness asserts (2.32a) and (2.32b) while another asserts

(2.32′) a′. John sang loudly
b. *This* made Pluto yelp,

they cannot both be giving the *same* cause for Pluto's yelping, namely, some situation supporting *both* John's kicking Pluto *and* his singing loudly. The notion of a sentence (*completely*) characterizing a described situation avoids these potential confusions. [[JOHN KICK PLUTO] ** e] and [[JOHN SING] ** e] will be inconsistent under any axiomatization of "kicking Pluto" and "singing" which makes them non-synonymous.

To put it schematically, the inference

$$[\Phi ** e_1],\ [\Psi ** e_2],\ [e_1 \text{ CAUSE OF } e_2] \vdash [[\Psi ** e_2] \text{ BECAUSE } [\Phi ** e_1]]$$

is sound, while

$$[\Phi * e_1],\ [\Psi * e_2],\ [e_1 \text{ CAUSE OF } e_2] \vdash [[\Psi * e_2] \text{ BECAUSE } [\Phi * e_1]]$$

is not.

Notice that episodic variables in EL are different from Davidsonian event variables in that they can be "attached" to any formula, whereas Davidsonian ones can be "attached" only to atomic ones. Thus, EL allows for episodes involving quantification, or episodes involving negation, which are not allowed in Davidson's method. As discussed, this is an important feature because such episodes are frequently cited as causal antecedents, anaphorically referred to, quantified over, etc. One may also note that the episodic operator, '**', is similar to Reichenbach's '[]*' in that both map sentences into situation predicates. In fact, an EL formula $[\Phi ** \eta]$ may be written as $[\Phi]^*(\eta)$ in Reichenbach's representation. However, '**' in EL is a two-place function operating on a sentence and an episode, unlike '[ ]*' that may have more than these two arguments, e.g., spatiotemporal individuals as extra arguments, as noted earlier. In EL, spatiotemporal information is represented as a property of episodes as hinted in (2.28′)–(2.32′).

2.3 Other Distinctive Features of EL

Besides its use of episodic (situational) variables and operators '*' and '**', EL has several other distinctive features. One is the incorporation of a DRT-like parameter mechanism; another is its use of an extremely liberal ontology in the semantics; and a third, its use of input-driven and goal-driven inference which can combine many "ordinary" deduction steps into one. I will discuss them in that order.

2.3.1 A DRT-like Parameter Mechanism

One point I did not comment on in discussing (2.32′b) is the *free* occurrence of variable e_9 outside the scope of its quantifier. Such free occurrences are permissible in EL in a sequence of conjoined sentences or in a conditional. The free variables behave much like the "parameters" of discourse representation theory (DRT; cf., [Heim, 1982; Kamp, 1981]). This is accomplished with a new twist on the semantics of '∃', making the variable it binds behave *referentially* if it already has a value under the current interpretation I, and as an ordinary *existential* variable otherwise. In other words, "prior" values of variables *preempt* existential quantifiers (and likewise 'The'). By defining the *parameters* of a formula or text as (roughly) its top-level ∃- (or The-) quantified variables, and allowing these to be preempted through an "external" binding, we obtain a treatment of anaphora and donkey sentences much like that of DRT. For instance, the conjunction of (2.32′a) and (2.32′b), would be interpreted as if the ∃-quantifiers had widest scope. At the same time, we account for the dual existential/referential character of indefinites (cf., [Fodor and Sag, 1982]). (See the semantics of '∃', 'The' and '∧' in Chapter 4).

The parameter mechanism is also the key to representing probabilistic conditionals (or, generic conditionals), such as "A wolf is (usually) gray," "A child (usually) loves his or her grandmother," or "When two strangers meet in a deserted region, they often greet." Probabilistic conditionals take the following form:

$$\Phi \;\rightarrow_{p, x_1, x_2, \ldots, x_k}\; \Psi ,$$

where Φ and Ψ are formulas involving free or ∃-/The-quantified variables $x_1, x_2, \ldots, x_k$, p is a numeric lower bound on frequency (statistical probability). $x_1, \ldots, x_k$ are called *controlled* variables. The semantics of the connective '$\rightarrow_{p, x_1, x_2, \ldots, x_k}$' essentially allows it to "take control" of the controlled variables by iterating over their denotations. ∃ and The are then again "preempted" through the externally supplied values. Thus,

(2.33) a. A wolf is usually gray (or, *Most wolves are gray*)
b. $(\exists x\ [x\ \text{WOLF}]) \rightarrow_{.8,x} [x\ \text{GRAY}]$
c. $[x\ \text{WOLF}] \rightarrow_{.8,x} [x\ \text{GRAY}]$

is interpreted as saying that at least 80% of wolves are gray. Here, (b) and (c) are equivalent. Similarly, to express that when a predator encounters a smaller non-predatory creature, it may attack it, a generalization involving the existentially quantified predator, non-predatory creature, and encounter episode in the antecedent would be used as below.

(2.34) a. When a predatory animal sees a non-predatory creature of comparable or smaller size, it may want to attack and eat it

b. $(\exists x\colon [x \text{ PREDATORY ANIMAL}]$
$(\exists y\colon [[y \text{ NON-PREDATORY CREATURE}] \wedge [y \text{ AS BIG AS OR SMALLER THAN } x]]$
$(\exists e_1 [[x \text{ SEES } y] ** e_1])))$
$\rightarrow_{.6,x,y,e_1} [(\exists e_2\colon [[e_2 \text{ BEGINS DURING } e_1] \wedge [e_1 \text{ CAUSE OF } e_2]]$
$[[x \text{ WANTS TO ATTACK } y] ** e_2]) \wedge$
$(\exists e_3\colon [[e_3 \text{ BEGINS DURING } e_1] \wedge [e_1 \text{ CAUSE OF } e_3]]$
$[[x \text{ WANTS TO EAT } y] ** e_3])]$

Probabilistic conditionals are often used in *causal* axioms. In particular, *predictive* causal axioms assume the occurrence of some particular type of episode e_1 in the antecedent, and predict another episode e_2 caused by e_1 in the consequent. The above is an example of a predictive axiom. Equally important are *explanatory* axioms, such as the following:

(2.35) a. If a creature wants to eat some food, it is likely to be hungry

b. $(\exists x\ (\exists e_1\ [[x \text{ WANTS TO EAT SOME FOOD}] ** e_1]))$
$\rightarrow_{.9,x,e_1} (\exists e_2\colon [[e_2 \text{ CAUSE OF } e_1] \wedge [e_2 \text{ SAME TIME } e_1]]$
$[[x \text{ HUNGRY} * e_2])$

Much of the world knowledge used in our experimentation is in fact stated as causal axioms like these. I will show more examples of probabilistic conditionals in Chapter 3.

2.3.2 A Liberal Ontology and Glimpses of Semantics

Besides being able to describe and relate situations, EL covers a wide range of English constructs. It can represent conjoined predicates by means of lambda abstraction (e.g., *person with a sound mind*); restricted quantifiers (e.g., *every graduate from the West Point* or *some aircraft manufactured by Boeing*); modal operators (e.g., "*Fortunately* John was not there"); attitudes (e.g., "Mary *believes* John is sincere"); perception (e.g., "Mary did not *hear* John leave"); predicate modifiers (e.g., *severe* damage); kinds of things, kinds of actions and kinds of events (e.g., the two kinds of fictitious creatures, *unicorns* and *dragons*; *the Boeing 737*; *failing after five tries*; *for a child to stay quiet*, etc.); facts and

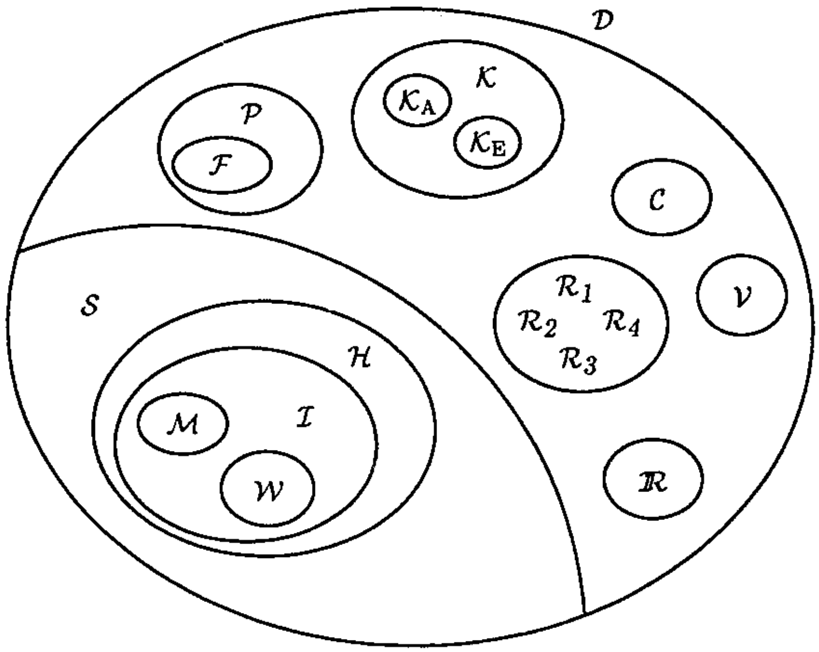

Figure 2.1: Ontology of Basic Individuals

propositions (e.g., "*That Bill follows every pretty girl* is well known" or "John claimed *that he was innocent*"); actions (e.g., *John's sending an open letter to the Dean* in "John sent an open letter to the Dean. They thought *it* was a brave action"); and other non-standard constructs, including the probabilistic conditionals introduced earlier. I will discuss in the next chapter how these are actually represented in **EL**. In this section, I first describe the ontological basis of this wide expressive range.

A Liberal Ontology

Model structures for episodic logic are based on a very liberal ontology of *possible individuals* $\mathcal{D}$. *Possible* individuals are meant to include not only real or actual individuals but also imaginary or fictitious ones. The aim is, like Hobbs', to include in $\mathcal{D}$ everything we can talk about. The reason for allowing "possible" individuals is that ordinary talk abounds with them, as in "Sherlock Holmes is a fictitious detective" and "Today's lecture has been cancelled" (due to Hirst [1991]). Note that the latter sentence, if true, refers to a nonexistent, but possible, event.

Figure 2.1 shows the assumed relations between the basic categories of *possible individuals*, $\mathcal{D}$. Besides "ordinary" individuals, $\mathcal{D}$ includes many unusual types of individuals. First, unlike situation semantics, EL allows *possible* situations $\mathcal{S}$. These are much like "partial possible worlds," in that predicate symbols are assigned partial extensions (and antiextensions) relative to them. Situations, i.e., episodes, occupy times and places, or, more generally, occupy spatiotemporal "trajectories" (regions).[14] Among the possible situations are informationally maximal *exhaustive situations* $\mathcal{H}$, and among the exhaustive situations are the spatially maximal *possible times* $\mathcal{I}$, which in turn includes the spatiotemporally maximal *possible worlds* $\mathcal{W}$ and the spatially maximal, temporally minimal moments of time $\mathcal{M}$.

The treatment of times and worlds as certain kinds of situations is unusual but quite plausible. Consider, for instance, "This week has been eventful," or "The present moment is the outcome of the entire history of the universe," suggesting that times such as *this week* or the *present moment* have episodic content. Times are distinguished from *clock times* in the episodic sense. (*Clock times* and *regions* are explained shortly.)

At this point, it is worth emphasizing that episodes, as the term is construed in EL, subsume events, state of affairs and other situations, circumstances or eventualities. I will generally use the term "events" for episodes in which something *happens,* as opposed to those in which some state or process persists. Note that *actions* or *activities* are not included in this list since, as mentioned earlier, they are not regarded as being of the same type as episodes. Actions are events paired with their agents. (More on this in Chapter 3.)

Disjointly from $\mathcal{S}$, we have not only ordinary individuals of our experience such as people, trees, rocks, clouds, rainbows, and computer networks, but also propositions $\mathcal{P}$, possible facts $\mathcal{F}$ (which we identify with consistent propositions), kinds of individuals $\mathcal{K}$ (including kinds of ordinary individuals, kinds of actions $\mathcal{K}_A$, and kinds of episodes, or situations, $\mathcal{K}_E$), the real numbers $\mathbb{R}$ (augmented with $-\infty$ and $+\infty$), and *1-D* regions $\mathcal{R}_1$ (or just $\mathcal{R}$), *2-D* regions $\mathcal{R}_2$, *3-D* regions $\mathcal{R}_3$, and *4-D* regions $\mathcal{R}_4$, containing subsets of $\mathbb{R}$, $\mathbb{R}^2$, $\mathbb{R}^3$, and $\mathbb{R}^4$, respectively.[15] Tentatively excluded from these are regions with arbitrarily closely packed "holes" and arbitrarily complex boundaries. A "boundary point" of region r is any point such that every open convex region containing it contains a point in r and a point not in r. Then in the one-dimensional case, it can be said that $r \in \mathcal{R}_1$ iff r is closed (contains its boundary points), and any portion of r lying within a finite interval contains at most finitely many boundary points of r (r is a *multi-interval*). Similarly, in

[14]It might be thought that the assignment of regions to all situations precludes the "located/unlocated" distinction in situation semantics. However, in situation semantics, it is not situations, but rather the "infons" they support, that are subject to that distinction. Similarly in EL, the *characterizations* of situations (in terms of sentence intensions or corresponding facts or situation types) are distinguishable in terms of their *persistence* properties. Fully persistent or "eternal" characterizations roughly correspond to unlocated infons.

[15]$\mathcal{R}_4$ is space-time trajectory that may not be connected. A "trajectory" of an episode is given by a total function,

$$Region : \mathcal{S} \rightarrow \mathcal{R}_4.$$

The clock time of an episode can then be expressed as the temporal projection of its trajectory, i.e.,

$$Clocktime(s) = \{w \mid \exists x, y, z \text{ such that } \langle w, x, y, z \rangle \in Region(s)\} \in \mathcal{R}_1.$$

Note that *Clocktime* above is metalinguistic, whereas `clock-time-of` is an *object* language EL functor.

the case of $r \in \mathcal{R}_2$, $r \in \mathcal{R}_3$, or $r \in \mathcal{R}_4$, it is required that r be closed and that no finite straight line cross the boundary of r infinitely often.

The reason for wanting to keep elements of $\mathcal{R}$ relatively "simple" is that they are intended respectively as the temporal and spatiotemporal "projections" of situations. For instance, by saying that the temporal projection of an episode is a multi-interval, we allow for repetitive or quantified events, even ones with infinitely many constituent subevents, but we prevent these subevents from being packed together infinitely densely over any finite time period. Examples of object-language functions whose interpretation relies on $\mathcal{IR}$ and $\mathcal{R}_1$ are `clock-time-of` and `interval`. Clock times of episodes are sets of real numbers (on some absolute time scale) comprising an interval, or, more generally, a multi-interval. The *interval* function provides a convenient way of mapping a pair of calendar times (expressed as 6-tuples) into the corresponding real interval. Thus, for a continuous episode e_1, for example, we may be able to specify its clock time interval by an EL equation like

$$(\text{clock-time-of } e_1) = (\text{interval } (\text{tuple } 1992\ 7\ 1\ 9\ 30\ 45)\ (\text{tuple } 1992\ 7\ 1\ 12\ 5\ 28)).$$

In the two calendar times, the 6 elements represent year, month, day, hour, minute, and second.

Finally, there are collections $\mathcal{C}$ and n-vectors (i.e., tuples) $\mathcal{V}$, $n = 2, 3, \ldots$, of all of these. Kind-, collection- and tuple-formation may be iterated, i.e., there are kinds of kinds, collections of collections, kinds of collections of tuples, kinds of kinds of collections, and so on.

Glimpses of Semantics

I conclude this subsection with some comments on the structure of situations and on the notions of interpretation and truth in EL. A more complete semantics is provided in Chapter 4.

Much as in situation semantics, sentences can be true, false or undefined in situations. At least as far as standard connectives and quantifiers are concerned, EL semantics is more or less standard, e.g., as in [Fenstad *et al.*, 1987], [Barwise, 1989] and [Devlin, 1991]. One difference is that we include standard negation, which is unproblematic in the set-theoretic approach of EL.

The metapredicates $\preceq$ and $\sqsubseteq$ respectively express *coextensive part of* and (general) *part of* relations between episodes. Thus, $s \preceq s'$ is a special case of $s \sqsubseteq s'$, with the additional stipulation that s and s' are coextensive (i.e., occupy the same spatiotemporal region). If $s \preceq s'$, any truth value assigned to a sentence at s is also assigned to it at s'. This is not in general true for $s \sqsubseteq s'$. For instance, the truth of [John popular] at a particular situation s does not guarantee its truth in more comprehensive situations, i.e., one that extends beyond s in time or in location, since John may not be popular when viewed over a longer time stretch or when his popularity is assessed over a wider regions. However, persistence of truth values to more comprehensive situations still holds, e.g., for atemporal/eternally true/unlocated sentences such as [e1 before Now1] or certain episodic/telic ones like [JOHN KICK PLUTO]. As seen later, all tensed English sentences are atemporal/unlocated in their final ELF-translation.

The persistence of truth through the $\preceq$ relation is called "upward persistence," and persistence of truth through the $\sqsubseteq$ relation "outward persistence" (to suggest the expansion of spatiotemporal purview). The reverse persistence of truth through the $\sqsubseteq$ relation, from spatiotemporally "larger" to "smaller" situations (provided the smaller situation settles the issue at all), is called "inward persistence." For example, if

(2.36) JOHN MARRY ANNE

(tenseless) is true in New York State in the first week of March 1992, then it is also true at every location that contains New York State in March 1992, in 1992, etc. But it is not necessarily true in Buffalo or Syracuse, or on March 1, 1992, or on March 2, 1992, etc. That is, JOHN MARRY ANNE is outward, but not inward, persistent. In contrast, if

(2.37) MARY IN-PATIENT

(again tenseless) is true at Strong Hospital at some time, then it is also true everywhere within Strong Hospital during that time, but not outside Strong Hospital, e.g., in the Genessee Hospital or in the city Rochester, or beyond that time (though it would be a different story for MARY IN-PATIENT SOMETIME SOMEWHERE). That is, MARY IN-PATIENT is inward but not outward persistent. Now, if

(2.38) MARY IN-PATIENT AT STRONG HOSPITAL

is true in the first week of March 1992, then it is also true everywhere on March 2, on March 3, etc., but not in 1992, in 1990's, etc. Temporal persistence properties are similar to spatial persistence properties. That is, if

(2.39) MARY IN-PATIENT IN 1ST WEEK OF MARCH 1992

is true in Rochester, then it is true at all times in every place containing Rochester.

However, formulas like

(2.40) JACK MAKE EXACTLY $6,000

do not allow us to draw any inference based on persistence. That is, (2.40) being true in the first week of March 1992 implies neither its being true on March 1, 1992, nor its being true in March 1992.

Factual formulas like

(2.41) WWI PRECEDE WWII

are both outward and inward persistent. Finally, negations of outward persistent sentences, e.g., $\neg$ [JOHN MARRY ANNE], are inward persistent, but negations of inward persistent sentences, e.g., $\neg$ [MARY SICK], are not necessarily outward persistent.[16]

In general, a situation can be part of many worlds, but an "exhaustive" situation (one with maximal "informational content") belongs to a unique world. This is because it supports not only statements about what *is* the case at that time in that location, but also what *was* and *will be* the case in that location, as well as all the unlocated (i.e., eternal) factual statements of a particular world. In other words, a situation that "settles all issues" also settles those pertaining to other times and other places. $\sqsubseteq$ is regarded as a special case of a transitive, reflexive relation $Actual \subset \mathcal{D}\times\mathcal{S}$, determining what individuals are actual with respect to a given situation. Thus, a part of a situation is always actual relative to it. As well, there is a relation $Nonactual \subset \mathcal{D}\times\mathcal{S}$, disjoint from *Actual*, determining the possible but nonactual individuals involved in a situation.

Note that people can talk about individuals that no longer "exist" or do not yet "exist," e.g., "Shakespeare is dead" or "The 6 billionth baby will be born in 1999." Being alive or dead (or having a present *physical* manifestation) have *nothing* to do with an individual

[16] A brief description of aspectual systems is in order. EL makes use of two systems of aspectual classes that are orthogonal to each other. The first consists of factual/stative/telic categories, and the second consists of bounded/unbounded categories.

A factual formula is an unlocated one, i.e., one that is true (or false) at any place at any time. For instance, formulas corresponding to "$2+2=4$" or "John finish*ed* the letter" (with spatiotemporal information incorporated) are factual. A stative or a telic formula may be true (or false) at some place at some time, but may not be so at other place or at other time. Roughly speaking, states and processes are statives, and achievements and accomplishments are telics.

A formula is bounded if it characterizes an episode that terminates in a distinctive state, i.e., characterizes an episode at the end of which there is a state change. Thus, factual formulas are necessarily unbounded, and telics are bounded. Statives may be either bounded or unbounded, however; e.g., Mary's being sick occurs unbounded in "*Mary was sick* when I saw her last time," but bounded in "*Mary was sick* twice last year." An application of some operators often transforms the aspectual class of their operands. I will discuss this further in Chapter 6.

The bounded-unbounded distinction plays an important role in computing the persistence of a formula. In general, unbounded formulas are inward persistent (modulo granularity), and bounded formulas are outward persistent. Polarized ones are exceptional, however. For instance, JACK MAKE AT MOST \$6,000 (negatively polarized) is *not* outward persistent despite its telicity, and JACK MAKE EXACTLY \$6,000 or JACK MAKE \$6,000 IN TOTAL is neither inward nor outward persistent (thus, bipolar).

Finally, for formula Φ, (perfect Φ), (progressive Φ) and (not Φ) are normally taken to be statives. On the other hand, some formulas do not belong to any of the aspectual categories, e.g., a conjunction of stative and telic formulas, such as [[MARY FIND WALLET] $\wedge$ [MARY HAPPY]] (as in "Mary found her wallet and was happy").

being *actual* in the technical sense. Shakespeare is actual in *any* situation in whose informational content he participates, no matter *when* that situation takes places (as long as we are talking about situations in the real world). That is because he was a real person rather than a fictitious one. Sherlock Holmes is nonactual in any situation in whose information content he participates (again, as long as we are talking about real situations). Individuals, logically speaking, have no spatial or temporal boundaries. They just have certain *properties* at certain times and places (more accurately, at certain situations). Some of these properties entail their being alive at those situations (e.g., the property *person*), others do not (e.g., *dead, famous*, etc.).

An *interpretation* I is defined so as to assign elements of $\mathcal{D}$ to all individual constants and to zero or more individual variables, and partial functions of type $\mathcal{D}^n \rightarrow (\mathcal{S} \rightarrow 2)$ to n-place predicate constants (where 2 is the set of truth values $\{0, 1\}$ and $\mathcal{D}^2 \rightarrow \mathrm{A}$ abbreviates $\mathcal{D} \rightarrow (\mathcal{D} \rightarrow \mathrm{A})$, $\mathcal{D}^3 \rightarrow \mathrm{A}$ abbreviates $\mathcal{D} \rightarrow (\mathcal{D} \rightarrow (\mathcal{D} \rightarrow \mathrm{A}))$, etc.). Note that the interpretation of predicates differs from the usual type in that the situational argument comes *last*, rather than first. It turns out that this allows us to dispense with Montagovian intension and extension operators. The syntactic combination of a predicate or function with a term is of course interpreted as function application in the semantics. By the definition of exhaustive situations, sentential truth values under an interpretation I are most fully determined at exhaustive situations $h \in \mathcal{H}$. Note it is possible that even at worlds (and hence times) not all sentences have truth values. Some possible candidates for truth-valueless sentences are (a) sentences with violated presuppositions, (b) vague sentences (admitting "borderline cases"), and (c) paradoxical sentences.

As an illustration of how an interpretation I is extended to a valuation $[\![\,]\!]_I$, the truth value (if any) of a negated sentence is determined at a situation s by

$$\begin{aligned}[\![\neg\Phi]\!]_I^s &= 1 \text{ iff } [\![\Phi]\!]_I^s = 0\,; \\ &= 0 \text{ iff } [\![\Phi]\!]_I^s = 1.\end{aligned}$$

The modal operator '$*$' can be seen to be a truth operator from its semantics (for s an exhaustive situation):

$$\begin{aligned}[\![\Phi * \eta]\!]_I^s &= 1 \;\text{ only if }\; \mathrm{Actual}([\![\eta]\!], s) \;\text{ and }\; [\![\Phi]\!]_I^{[\![\eta]\!]} = 1\,; \\ &= 0 \;\text{ only if }\; \mathrm{Nonactual}([\![\eta]\!], s) \;\text{ or }\; [\![\Phi]\!]_I^{[\![\eta]\!]} \neq 1\,,\end{aligned}$$

where the 'only if's become 'iff's for $s \in \mathcal{H}$ (i.e., s an exhaustive situation). The semantics of '$**$' strengthen the conditions for truth, as will be seen in Chapter 4.

I mention only two further points, both concerning the semantics of quantification. The conception of "quantified episodes" in `EL`, i.e., episodes characterized by quantified sentences, is to some extent novel. An episode *characterized* by a sentence with quantifier $\forall$, `Most`, etc., is the *join* of a set of subepisodes of the type quantified over. The temporal projection of such a quantified episode is typically a multi-interval. For instance, in the sentences, "Every man shaved. *This* took two hours since there were only two razors," the shaving subepisodes together comprise the quantified episode (rather than merely lying within it). Note that by directly providing an episode corresponding to the quantified sentence, the interpretation of the anaphor in the second sentence is unproblematic, and the claim made about overall duration is readily interpretable.[17]

Finally, certain subtleties in the temporal reference of quantified nominals which were previously handled in the semantics (cf., [Schubert and Hwang, 1990a]) are now approached pragmatically. For instance, consider

(2.42) A child grew up here.

The simplest interpretation possible takes the form

$$(\exists e_1{:}\,[e_1 \text{ BEFORE } Now1]\ [(\exists x{:}\,[x \text{ CHILD}]\,[x \text{ GROW UP HERE}])\ {**}\ e_1])\,.$$

In view of the scope of the '**' operator here, the *child* predicate will be interpreted relative to e_1, the "growing up" episode. *Child* is inward persistent, so that x is a child at all times *during* the growing-up episode, but not necessarily at the end points, on the notion of inward persistence adopted in EL.[18] This interpretation seems quite reasonable here, but not in the analogous sentence "A famous composer grew up here." In this case "temporal displacement" between the episode of being a famous composer and that of growing up needs to be allowed. One possibility is to simply give the subject existential wider scope than '**' as below.

$$(\exists e_1{:}\,[e_1 \text{ BEFORE } Now1]\ (\exists x{:}\,[x \text{ FAMOUS COMPOSER}]\,[[x \text{ GROW UP HERE}]\ {**}\ e_1]))$$

For this to be true now, x needs to be a famous composer now (at the time of evaluation). x need not be alive now, since "famous composer" is commonly applied posthumously (e.g., "Mozart is a famous composer"). The two scopings are allowed in EL, thanks to the treatment of tense operators as ambiguously scoped (see Chapter 6). However, even

[17]There are actually two readings, depending on how one measures duration. One way to measure duration is as elapsed time, i.e., the length of time between the end of the last subepisode and the beginning of the first; the other is as "utilized time," i.e., the sum of lengths of shaving subepisodes (assuming they were temporally disjoint). This sort of ambiguity is more apparent in a sentence such as "He graded all the tests; It really took only two hours, but he was at it all day, since he was constantly interrupted by long distance phone calls."

[18]This also explains why "A child grew up" cannot be paraphrased as "An adult grew up," and why "A child lived here for 30 years" seems nonsensical on a normal understanding of the duration of childhood.

scope ambiguity is insufficient to account for the temporal reference of all nominals (cf., [Enç, 1981]). For instance, a reminiscence about someone's boyhood might culminate in the statement "The boy is now an old man." Also sentences with event nominals are apt to involve displacement: "The incident still haunts John after all these years," "The completion of the project is in sight," etc. For these cases, the nominals may be treated as implicitly modified by tense-like, but scoped, predicate modifiers, meaning something like "former," "erstwhile," "present," "future," or "sometime." (Note that these English words *do* function as modifiers of nominals.) The choice of implicit modifier—in the absence of an explicit one—is then a matter of pragmatics.

2.3.3 Inference

Here is a preliminary sketch of inference in EL. One very general inference rule, resembling those used in expert systems, is called *Rule Instantiation* (RI). In some cases, this can be thought of as a general form of *modus ponens* and *modus tollens* with universal instantiation and use of multiple minor premises (instantiating the antecedent and/or consequent of a universally quantified conditional), as in the following example.

$$\frac{\begin{array}{c}(\forall x\colon [[x \text{ FRIENDLY}] \wedge [x \text{ PERSON}]]\ (\exists y\ [y \text{ IS A FRIEND OF } x])),\\ [\text{JACK PERSON}],\quad \neg\,(\exists z\ [z \text{ IS A FRIEND OF JACK}]\end{array}}{\neg\,[\text{JACK FRIENDLY}]}$$

This amounts to inferring "Jack is not friendly," given a rule "Every friendly person has a friend" and a fact "Jack does not have a friend" (and a subsidiary fact that "Jack is a person"). GC (Goal chaining) is the dual of RI, and is mainly for goal-driven inference used in question-answering. Consider the following example:

$$\frac{\begin{array}{c}(\forall x\colon [x \text{ DOG}]\ (\exists y\colon [y \text{ BONE}]\ [x \text{ HAS } y])),\\ ?\,(\exists z\colon [z \text{ TOY}]\ [\text{PLUTO HAS } z])\end{array}}{?\,[[\text{PLUTO DOG}] \wedge (\forall z\colon [z \text{ BONE}]\ [z \text{ TOY}])]}$$

Here, GC reduces the goal (or question) of "Pluto has a toy" into two subgoals of "Pluto is a dog" and "Every bone is a toy," given the knowledge that "Every dog has a bone." RI and GC are formally discussed in Chapter 5.

Note that RI and GC also allow instantiation of probabilistic conditionals such as the predictive and explanatory axioms shown earlier. As a simple example, RI allows the probabilistic inference

$$\frac{(\exists x\,[x \text{ WOLF}]) \rightarrow_{.8,x} [x \text{ GRAY}],\quad [\text{W WOLF}]^{.9}}{[\text{W GRAY}]^{.72}}$$

The superscripted numbers are interpreted as lower bounds on *epistemic* probabilities (in contrast with the *statistical* interpretation of probabilities modifying the connective in probabilistic conditionals). Note that successive inference steps of this type will lead to attenuation of the probabilities assigned to successive conclusions. Such inference chaining—with safeguards against certain fallacies such as circular reasoning—is done routinely in our implementation. A more problematic issue is the "parallel" combination of evidence. This issue arises when several generalizations (or inference chains) assign different epistemic probabilities to the same formula. A partial solution, applicable when the antecedent of one generalization is more specific than the other, i.e., entails the other, is to apply only the more specific rule (*cf.*, [Bacchus, 1988; Bacchus, 1990; Kyburg, 1983]). However, this leaves open the question of how to combine logically independent (or only "probabilistically dependent") bits of evidence. It is hoped that methods similar to those of [Pearl, 1988] may be applied; however, these are aimed essentially at sentential reasoning and so will require considerable extension. At this point, probabilities can be combined serially, but not in parallel (at least not in a principled way).

2.4 Summary

I discussed the need for admitting situations as individuals in the ontology for NLU, and discussed some of the previous works showing how they treated events and situations. I then showed some of the basic features of EL. Namely, in EL, sentences provide characterizations of situations, and '**' is used to indicate the characterization relation such that for sentence Φ and situation η, $[\Phi ** \eta]$ means that Φ characterizes η. The characterization of situations may involve quantifiers, negation, connectives, and other constructs. The ontology employed in EL is very permissive facilitating the interpretation of a large subset of English constructs in EL. I also provided some preliminary glimpses of the intensional semantics of EL, including some examples of truth conditions for ELFs, and discussed briefly how inference is carried out in EL (via general inference rules RI and GC). The following chapters will discuss the syntax, semantics and inference rules of EL in more detail.

Chapter 3

The Logical Syntax of Episodic Logic

One of the distinctive features of Episodic Logic (EL) is its NL-like *expressiveness*, which makes it easy to derive EL-translations from English sentences. In the approach taken in this thesis to interpreting English discourse, the representation of an input sentence is obtained in several — possibly interleaved — processing stages. Initial representations are in general ambiguous in that they may still contain various unscoped operators (e.g., quantifiers and coordinators) and indexical operators (e.g., tense operators). This preliminary unscoped, indexical logical form is called ULF. The next scoped, yet indexical logical form is referred to as LF. The *final* episodic logical form, ELF, is the context-independent representations which are ultimately "committed to memory," ready for use in inferential processes. I will loosely refer to the various levels of representations as logical forms. The emphasis in this chapter, however, is on the ELF. I will first show the syntax of EL, and illustrate it with sample English sentences, which in fact would serve as intuitive motivation for both syntax and the eventual formal semantics. The derivation of the logical form from English input, via various processing stages, is described in Part II of the thesis. Interpretation of ELF representations are discussed in the next chapter.

3.1 The Episodic Logical Form

EL is a first order intensional logic with λ-abstraction, nominalization, and various other extensions. EL syntax allows restricted quantifiers, modal operators, predicate and sentence modifiers, action abstraction, generic conditionals, constructs involving events, actions, facts, kinds, attitudes, donkey anaphora (via a DRT-like parameter mechanism as discussed in Chapter 2), and other non-standard constructs. Before giving the formal syntax, I will first indicate its "flavor" with examples. Additional, less trivial examples are provided in the section following the formal syntax.

In EL, the initial logical form closely follows surface structure, and even the final form is rather English-like. It "mimics" noun phrases in its use of restricted quantifiers, and follows English sentence syntax by having the "subject" of a predication precede the predicate. For example, consider the following.

(3.1) a. John kicked Pluto

b. [John <past kick> Pluto]

c. (past [John kick Pluto])

d. (∃ e1: [e1 before Now1] [[John kick Pluto] ** e1])

Sentence (3.1 a) is initially translated into an unscoped, indexical logical form shown in ULF (3.1 b). After scoping of the 'past' operator,[1] we get LF (3.1 c), which is then deindexed to ELF (3.1 d). In the logical form, square brackets indicate predicate infix expressions; round brackets, prefix expressions; and angle brackets, unscoped operators. Infix notation is used for readability, with the last argument wrapped around to the position preceding the predicate. That is, for π a predicate (interpreted as a "curried" function), $\pi(\alpha)(\beta)$ is written as $[\beta\ \pi\ \alpha]$. In the case of connectives, the conventional way of positioning arguments is observed except that the formula is enclosed in [ ]; e.g., for Φ, Ψ formulas, $\Phi \rightarrow \Psi$ is written as $[\Phi \rightarrow \Psi]$. In the case of the episodic, modal operator '**', $**(\Phi)(\eta)$ is written as $[\Phi ** \eta]$, rather than $[\eta ** \Phi]$, as seen in the matrix of (3.1 d). Also, note in (3.1 d) that restricted quantifiers of form $(\mathtt{Q}\alpha{:}\Phi\Psi)$ are used, where Q is a quantifier, α is a variable, and restriction Φ and matrix Ψ are formulas. That is, $(\forall\alpha{:}\Phi\Psi)$ and $(\exists\alpha{:}\Phi\Psi)$ are equivalent to $(\forall\alpha)[\Phi \rightarrow \Psi]$ and $(\exists\alpha)[\Phi \wedge \Psi]$, respectively. When there is no restriction Φ, we write $(\mathtt{Q}\alpha\,\Psi)$.

As a final point, note that the 'past' operator in (3.1 c) has been reduced to [*e1 before Now1*] in (3.1 d), with certain simplifications. (For example, an "orienting" relation is omitted. Also missing in (3.1 d) is a speech act. These will be discussed in Chapter 8.) *Now1* is a term that denotes the utterance time of (3.1 a). This reduction of 'past', as well as the introduction of episodic variables, is done in the deindexing stage.

At this point, I want to remind the reader that

[[John kick Pluto] ** e1]

implies that e1 is a part (in an informational sense) of some episode e2, coextensive with e1, such that

[1] For scoping of quantifiers and other operators, see, for example, [Schubert and Pelletier, 1982; Hurum and Schubert, 1986; Hurum, 1987].

[[John near Pluto] * e2],
[(∃x: [[x leg] ∧ [x part-of John]] [John move-toward Pluto x]) * e2],
[[John touch Pluto] * e2], etc.

Note that like '**', '*' takes sentential argument first, semantically as well as syntactically.

The following is a slightly more complex (and realistic) example.[2] I omit intermediate logical forms.

(3.2) a. An object bumped into the left wing of Airplane345, causing it to get a crack.

b. (The x: [[x ((attr left) wing)] ∧ [x part-of Airplane345]]
(∃ e_1: [[e_1 before *Now2*] ∧
[(∃y:[y object] [y bump-into x]) ** e_1]]
(∃e_2: [e_1 cause-of e_2]
[(∃z: [z crack] [x get z]) ** e_2])))

Notice in the final ELF (3.2 b) two episodes that have been introduced: e_1, an episode of "some object y bumping into the left wing x," and e_2, an episode of "wing x getting a crack z." Also, notice the clause [e_1 cause-of e_2] which shows the causal relationship between the two episodes. (`attr` is an operator that transforms a predicate into an *attr*ibutive predicate modifier. It will be discussed shortly.)

3.2 The Logical Syntax

3.2.1 Prettifying Rules

The EL syntax allows for both "basic" and "prettified" forms. The basic predicate and function syntax uses prefix form, and is explicitly "curried," i.e., predicates and functions are applied to one argument at a time. The prettified syntax "flattens" the curried notation and infixes predicates and certain 2-place functions, for improved readability. The relation between the flattened syntax and basic syntax is similar to the relation between dot-free lists in Lisp and their dotted-pair equivalents. For instance, formulas [*e1 before Now1*] and [*John kick Pluto*] in (3.1 c) and (3.1 d) shown earlier are the prettified forms of ((*before Now1*) *e1*) and ((*kick Pluto*) *John*).

Since examples of all types of expressions will be given in prettified form, I first state the *prettifying rules*: *flattening rules* and *infixing rules*. The rules give a unique "maximally pretty" result, irrespective of their order of application. However, as with the dot-free notation in Lisp, the prettified syntax may be arbitrarily mixed with the basic syntax

[2]This example is from the ARMS application domain [Namioka *et al.*, 1991; Namioka *et al.*, 1992].

in the specification of axioms and inference rules. (Thus, matching operations used for inference, such as unification, are also assumed to allow for the mixed syntax.)

Flattening Rules

- For π an n-place function or predicate ($n \geq 2$) and $\tau_1, \ldots, \tau_i, \ldots, \tau_k$, terms ($1 \leq i < k \leq n$):

  $((\pi\ \tau_1\ \ldots\ \tau_i)\ \tau_{i+1}\ \ldots\ \tau_k) \Rightarrow (\pi\ \tau_1\ \ldots\ \tau_k)$.

  That is, (*distance x y*) is preferred to ((*distance x*) *y*), where x, y are terms, and (*give Mary Fido*) to ((*give Mary*) *Fido*); by contrast, ((adv-q *extreme*) *intelligent*) cannot be rewritten as (adv-q *extreme intelligent*) since adv-q operates on predicates rather than on terms (adv-q stands for *qualifying adverbial*; this will be discussed shortly).

- For π an n-place predicate and $\tau_1, \ldots, \tau_n$, terms ($n \geq 2$):

  $[\tau_n\ (\pi\ \tau_1\ \ldots\ \tau_{n-1})] \Rightarrow [\tau_n\ \pi\ \tau_1\ \ldots\ \tau_{n-1}]$.

Infixing Rules

- **Predicate infixing.** For π an n-place predicate, and $\tau_1, \ldots, \tau_n$ terms ($n \geq 1$):

  $(\pi\ \tau_1\ \ldots\ \tau_n) \Rightarrow [\tau_n\ \pi\ \tau_1\ \ldots\ \tau_{n-1}]$.

  Note that this convention places the *last* argument of a predicate in "subject" (initial) position. The square brackets are used to make infix formulas distinctive.[3] No syntactic ambiguity arises as long as all atomic symbols are typed.

  Following Reichenbach [1947], we take the sentential argument Φ in $[\Phi\ **\ \eta]$ and $[\Phi * \eta]$ to be first semantically as well as syntactically. Also, the pairing function "$\mid$" and dyadic arithmetic functions may be infixed, with the function symbol following its *first* argument as below.

- **Episodic operator infixing.** For Φ a formula, τ a term, and $op \in \{*, **\}$:

  $(op\ \Phi\ \tau) \Rightarrow [\Phi\ op\ \tau]$.

- **Pairing function infixing.** For τ_1, τ_2 terms:

  $(\mid\ \tau_1\ \tau_2) \Rightarrow [\tau_1\ \mid\ \tau_2]$.

- **Arithmetic function infixing.** For τ_1, τ_2 terms and $\pi \in \{+, -, \times, /, \ldots\}$:

  $(\pi\ \tau_1\ \tau_2) \Rightarrow [\tau_1\ \pi\ \tau_2]$.

3.2.2 The Basic Logical Syntax

This subsection provides the basic, i.e., non-prettified, logical syntax for EL in Backus-Naur form. The constructs provided appear adequate for — and easily computable from — a large subset of English, but it is anticipated to have some future modifications and expansions. In the following, it is assumed that whenever specific atoms are given as options on the RHS of a BNF rule, these are flagged as being of the category on the LHS. In other words, atoms are *typed*. Furthermore, the three continuation dots '...' given as an option always mean "⟨alphanumeric atom, flagged as being of the category of the LHS⟩." Flags (types) of atoms are required to be unique except, possibly, with respect to *function adicity*; i.e., functions of variable adicity (such as +, *max*, *set-of*, etc.) are allowed. Also, note that $n \geq 1$, round brackets indicate prefix notation, square brackets indicate infix notation, and { } indicate optional constituents.

⟨sentence⟩ ::= ⊤ | ⊥ | (⟨1-place-pred⟩ ⟨term⟩) |
(⟨quantifier⟩ ⟨var⟩ {:⟨sentence⟩} ⟨sentence⟩) |
(⟨sentence-op⟩ ⟨sentence⟩) |
[⟨sentence⟩ ⟨episodic-op⟩ ⟨term⟩] |
[⟨sentence⟩$_1$ ⟨n-place-connective⟩ ⟨sentence⟩$_2$ ... ⟨sentence⟩$_n$]

⟨term⟩ ::= ⟨var⟩ | ⟨const⟩ | (⟨1-place-function⟩ ⟨term⟩) |
(⟨pred-nominalization-op⟩ ⟨1-place-pred⟩) |
(⟨sentence-nominalization-op⟩ ⟨sentence⟩)

⟨0-place-pred⟩ ::= ⟨sentence⟩

⟨n-place-pred⟩ ::= ⟨n-place-pred-const⟩ |
(λ⟨var⟩ ⟨($n-1$)-place-pred⟩) |
(⟨($n+1$)-place-pred⟩ ⟨term⟩) |
(⟨n-fold-pred-modifier⟩ ⟨1-place-pred⟩)

⟨n-place-function⟩ ::= ⟨n-place-function-const⟩ |
(⟨($n+1$)-place-function⟩ ⟨term⟩)

⟨var⟩ ::= ⟨alphanumeric atom⟩

⟨const⟩ ::= ⟨alphanumeric atom⟩ | ⟨numeric const⟩ | ⟨quoted string⟩

⟨quantifier⟩ ::= ∀ | Most | Few | No | ∃ | The | ...

[3]However, square brackets are replaced by ordinary round brackets in the current EPILOG implementation.

⟨sentence-op⟩ ::= ¬ | pres | past | futr | perf | prog | □ | perhaps | ...
(adv-e ⟨1-place-pred⟩) | (adv-f ⟨1-place-pred⟩) |
(adv-p ⟨1-place-pred⟩) | ...

⟨1-fold-pred-modifier⟩ ::= ⟨1-fold-pred-modifier-const⟩ | (attr ⟨1-place-pred⟩) |
(nn ⟨1-place-pred⟩) | (na ⟨1-place-pred⟩) |
(adv-a ⟨1-place-pred⟩) | (adv-q ⟨1-place-pred⟩) | ...

⟨2-fold-pred-modifier⟩ ::= ⟨2-fold-pred-modifier-const⟩ | (rel ⟨2-place-pred⟩) | ...

⟨1-place-pred-const⟩ ::= happy | person | certain | probable | ...

⟨1-fold-pred-modifier-const⟩ ::= plur | very | former | almost | in-manner | ...

⟨2-fold-pred-modifier-const⟩ ::= trans | de-nom | ...

⟨episodic-op⟩ ::= ** | *

⟨2-place-connective⟩ ::= ∧ | ∨ | → | $\rightarrow_{<prob><var>_1 \ldots <var>_n}$ | because | ...

⟨n-place-connective⟩ ::= ∧ | ∨

⟨prob⟩ ::= ⟨numeric const, with value between 0 and 1⟩

⟨1-place-function-const⟩ ::= fst | rst | time-of | clock-time-of | − | ...

⟨2-place-function-const⟩ ::= | (“pairing function”) | tuple | + | − | × | interval | ...

⟨$(n+2)$-place-function-const⟩ ::= tuple | + | × | ...

⟨pred-nominalization-op⟩ ::= K | Ka

⟨sentence-nominalization-op⟩ ::= Ke | That | YN-q

Remarks

I now briefly discuss some of the unusual constructs. They will be explained later in more detail with examples.

An n-fold predicate modifier uniformly maps 1-place predicates into n-place predicates, $n \geq 1$. (Tentatively, $n \leq 2$, as well.) 1-fold modifiers, mapping monadic predicates to monadic predicates, are heavily used; 2-fold ones less so. Some derived lexemes are treated this way. For example, “He *sneezed* a hearty sneeze” might involve (*trans* sneeze), where `trans` is a 2-fold predicate modifier; and denominal verbs, as in “He treed the cat” or “He grasses the yard” might involve (*de-nom* tree), etc. Another example might be *measure* terms like *age*, i.e., (*age* year) might give age-in-years, a 2-place relation. `rel` is an operator that maps 2-place predicates such as "see" and "consider" into 2-fold predicate modifiers, which may then apply to 1-place predicates; e.g., ((rel see) swim) or ((rel consider) intelligent) as in "John saw Mary swim" or "John considers Mary intelligent."

Sentence operators include: logical ones like negation; indexical ones like `decl` (indicating declarative speech acts), `pres`, `past` and `perf` (from English tense and aspect), and `futr` (a futural modal operator); modal operators like `prog` (for progressives); and various operators corresponding to those English adverbials that modify sentence meanings. Such adverbials are translated into operators of form (*adv-*δ π), as in the following.

(3.3) a. John saw Mary *in California*
b. (past ((adv-e (in-loc California)) [John see Mary]))

(3.4) a. John *regularly* saw Mary
b. (past ((adv-f regular) [John see Mary]))

`adv-e` and `adv-f` are functions that uniformly map predicates over episodes and sequences of episodes into sentence modifiers. There are also adverbials that operate on predicates such as "with a hammer," "toward the car," and "around the world." These adverbials are typically translated into operators of form (*adv-a* π), where `adv-a` is a function that uniformly maps predicates over actions and attributes into predicate modifiers. Similarly, `adv-p` and `adv-q` are used to translate modal (propositional) adverbials and quality adverbs, e.g., "certainly" and "extremely," respectively. That is, "John was certainly extremely happy" will be

((adv-p certain) (past [John ((adv-q extreme) happy)])).

Predicate modifiers may also be formed by applying a function `attr` to an adjectival predicate. For instance, the NP "white wine" is translated into (K ((*attr white*) *wine*)), where `K` is a kind-forming operator to be explained shortly. na and nn map nominal predicates (i.e., nouns) to adjectival predicate modifiers and nominal predicate modifiers, e.g., ((na sodium) free) for "sodium free" and ((nn car) dealer) for "car dealer."

One limitation of the above syntax is that there are no quantifier modifiers to handle constructs like "almost every," "very few," "nobody else," etc, but this issue will not be pursued here.

In the list of 2-place connectives, the "implication" sign with a probability and a list of (controlled) variables attached allows for probabilistic conditionals, i.e., generic conditionals. Development of a detailed semantics for such conditionals is beyond the scope of this thesis, but I should mention that probabilistic inferences can be made in EL based on generic conditionals. Intuitively, this means in at least a fraction p of cases where the antecedent holds, the consequent holds also (at least this is the *extensional* part of the meaning).

Various function constants (e.g., 'fst', '|', etc.) will be explained later when examples involving them are seen.

A predicate nominalization, such as (*K snow*) or (*K* (*kick Pluto*)), yields a term denoting an abstract individual; in this case, the kind of stuff, snow, or the kind of thing that kicks Pluto (a generic Pluto-kicker, as it were). (*Ka* (*kick Pluto*)) similarly forms an abstract individual, but here it is the *action* — or *attribute* or *property* — of kicking Pluto. Thus, the abstract kind, 'Pluto-kickers', is distinguished from the abstract property, 'kicking Pluto'. And so it should be; for instance, contrast "Pluto-kickers are a pitiful species" with "The property of kicking Pluto is a pitiful species."

Sentence nominalization operators likewise form abstract individuals. Both will be illustrated later, but it should noted here that in EL the individuals formed by That are considered as *propositions*. Such propositions are objects of attitudes, and are *not* situations.

3.3 Illustrating the Syntax; Motivating the Semantics

The syntactic and semantic sketches so far give a general idea of the resources of EL. The following examples illustrate how some of the resources come into play in the representation of anaphora, donkey sentences, attitudes, etc. Note that tense is either omitted or only very roughly translated in this section. Formulas with tense incorporated will be discussed in detail in Chapters 8 and 9.

3.3.1 Restricted Quantifiers

EL "mimics" noun phrases in its use of restricted quantifiers. Consider the following sentences and their translations (minus tense).

(3.5) a. Everyone looked at Mary
b. [<∀ person> look-at Mary]
c. (∀x: [x person] [x look-at Mary])

(3.6) a. Most girls wore a dress
b. [<Most girl> wear <∃ dress>]
c. (Most x: [x girl] (∃y: [y dress] [x wear y]))

(3.7) a. No one left a message
b. [<No person> leave <∃ message>]
c. (No x: [x person] (∃y: [y message] [x leave y]))

(3.8) a. Few things will cheer up Mary

b. [<Few thing> cheer-up Mary]

c. (Few x: [x thing] [x cheer-up Mary])

In the above, (b)-formulas are the initial representations computed from the surface form, i.e., ULFs with quantifier and predicate ambiguities unresolved. (c)-formulas are LFs after predicate disambiguation and quantifier scoping.

Quantifiers may be classified into "nonpreemptable" and "preemptable" ones. ∀, No, Most, Few, etc., are nonpreemptable, and ∃, The, This, etc., are preemptable. Preemptable quantifiers are also called *determiners*. Roughly, the distinction lies in the ability of determiners to behave referentially, as will be seen in the formal semantics. Note that numbers are not taken as quantifiers. Also, it might be worth noting here that some newer versions of situation theory move toward adopting quantified infons that resemble the restricted quantification of EL as shown below:

(i) $\exists x^{\tau(x)}\ \sigma(x)$; $\forall x^{\tau(x)}\ \sigma(x)$ [Barwise, 1989]

(ii) $(\exists \dot{x} \in u)\,\sigma$; $(\forall \dot{x} \in u)\,\sigma$ [Devlin, 1991]

In (i) above, $\tau(x)$ and $\sigma(x)$ are infons dependent on the parameter x; the former acting as a restriction, the latter acting as a matrix. Similarly in (ii), u and σ serve as restriction and matrix (though u is a set rather than an infon). This trend seems to indicate that the need for restricted quantifiers is widely recognized.

As will be seen later in the thesis, if tense is incorporated, the final representation involves the operator '**' relating a sentence to the episode it characterizes. For example, (3.5 b, c) will look like the following (with some simplifications), once tense is incorporated:

(3.5′) b. [<∀ person> <past look-at> Mary]

c. (∃e: [e before *Now*]
    [(∀x: [x person] [x look-at Mary]) ** e])

Notice episode e and its characterization (∀x: [x person] [x look-at Mary]) that are connected by the episodic operator '**'. Note, however, that this formula requires further refinement, especially of the quantifier restriction *person*, if it is to be regarded as literally true. We would not ordinarily interpret (3.5 a) as asserting that "Every person in the world looked at Mary," just as we would not ordinarily interpret "Every general graduated from West Point" as asserting that "Every general in the world graduated from West Point." The intended meaning may just be that "Every U. S. general in recent history graduated from West Point." The salient spatiotemporal frame from the context is essential for correctly interpreting quantifier restrictions. (In fact, every atom in the logical expression potentially requires context-dependent refinement. In other words, the 'predicates' and other atoms of the immediate LF really stand for classes of possible meanings, requiring contextual disambiguation.) After such refinement, (3.5′ c) might look like

(3.5′) d. (∃ e1: [e1 before Now1]
[(∀x: [[x person] ∧ [x in-loc *Room17*]] [x look-at Mary]) ** e1]).[4]

As a final point, the current version of EL does not have an operator that explicitly joins the individuals that satisfy the restriction clause. It appears that use of a set obviates the need for such an operator, as shown in the following example.

(3.9) a. Everyone who wanted to see the actor gathered at the hall

b. (The-least x: [[x set] ∧
[x contains <∀ λy [[y PERSON] ∧ [y WANTS TO SEE THE ACTOR]]>]]
[x GATHERED AT THE HALL])

c. (The-least x: [[x set] ∧ (∀y: [[y PERSON] ∧ [y WANTS TO SEE THE ACTOR]]
[x contains y]))]
[x GATHERED AT THE HALL])

3.3.2 λ-Abstraction

Another feature that leads to close conformity between the surface form and the episodic logical form is λ-abstraction. This is illustrated in the examples below, with tense again neglected. Note that (3.11 b) and (3.12 b) are in unscoped form.

(3.10) a. Canada is very distant from Australia
b. [Canada (very λx[x distant-from Australia])]

Note: This is equivalent to [Canada (very (distant-from Australia))].

(3.11) a. A man with a suitcase
b. <∃ λx[[x man] ∧ [x with-accomp <∃ suitcase>]]>

(3.12) a. the brother of Mary who is a doctor
b. <The λx[[x brother-of Mary] ∧ [x doctor]]>

In (3.10 b), the predicate modifier *very* is a function which, when applied to a predicate, yields another, more restricted predicate.

[4]Strictly, we should use distinct symbols for the "person" property in the original and "refined" LFs, since the former implicitly means "person in Room17," while the latter means simply "person." In [Schubert and Pelletier, 1982], all atoms of the original LF are given *ad hoc* indices to indicate their more or less indeterminate meaning, which still needs to be refined as a function of context.

3.3.3 Predicate Modifiers

Yet another feature that lets episodic logical form look English-like is predicate modification. Examples of predicate modifiers are *almost, fake, former*, complex modifiers like (*attr* interesting), (*nn* truck) (*adv-a* (for-benef Mary)) and (*adv-a* (in-manner abrupt)), as shown below.

(3.13) a. The former president of America visited Russia
b. (The x: [x (former (president-of America))] (past [x visit Russia]))

(3.14) a. Mary watched an interesting movie
b. (past (∃x: [x ((attr interesting) movie)] [Mary watch x]))

(3.15) a. The truck driver disappeared
b. (past (The x: [x ((nn truck) driver)] [x disappear]))

(3.16) a. John bought a fishing boat
b. (past (∃x: [x ((nn fishing) boat)] [John buy x]))

(3.17) a. This drink is alcohol free
b. (pres (This x: [x drink] [x ((na alcohol) free)]))

(3.18) a. John bought the boat for Mary
b. (past (The x: [x boat] [John ((adv-a (for-benef Mary)) (buy x))]))

(3.19) a. Mary abruptly stood up
b. (past [Mary ((adv-a (in-manner abrupt)) stand-up)])

In the above, *former* is a (non-intersective) predicate modifier.[5] `attr` is a function that uniformly maps 1-place (adjectival) predicates, e.g., *interesting* in (3.14 b), into predicate modifiers. nn is a function that maps 1-place (nominal) predicates, e.g., *truck* and *fishing* in (3.15 b) and (3.16 b), into predicate modifiers.[6] `na` is a function that maps 1-place (nominal) predicates, e.g., *alcohol* in (3.17 b), into predicate modifiers that will apply to adjectival predicates like *free*. nn and na are context-charged operators, i.e., they need to be disambiguated with world knowledge, etc. As mentioned, `adv-a` is an operator that forms a predicate modifier from a predicate over actions and attributes.

---

[5]One may recall that this is in contrast with Hobbs' approach in which predicate modifiers are translated as properties of events and conditions. E.g., "Reagan was the former president" would be translated into $former(E) \wedge president'(E, Reagan)$.

[6]nn is analogous to Hobbs *et al.*'s [1988] predicate variable *nn*.

3.3.4 Anaphoric Variables

Consider the following two successive sentences and their logical forms.

(3.20) a. Every man shaved
b. (past (∀ x: [x man][x shave]))
c. (∃ e1: [e1 before Now1] [(∀ x: [x man] [x shave]) ** e1])

(3.21) a. This delayed dinner
b. (past (The y: [y dinner] [*This-thing* delay y]))
c. (∃ e2: [e2 before Now2] (The y: [y dinner] [[e1 delay y] ** e2]))

(3.20 b) and (3.21 b) are preliminary, indexical LFs, and (3.20 c) and (3.21 c) are deindexed ELFs. Notice that *This-thing* in (3.21 b) has been resolved to e_1 in (3.21 c), so that e_1 now occurs *outside* the scope of its ∃-quantifier in (3.20 c). The interpretation of such free variables is much like that of *parameters* in DRT. That is, semantically, this is equivalent to conjoining (3.20 c) and (3.21 c) and widening the scope of ∃-quantifiers. However, in conditional contexts, the interpretation of free variables is different from any reading obtainable by scope-widening (again as in DRT). This is crucial since in such contexts, scope widening does not preserve meaning. For example, in "If I find *a quarter*, I will give *it* to you,"[7]

$$(\exists x\!: [x \text{ quarter}]\, [I \text{ find } x]) \rightarrow [I \text{ give } You\ x]$$

is not equivalent to

$$(\exists x\!: [x \text{ quarter}]\, [[I \text{ find } x] \rightarrow [I \text{ give } You\ x]]).$$

The importance of episodes in causal relations becomes evident if we replace (3.21 a) by "This tied up the bathroom and caused an outburst from Mary."

3.3.5 Donkey sentences

As just discussed, parameters permit a DRT-like treatment of indefinites. Note in (3.22 c) below, the occurrence of variable y that appear outside its scope, as a result of resolving 'It' in (3.22 b).

[7]This example is from [Schubert and Pelletier, 1989].

(3.22) a. Every boy who owned a dog kicked it

b. (past (∀x: [[x boy] ∧ (past (∃y: [y dog][x own y]))] [x kick *It*]))

c. (∃e1: [e1 before Now1]
[(∀ x: [[x boy] ∧ (∃e2: [e2 at-or-before e1] [(∃ y: [y dog] [x own y]) ** e2])]
[x kick y])
** e1])

Here, again, the complete analysis requires the salient context, that is, "every boy in some salient setting that owned a dog" kicked it.

3.3.6 Probabilistic Conditionals (Generic Conditionals)

Closely related to donkey sentences are probabilistic conditionals (extensionally interpretable generic conditionals) such as "A boy who owns a dog usually loves it" or "If a boy owns a dog, he usually loves it." Consider the following examples.

(3.23) a. If a boy owns a dog, he usually loves it

b. $(\exists x\colon [x \text{ boy}]\ (\exists y\colon [y \text{ dog}]\ (\exists e_1\ [[x \text{ own } y] ** e_1])))$
$\rightarrow_{.8,x,y,e_1} (\exists e_2\colon [e_2 \text{ same-time } e_1]\ [[x \text{ love } y] ** e_2])$

(3.24) a. If an aircraft that is less than 3 years old has a crack, usually the crack is not due to corrosion[8]

b. $(\exists x\colon [[x \text{ aircraft}] \wedge (\exists n\colon [n \text{ number}]\ [[x \text{ (age year) } n] \wedge [n < 3]])]$
$(\exists y\colon [y \text{ crack}]\ [y \text{ located-on } x]))$
$\rightarrow_{.8,y} (\neg\ [y \text{ due-to (K corrosion)}])$

In (b)-formulas, '.8' attached to the conditional is a lower bound on the statistical probability, and x, y and e_1 and y are controlled variables. These rules say, roughly, in at least 80% of the situations in which the antecedent is true, the consequent will also be true. Different choices of controlled variables lead to different readings. Suppose, for instance, there are 99 boys owning one dog each and one boy owning 401 dogs. Then, what do we mean by (3.23 a)? The "cases" we have in mind may be the boys (at least 80 boys out of 100) or the dogs (at least 400 of 500) or owning situations (at least 400 out of 500). Note that (3.23 b) provides the third reading, which might not be the most natural. For example, if only the one boy who owns 401 dogs loves them and each of the 99 boys who own a dog does not love his dog, we would be inclined to deny (3.23 a),

[8] As mentioned earlier, *age* in (3.24 b) is a 2-fold predicate modifier that transforms a 1-place predicate (*year*, in this case) into a 2-place relation.

while (3.23b) would be true. (Taking only the boy, x, who happens to be the subject both in the antecedent and the consequent of the original sentences, as a controlled variable will give us a reasonable reading in this particular case.) This is known as the proportion problem and has been the subject of much investigation (*cf.*, initially discussed in [Heim, 1982], and subsequently discussed in, e.g., [Root, 1986; Kadmon, 1987; Schubert and Pelletier, 1989]). The area of generic and habitual sentences is a complex and problematic one, but it appears that for a quantificational conditional sentence, a representation in terms of a probabilistic conditional with control over all existentials in the antecedent that occur anaphorically in the consequent usually leads to intuitively reasonable uncertain inferences. In particular, in usual situations with no disproportionate distribution, this gives reasonable approximations of probability. See Chapter 4 for more discussion and a "first cut" formal semantics.

As discussed in Chapter 2, probabilistic conditionals are often used in causal axioms. Let us consider the previous example (2.34) again, rewritten here as (3.25), this time with its "real" ELF (with rough translation of comparatives).

(3.25) a. When a predatory animal sees a non-predatory creature of comparable or smaller size, it may want to attack and eat it

b. $[(\exists x\!: [x \text{ ((attr predatory) animal)}]$
$(\exists y\!: [[y \text{ creature}] \wedge (\neg [y \text{ predatory}]) \wedge$
$[[y \text{ as-big-as } x] \vee [y \text{ smaller-than } x]]]$
$(\exists e_1 [[x \text{ see } y] ** e_1])))$
$\rightarrow_{.6, x, y, e_1} [(\exists e_2\!: [[(\text{begin-of } e_2) \text{ during } e_1] \wedge [e_1 \text{ cause-of } e_2]]$
$[[x \text{ want (Ka (attack } y))] ** e_2]) \wedge$
$(\exists e_3\!: [[(\text{begin-of } e_3) \text{ during } e_1] \wedge [e_1 \text{ cause-of } e_3]]$
$[[x \text{ want (Ka (eat } y))] ** e_3])]]$

As will be discussed shortly, Ka is an operator that forms a kind of action from an action predicate intension.

3.3.7 Actions

In EL, actions are represented as 'agent-event' pairs, as shown below.

(3.26) a. If a person kicks an animal, *that*'s a wicked action

b. [(∃x: [x person] (∃ e1 [(∃y: [y animal] [x kick y]) ** e1]))
→ [*That-thing* ((attr wicked) action)]]

c. [(∃x: [x person] (∃ e1 [(∃y: [y animal] [x kick y]) ** e1]))
→ [[x | e1] ((attr wicked) action)]]

(3.27) a-1. Mary stepped on John's foot.

2. John thought *it* was intentional.

b-1. (∃ e2: [e2 before Now2]
[(∃x: [[x foot] ∧ [x part-of John]] [Mary step-on x]) ** e2])

2. (∃ e3: [e3 before Now3]
[[John think (That [[Mary | e2] intentional])] ** e3])

Notice that *That-thing* in (3.26 b) is resolved to the ordered pair [x | e1], namely, x's action of kicking an animal, in (3.26 c). '|' is a pairing function applicable to individuals and tuples. (As in Prolog, an individual paired with an n-tuple gives an $(n+1)$-tuple headed by the individual.) In (3.27 a-2), 'it' refers to Mary's action of stepping on John's foot, and is resolved to [Mary | e2] in (3.27 b-2). Note that it was not the *kind* of action, "stepping on someone's foot," that John thought to be intentional, but a specific *instance* of it, performed by Mary as part of event e2.

The agent-event pair representation of actions is motivated by the observation that actions are distinguished from events or episodes in that they have well-defined *agents*. Thus, one may *do* or *perform* an action, but *not* do or perform an episode or event; likewise, as the above examples illustrate, it makes sense to talk about "wicked actions" or "intentional actions," but not "wicked events" or "intentional events." It also seems that the criteria for individuating actions are different from those for individuating episodes. For example, as briefly mentioned in Chapter 2, it seems that (3.28) and (3.29) below may describe the same episode or event (an exchange of a car for a sum of money), but different actions (a buying and a selling).

(3.28) John bought the car from Mary.

(3.29) Mary sold the car to John.

Note, in particular, that the buying in (3.28) may have been performed *reluctantly* and the selling in (3.29) *eagerly*, but it would be very odd to say that the *events* described in (3.28) or (3.29) were reluctant, or eager, or occurred reluctantly or eagerly. Events simply do not have such properties.

As discussed in Chapter 2, our approach is different from the one taken by Davidson, Hobbs, etc., who regard actions as events. But there are also approaches that distinguish actions and events as we do. For instance, Jacobs [1987] regards actions as VIEWs of events, where a VIEW is a structure association used to represent knowledge about concepts that may be used in expressing other concepts. For example, he considers the *transfer-event* concept as related to the concepts *giving-action* and *taking-action* by a VIEW. Although our conception of actions as agent-event pairs is somewhat different from Jacobs', both are based on the intuition that events and actions are different, though closely related.

It appears that specific actions determine specific events (namely, the *performances* of those actions). For instance, when John performed the action of kissing Mary, this performance amounts to a specific event of John kissing Mary. If specific actions determine specific events, and actions have well-defined agents while the events they determine do not, then the simplest possible theory of their relationship is the one illustrated: actions are ordered agent-event pairs. Thus, for an event e which is a performance of an action a by agent x, $a = [x \mid e]$, where '$\mid$' is the pairing function. The agent of an action is then just the first element of the action, written (fst a), and the event of the agent performing the action is the second element of the action, written (rst a). This view turns out to be very helpful in the analysis of adverbials — especially those that modify actions. (More on this later.) Also, as will be seen in Chapter 10, the distinction has resolved some persistent difficulties we encountered in reasoning about actions, such as the "wicked" actions of the wolf in the story of *Little Red Riding Hood.*

I should remark here that while actions, on our account, are agent-event pairs, the converse is not in general true: many agent-event pairs are not actions. For instance, if e is the event of the sun rising, then [John $\mid$ e] is certainly not an action. Whether or not an agent-event pair is an action depends entirely on the characterization of the event, a dependence that is to be captured by meaning postulates.

3.3.8 Kinds or Property Abstraction

Several operators for nominalizing (reifying) sentence or predicate intensions are also available in EL, including a proposition-forming operator That and kind-forming operators K, Ka and Ke. Our approach here owes much to Carlson [1982] and Chierchia and Turner [1988]. We will first consider K, a kind-forming operator, for mapping predicates to individuals. (Ka and Ke will be discussed in the next section.)

(3.30) a. Snow is white
b. [(K snow) white]

(3.31) a. The dog is a mammal
b. [(K dog) mammal]

(3.32) a. Wolves are warm blooded
b. [(K wolf) warm-blooded]

The K operator is used in the interpretation of mass nominals like *snow* and non-numeral bare plurals like NP *wolves* in (3.32 b).[9]

[9] This translation of bare plurals is probably oversimplified. In a refinement still being worked out, the translation of (3.32 a) would be [(K (plur wolf)) (plur warm-blooded)]. Here, plur is an operator that

As an expedient not expected to be retained in the long run, EL syntax also includes an operator K1, for use in the interpretation of indefinite count singulars like *a dog* and bare numeral plurals like *twelve eggs*, as in [Schubert and Pelletier, 1989]. For instance, K1 is used as follows (with rough translations of numbers).

(3.33) a. Twelve eggs cost one dollar
b. [(K1 (twelve egg)) cost (K1 (one dollar))]

The distinction between K and K1 is based on the following contrasts:

An egg is cheap
Eggs are cheap
A dozen eggs (in a carton) are cheap
Cartons of a dozen eggs are cheap
Cartons of a dozen eggs are widespread
*An egg is widespread
*Twelve eggs (in a carton) are widespread
*A dozen eggs (in a carton) are widespread

Eggs are a staple of the American diet
The egg is a staple of the American diet
*An egg is a staple of the American diet
*Two eggs are a staple of the American diet

The "widespread" and "staple" examples suggest that the K operator, while suitable for forming kind-level entities corresponding to non-numeral bare plurals and generic definites, is unsuitable for indefinite a(an)-singulars and numeral plurals. For these we invoke the K1 operator, which forms kind-like entities that somehow fail to have genuine kind-level properties. This leads to simple translations of many sentences involving indefinites, and particular consequences can be derived by meaning postulates. (However, a workable semantics for K1 operator that explains the above contrasts has not been devised yet.) Moreover, if one takes "An egg is cheap" or "Twelve eggs cost one dollar" as analogous to "A cat lands on its feet," or "A cat hates a dog," etc., then one would like to treat such sentences as generic conditionals. Thus, one would like to represent sentences like (3.33 a) as probabilistic conditionals rather than as kind-level predications, as in the following sort of way (with certain simplifications in representing numbers and plurals):

(3.33) c. $(\exists x\,[x \text{ (twelve egg)}]) \rightarrow_{.8,x} [x \text{ cost (K1 dollar)}]$

uniformly maps predicates applicable to (non-collective) individuals to predicates applicable to collections. I.e., (plur P) is true of a collection just in case P is true of each member (cf., [Link, 1983, 1987]). See [Hwang and Schubert, 1992b] for details.

But note that K1 is still required in the translation of "a dollar," since an existential translation of the VP "cost one dollar" into

$$\lambda x\,(\exists y\!:[y \text{ (one dollar)}]\ [x \text{ cost } y])$$

seems wrong (amounting to "costs some dollar or other"). It appears that a resolution of these difficulties will require a subtler approach (cf., [Krifka, 1991]). Here we retain the simplicity of unitary logical forms, while setting aside the unsatisfactory analysis of kinds and generic sentences.

3.3.9 Kinds of Actions and Events

The two remaining kind-forming operators are Ka and Ke, forming kinds of actions (more generally, *attributes*) and kinds of events, respectively. Ka transforms a 1-place predicate into a *kind of action* or *attribute*; and Ke transforms a formula (semantically, a sentence intension) into a *kind of event/situation*. That is, Ka in (3.34–3.38) below is a property forming operator that maps monadic predicate intensions to abstract types of *a*ctions and *a*ttributes. Ke in (3.39–3.40) below is a sentence nominalization operator, which forms (reified) types of *e*vents from sentence intensions.

(3.34) a. Mary loves *skiing*
b. (pres [Mary love (Ka ski)])
c. (∃ e1: [e1 at-about Now1] [[Mary love (Ka ski)] ** e1])

(3.35) a. *To kiss Mary* was fun
b. (past [(Ka (kiss Mary)) fun])
c. (∃ e1: [e1 before Now2] [[(Ka (kiss Mary)) fun] ** e1])

(3.36) a. John likes *to kick Pluto*
b. (pres [John like (Ka (kick Pluto))])
c. (∃ e1: [e1 at-about Now3] [[John like (Ka (kick Pluto))] ** e1])

(3.37) a. The boy wanted *to eat an apple*
b. (The x: [x boy] (past [x want (Ka λy($\exists z$: [z apple] [y eat z]))]))
c. (The x: [x boy] ($\exists e_1$: [e_1 before Now4]
[[x want (Ka λy($\exists z$: [z apple] [y eat z]))] ** e_1]))

(3.38) a. Everyone wishes *to be happy*
b. (past ($\forall x$: [x person] [x wish (Ka happy)]))
c. ($\exists e_1$: [e1 same-time Now5] [($\forall x$: [x person] [x wish (Ka happy)]) ** e_1])

(3.39) a. *For Mary to dance* was rare

b. (past [(Ke [Mary dance]) rare])

c. (∃ e1: [e1 before Now6] [[(Ke [Mary dance]) rare] ** e1])

(3.40) a. Mary suggested *that John leave*

b. (past [Mary suggest (Ke [John leave])])

c. (∃ e1: [e1 before Now7] [[Mary suggest (Ke [John leave])] ** e1])

In the above, "skiing," "to kiss Mary," "to kick Pluto," etc., are kinds of *actions*, while "to be happy" is a kind of *attribute*. Thus, for example, (3.34 a) says that Mary loves that kind of action ("skiing"), and (3.35 a) says that the kind of action "kissing Mary" is fun; while (3.38 a) says that the attribute of being happy, or happiness, is wished by everyone.

Next, "for Mary to dance" is a kind of *event*, and (3.39 a) asserts that this kind of event is rare.[10] In (3.40), what Mary suggested is the kind of event in which John leaves. Incidentally, to be more accurate, it may be necessary to use generic tense `gpres` or `gpast` in (3.34–3.36) and (3.38–3.39) to capture the generic temporal reference. Note that `Ka` and `Ke` always apply to predicates and sentence intensions *without tense*, though possibly with aspect (e.g., " 'Tis better to *have* loved and lost than never have loved before").

Note that formulas (3.36 c) and (3.39 c) are equivalent, by definition of `Ka` and `Ke`, to (3.41) and (3.42) below, which involve the more basic kind-forming operator `K`. As mentioned in Section 3.3.7, `fst` and `rst` are operators that, when applied to a pair or tuple, pick the "first" element and "rest" of the pair or tuple, respectively.

(3.41) (∃e1: [e1 at-about Now1]
[[John like (K λa[[(fst a) kick Pluto] ** (rst a)])] ** e1])

(3.42) (∃e1: [e1 at-about Now1]
[[(K λe[[Mary dance] ** e]) rare] ** e1])

However, note that not all the infinitives may be translated into kinds of actions/attributes. Consider the following examples.

(3.43) a. John made Mary leave

b. (past [John ((rel make) leave) Mary])

(3.44) a. John helped Mary leave

b. (past [John ((rel help) leave) Mary])

[10] Note that the *for–to* construct is semantically *ambiguous* between a kind-of-event reading and a kind-of-action reading. For instance, in "*For Mary to buy things from John* was a mistake," the *for–to* construction need be interpreted as denoting a kind of action (or, more precisely, realization of a kind of action), not a kind of event. Otherwise, we could end up interpreting the sentence meaning the same thing as "*For John to sell things to Mary* was a mistake."

(3.45) a. John considers Mary (to be) intelligent
b. (past [John ((rel consider) intelligent) Mary])

(3.46) a. John found the story (to be) convincing
b. (past [John ((rel find) convincing) Story])

In the (b)-formulas above, (*rel* π), where π is a 2-place predicate, is a 2-fold predicate modifier that uniformly transforms monadic predicates into relational predicates. Note that in the above examples, the infinitival complements cannot be naturally replaced by accusative or dative NPs, i.e., object-like NPs, and are not translated into kinds. Instead, they form predicate-like phrases together with main verbs. One may want to compare the similarity of the above constructs, (3.43)–(3.46), with the following:

(3.47) a. They elected John president
b. (past [*They* ((rel elect) president) John])

Also, infinitival or participial complements of perception verbs, e.g., "John saw Mary *swim*" or "John saw Mary *swimming*," seem to call for the same analysis as the above. This will be discussed shortly in Section 3.3.12.

As a final point, note that although most verbs take both infinitives and gerunds interchangeably as objects, some do not. For instance, "John remembered asking Mary for help" and "John remembered to ask Mary for help" are different. The former says John remembered a specific episode (or action) of his asking Mary for help at some earlier time; the latter says John remembered he was supposed to do an action of the kind, "ask Mary for help." However, in EL, both are translated uniformly as kinds of actions/attributes. The distinction may then be made by using different predicates for main verbs, e.g.,

(3.48) a. John remembered asking Mary for help
b. (past [John remember_1 (Ka (ask-for Mary (K help)))])

(3.49) a. John remembered to ask Mary for help
b. (past [John remember_2 (Ka (ask-for Mary (K help)))])

Later, by meaning postulates, one gets from (3.48b) the inference that John asked Mary for help at some earlier time.

3.3.10 Attitudes and Propositions

Note the proposition-denoting terms headed by That below.

(3.50) a. *That Mary smoked* cannot be true

b. (pres (cannot [(That (past [Mary smoke])) true])),

with a rough translation of "cannot"

(3.51) a. *That Jack went after every pretty woman* is well-known

b. (pres [(That (past ($\forall x$: [x ((attr pretty) woman)] [Jack go-after x])))
well-known])

(3.52) a. Mary knows *that John is intelligent*

b. (pres [Mary know (That (pres [John intelligent]))])

(3.53) a. Mary told Jack *that John kicked Pluto*

b. (past [Mary tell Jack (That (past [John kick Pluto]))])

(3.54) a. The inspector believed *Crack11 indicated danger*

b. (past (The x: [x inspector]
[x believe (That (past [Crack11 indicate (K danger)]))]))

As mentioned earlier, in EL, the objects of attitudes are considered propositions, not situations.[11] Recall that we take propositions as subsuming possible facts. Possible facts are just consistent propositions — there are self-contradictory propositions (and these may, for instance, be objects of beliefs, etc.), but there are no self-contradictory possible facts.[12] We also assume that modal verbs with *that*-complements are always interpreted as if they contained an implicit "to be true." Thus, for example, "John discovered that Mary left" is interpreted as "John discovered it to be true that Mary left." Note that "the fact that Φ" is *not* a *that*-complement, but an ordinary NP.[13]

---

[11]This is in contrast with Hobbs' approach, which takes events as objects of attitudes. For instance, (3.53 a) is translated into

$tell'(E1, Mary, Jack, E2) \wedge kick'(E2, John, Pluto)$,

i.e., $E2$ is John's Pluto-kicking event, and Mary tells Jack the event $E2$. Similarly, in his method, one *knows/believes/denies* events. This does not seem to agree with our intuition, however. On the other hand, in situation semantics, the objects of attitudes are the support relation between situations and infons. For instance, (3.53 a) will be represented as something like

(s $\models$ [Mary tell Jack (s' $\models$ [John kick Pluto])]).

This representation is not very distant from the one in EL.

[12]That is, there are no situations supporting an inconsistent proposition that says, e.g., a person is a high school graduate, and yet is not one. But even such inconsistent propositions exist as individuals, as actual individuals in all the worlds. They are after all derived (by nominalization) from well-defined mathematical objects such as $[\![\top \wedge \bot]\!]$. This is a partial function on situations *which is false whenever it has a value*. ($[\![\,]\!]$ is a valuation function discussed in Chapter 4.) But since it need not be *defined* in exactly the situations where other logical falsehoods are defined, it can have a different intension from them. Thus, $[\![$(That $[\top \wedge \bot]$)$]\!]$ can likewise be a different object from, say, $[\![$(That $[A \neq A]$)$]\!]$, etc. This is *very* different from, e.g., "the round square," or from "the number which equals both 0 and 1," which *are* truly inconceivable (on the usual understanding of the predicates involved).

[13]Sentences like "John discovered the fact that Φ" or "That Φ is a fact" may be initially translated as follows (minus tense):

Note that in contrast to Ke that applies to untensed sentences (see Section 3.3.9), That always applies to tensed sentences. Thus, for example, (3.53 b) and (3.54 b) will look like the following (with certain simplifications) once tense operators are deindexed:

(3.53) c. ($\exists e_1$: [e_1 before Now1]
[[Mary tell Jack (That ($\exists e_2$: [e_2 before e_1] [[John kick Pluto] ** e_2]))]
** e_1])

(3.54) c. ($\exists e_3$: [e_3 before Now2]
[(The x: [x inspector]
[x believe (That
($\exists e_4$: [e_4 at-about e_3] [[Crack11 indicate (K danger)] ** e_4]))])
** e_3])

In addition to Ke and That, EL syntax will include other nominalization operators, Answer-to, YN-q, etc., in the future, as in the following (with tense neglected).

(3.55) a. John knows whether Mary loves Bill
b. [John know (Answer-to (YN-q [Mary love John]))]

(3.56) a. John knows who admires whom
b. [John know (Answer-to (Wh x (Wh y [x admire y])))]

(YN-q is from *yes-no-question.*) There is no firm semantics for these operators yet.[14]

3.3.11 Modal Operators

I consider only modal adverbs and phrasal modal operators here. Some modal adverbs are decomposable into an adjective and a suffix, e.g., *probably, certainly*, etc., while some are undecomposable, e.g., *perhaps, likely*, etc. (Sentential modal operators, e.g., "As Mary suspected," are not considered in this thesis; futural modal operator *will* is discussed in Part II of the thesis.) Here are some examples.

(3.57) a. Perhaps Pluto is sick
b. (perhaps (pres [Pluto sick]))
c. (perhaps ($\exists e_1$: [e_1 at-about Now_1] [[Pluto sick] ** e_1]))

(3.58) a. John certainly loves Mary
b. (((adv-p) certain) (pres [John love Mary]))
c. [(That ($\exists e_2$: [e_2 at-about Now_2] [[John love Mary] ** e_2])) certain]

[John discover <The λx[[x fact] $\wedge$ [x = (That Φ)]]>]; [(That Φ) fact].

[14] Some relevant work on this is [Ginzburg, 1991].

(3.59) a. According to Jack, the inspector detected the leak

b. (((adv-p) (according-to Jack))
(past (The x: [x inspector] (The y: [y leak] [x detect y]))))

c. [(That ($\exists e_3$: [e_3 before Now_3]
[(The x: [x inspector] (The y: [y leak] [x detect y])) ** e_3]))
according-to Jack]

(b)-formulas above are indexical LFs; (c)-formulas are deindexed ELFs. Note that "Certainly Φ" is treated equivalent to "It is certain that Φ" or "That Φ is certain." And similarly so for "according to Jack."

3.3.12 NI Perception Statements

We treat NI perception verbs as operators that uniformly map monadic predicates into relational predicates as illustrated below.

(3.60) a. Mary heard Pluto yelp

b. (past [Mary ((rel hear) yelp) Pluto])

c. (∃e1: [e1 before Now1]
[[Mary ((rel hear) yelp) Pluto] ** e1])

(3.61) a. Mary saw John kick Pluto

b. (past [Mary ((rel see) (kick Pluto)) John])

c. (∃e2: [e2 before Now2]
[[Mary ((rel see) (kick Pluto)) John] ** e2])

(3.62) a. Pluto smelled something burning

b. (past (∃x: [x thing] [Pluto ((rel smell) λy(prog [y burn])) x]))

c. (∃e3: [e3 before Now3]
[(∃x: [x thing] [Pluto ((rel smell) λy(prog [y burn])) x]) ** e3])

We saw in Section 3.3.9 that `rel` is an operator that forms a 2-fold predicate modifier from a 1-place predicate. That is, it transforms 1-place predicates — *hear, see,* and *smell,* above — into a function that uniformly maps 1-place predicates — *yelp*, (*kick Pluto*), and λy(*prog* [y *burn*]), above — into relational predicates corresponding to *hear yelp, see kick Pluto,* and *smell burn* that take as arguments Pluto and Mary, John and Mary, and something x and Pluto, respectively. Thus, for example, (3.62 b, c) simply say that the relation "smell burning" held between something and Pluto. This is in contrast with the approach taken by situation semanticists, e.g., [Barwise, 1989], who would analyze (3.62 a) as "Pluto smelled a situation in which something was burning," which is intuitively unnatural. This analysis may also be applied to perception verbs with adjectival complements. For instance, "John saw Mary asleep," "John saw Mary drunk," etc., may be translated into [John ((rel see) asleep) Mary], [John ((rel see) drunk) Mary], etc.[15]

One may question the analysis described above, on the grounds that (3.60 a) may be followed by a sentence referring anaphorically to the event or action described. For example, (3.60 a) may be followed by

(3.60′) Bill heard *it*, too

indicating that the object of *heard* in both (3.60 a) and (3.60′) must be an individual, i.e., a situation or action (Pluto's yelp or yelping). However, in the first place, (3.60 a) and (3.60′) involve distinct verb subcategorization patterns: in (3.60 a) there are NP and VP complements, while in (3.60′) there is only an NP′ object. Thus, even if (3.60′) involves a situational (or action) object (a yelp or yelping of Pluto), we should not expect to find the same object in the LF of (3.60 a). Furthermore, it is commonly held that referents of anaphoric pronouns need not appear explicitly in the LFs of the preceding discourse. Rather, it appears that anaphoric referents are often obtained by inferential or constructive processes from the prior discourse, as the following examples illustrate.

a. John did the dishes. Mary didn't want to do *it*.
b. The three boys each ordered a large anchovy pizza.
Because of the heavy traffic, *they* were delivered cold.

In a, the antecedent of 'it' is *doing the dishes*, a kind of action which may not be located directly in the LF of the first sentence. In *b*, the antecedent of 'they' is the collection of three pizzas ordered by the three boys, which in all likelihood is not an explicit constituent of the LF of the first sentence. The latter example is due to Webber who discussed this and many related phenomena in discourse anaphora [1978; 1983; 1991].

Finally, note that from (3.61 c), for instance, one readily gets the following kinds of inferences by meaning postulates (`coexten-subep-of` is an object language equivalent of '$\preceq$'):

($\exists$e2: [e2 coexten-subep-of e1] [[Mary see John] ** e2]),
($\exists$e3: [e3 coexten-subep-of e1] [[John kick Pluto] ** e3]).

[15]The adjective has to denote a temporary (rather than enduring) state; cf., **John saw Mary tall*, **John saw Mary happy*, etc. This seems related to Carlson's [1982] stage-level versus individual-level predicate distinction.

3.3.13 Intensional Verbs

I now briefly discuss how intensional verbs are treated in EL. Here are some sentences involving intensional verbs and their EL translations. (Note that (3.63 b) and (3.64 b) below are semi-scoped.)

(3.63) a. John will *design* the house
b. (pres (futr [John (design λ*x*[*x* = <The house>])]))
c. (pres (futr [John (design λ*x*(The *y*: [*y* house] [*x* (design *y*)]))]))

(3.64) a. Mary *wrote* a letter
b. (past [John (write λ*x*[*x* = <∃ letter>])])
c. (past [John (write λ*x*(∃*y*: [*y* letter] [*x* (write *y*)]))])

(3.65) a. Mary is *making* orange juice
b. (pres (prog [Mary (make λ*x*[*x* = (K ((nn orange) juice))])]))

As shown above, intensional verbs are translated as predicate modifiers. This is in analogy with the treatment of NI-perception verbs, but there is a difference in the treatment of direct objects. The objects of perception are normally treated as actual,[16] and hence should be extensionally interpreted; for objects of intensional verbs, there is generally no presupposition of actual existence — at least not in the "opaque" (*de dicto*) reading. That is, "the house" and "a letter" in the above examples do not necessarily exist, in the world wherein the sentences are evaluated. That is why they are scoped under the intensional verbs in (3.63 c) and (3.64 c). The "transparent" (*de re*) readings can be obtained by choosing wide scope for the unscoped terms <The house>, <∃ letter>, i.e., just inside the tense operators, but outside the intensional verbs. In sentence (3.65 a), borrowed from the TRAINS domain, the object is treated as predicate denoting a "kind." This can be transformed later, by meaning postulates, to a "realization" of the kind "orange juice," as below.

... [Mary (make λ*x*(∃*y*: [*y* realizes (K ((nn orange) juice))] [*x* = *y*]))] ...

which happens to be equivalent to

... [Mary (make ((nn orange) juice))] ...

[16]That is, setting aside hallucinations and the like.

3.3.14 Adverbials

Semantically, adverbials may be classified into two classes: ones that operate on sentences and ones that operate on (1-place) predicates. Typically, those that modify actions are predicate modifiers, and those that modify episodes are sentential operators. Consider the following example.

(3.66) a. John walked with Pluto in Disneyland

b. (past ((adv-e (in-loc Disneyland))
[John ((adv-a (with-accomp Pluto)) walk)]))

c. (∃e1: [e1 before Now1]
[[e1 in-loc Disneyland] ∧ [[John | e1] with-accomp Pluto] ∧
[John walk]]
** e1])

In (3.66 a), the meaning of "in Disneyland" modifies the *episode* described by "John walk," or, more specifically, its spatial location. "With Pluto," on the other hand, is about the *action* of John's walking. As mentioned earlier, episode modifiers are in the form of (*adv-e* π), where π is a predicate over episodes such as (*at-time Noon*), "at noon," and (*lasts* (*K1* (*two hour*))), "for two hours." Action modifiers take the form (*adv-a* π), where π is a 1-place predicate over actions/attributes such as (*for-benef Mary*), as in "John bought a flower *for Mary*," and (*in-manner abrupt*), as in "John *abruptly* opened the closet."

Applying appropriate meaning postulates, we obtain from (3.66 c):

(3.66) d. (∃e1: [e1 before Now1]
[[[e1 at-loc Disneyland] ∧ [[John | e1] with-accomp Pluto] ∧
(∃e2: [e2 coexten-subep-of e1] [[John walk] ** e2])]
** e1]).

Then, from (3.66 c) and (3.66 d), we get (skolemizing E1/e1 and E2/e2):

[E1 before Now1],
[E1 at-loc Disneyland],
[[John | E1] with-accomp Pluto],
[[John walk] * E1],
[E2 coexten-subep-of E1],
[[John walk] ** E2],
[[[E1 at-loc Disneyland] ∧ [[John | E1] with-accomp Pluto] ∧ [John walk]]
** E1].

The treatment of temporal and locative adverbials in EL views these as providing conjunctive information about the described episode. In this respect, it follows Davidson

[1967], Reichenbach [1947], and Dowty [1982], rather than Priorean [1967] tense logic approaches.

3.3.15 Tense and Aspect

I now show an example that involves tense, aspect and a temporal adverbial.

(3.67) a. John has been walking Pluto for ten minutes

b. (pres (perf ((adv-e (lasts (K1 (ten minute))))
(prog [John walk2 Pluto]))))

c. ($\exists$e1: [e1 at-about Now1]
[($\exists$e2: [e2 until e1]
[[[e2 lasts (K1 (ten minute))] $\wedge$ (prog [John walk2 Pluto])] ** e2])
** e1]),

with rough translation of numbers

In the indexical LF (3.67 b), perf is an indexical sentential operator indicating perfect aspect, and prog is a sentential, modal operator that yields the progressive aspect of its operand. In ELF (3.67 c), pres and perf are deindexed, introducing predications [e_1 at-about Now1] and [e_2 until e_1]. (Deindexing of indexical operators such as tense and aspect is discussed in detail in Part II of the thesis.) Note that e_2 is a ten minute-long episode over which John is walking Pluto. In other words, its characterization consists of two descriptions: that *it is ten minutes long* and that *John is walking Pluto throughout it.*

3.4 Summary and Some Comparisons

I have shown the syntax of EL and illustrated it with examples. Among others, EL is capable of representing restricted quantifiers, λ-abstracts, predicate modifiers, actions, kinds, attitudes and propositions, modal operators, and perception statements. Most importantly, however, it makes implicit time and situation dependencies explicit through the use of episodic variables, and admits unbound "anaphoric" variables and the representation of generic conditionals. The semantics of the expressions shown in this chapter is provided in Chapter 4, and their derivation from English is discussed in Chapters 6–9.

Although omitted in this chapter, speech acts are also made explicit in EL logical forms. Basically, declarative sentences are translated using speech act *tell*, questions are translated using *ask*, and imperatives are translated using *instruct* or *request*, roughly as shown below (Speaker and Hearer are the speaker and hearer arguments of the utterance context):

[Speaker tell Hearer (That Φ)]
[Speaker ask Hearer (YN-q Φ)]
[Speaker instruct Hearer (Ka π)]

This will be discussed again in Chapter 8. A few more categories are still needed in EL though, including *wh*-formulas (and corresponding nominalization operator *Answer-to*), and the operators *degree, -er, -est,* and $rank_n$ related to formation of comparison predicates. The semantics of these are still more or less open.

We believe, however, that EL is the most expressive knowledge and semantic representation yet brought to bear on the problem of narrative understanding. It goes well beyond the current state of the art as represented by such works as [Hobbs, 1985b; Hobbs *et al.*, 1986] and [Alshawi and van Eijck, 1989; Alshawi, 1990]. Both of these approaches use extensional first-order logic, with some extensions, for semantic representation. Like EL, they both admit events and states in their ontology. However, using Davidsonian event variables, they are unable to deal with non-atomic events involving quantification, logical compounds, etc. As well, neither of them distinguishes events and actions.

Let us look at some specifics by way of comparison with EL. Like EL, the *Core Language Engine* (CLE) of Alshawi and van Eijck [1989] has multiple levels of semantic representation: QLFs (quasi-logical forms which may involve unscoped expressions and unresolved referential expressions) and LFs (fully resolved logical forms). Below are some of their examples with fully resolved LFs.

(3.68) a. Every representative voted
b. quant(forall, x, Representative(x),
past(quant(exists, e, Event(e), Vote(e, x))))

(3.69) a. At least three but less than seven representatives voted
b. quant($\lambda m \lambda n.(n \geq 3 \wedge n < 7)$, x, Representative(x),
past(quant(exists, e, Event(e), Vote(e, x))))

(3.70) a. John left suddenly
b. past(quant(exists, e, Event(e), Leave(e, john) ∧ Sudden(e)))

(3.71) a. John designed a house in Cambridge
b. quant(exists, h, House(h) ∧ In_location(h, cambridge),
past(quant(exists, e, Event(e), Design(e, john, h))))
c. quant(exists, h, House(h) ∧
past(quant(exists, e, Event(e), Design(e, john, h) ∧
In_location(e, cambridge))))

(3.72) a. John invented paperclips
b. past(quant(exists, e, Event(e), Invent(e, john, kind(p, Paperclip(p)))))

As mentioned above, they use Davidsonian event variables as extra arguments of predicates. Also, note that their final LFs are still indexical, involving tense operators, and like Hobbs *et al.*, they do not specify how to deindex them. However, unlike Hobbs *et al.* who take tense operators as predicates over events, they take them as sentence operators.

Let us now consider the formulas one by one. (3.68 b) shows the restricted quantification, which is quite similar to the one in EL. More interesting is the generalized quantifier that appears in (3.69 b). They characterize a generalized quantifier with restriction set A and intersection set $A \cap B$ with a function $\lambda m \lambda n.\mathbf{Q}(m,n)$, where $m = \mid A \mid$ and $n = \mid A \cap B \mid$. In EL, numbers are treated as predicate modifiers, not quantifiers. For instance, (3.69 a) might be translated into something like the following.

$$
\begin{aligned}
&(\exists e_1 \colon [e_1 \text{ before } Now_1] \\
&\quad [(\exists x \colon (\exists n \colon [[n \text{ number}] \wedge [n \geq 3] \wedge [n < 7]] \; [x \; ((\text{num } n) \text{ representative})]) \\
&\qquad [x \text{ vote}]) ** e_1])
\end{aligned}
$$

Here, **num** is an operator that uniformly maps numeric individuals, i.e., numbers, into predicate modifiers. Next, (3.70 b) and (3.71 b) show the treatment of adverbials. They translate adverbials uniformly as predication over events. Although temporal or spatial adverbials (e.g., "in Cambridge" above) are properly translated as such, taking manner and other action modifying adverbials (e.g., "suddenly" above) as predications over events is untenable as pointed out in Chapter 2. Though "sudden events" may not sound odd, "reluctant events," "hesitant events," "intentional events," etc., do. Such adverbials need to be taken as predications over actions. Also, if we look closely at the two readings provided for (3.71 a), (3.71 b) asserts that the house designed is in Cambridge, and (3.71 c) asserts that the designing took place in Cambridge. In the latter case, the house may have never been built, but the translation insists that it exists. Thus the translation fails to capture the intensionality of "designing." In EL, such intensional verbs are treated as predicate modifiers as shown in Section 3.3.13; so, for instance, "design a house" would be translated into (design $\lambda x[x = \langle \exists \text{ house} \rangle]$), or (design $\lambda x(\exists y : [y \text{ house}]\,[x = y])$) after scoping. The point of interest in (3.72 b) is the operator `kind`, intended to form a "natural" kind from properties. `kind`(x, P(x)) is interpreted as the "typical individual satisfying P." Thus, apparently, `kind` is a variable-binding operator, since its first argument is not externally bound. However, no means are provided for representing kinds of actions or events, and it is not clear how to represent sentences like "Swimming is fun" or "For John to be late is rare." Also, the proposed language is not capable of representing donkey sentences or probabilistic conditionals (i.e., generic conditionals).

Thus, although CLE has various extensions of ordinary first-order logic, allowing events, collections, kinds, etc., to be expressed, it still omits many important English constructs, most significantly, intensional verbs. (Some modal predicates are handled, but the semantics of these remain unclear.) So CLE retains many of the weaknesses of Hobbs' neo-Davidsonian approach. Finally, where they extend FOL, they do not yet have formal semantics.

I now close the discussion by showing an example from [Hobbs, 1985b], with slight simplification, and comparing his translation with an EL representation.

> *The government has repeatedly refused to deny that Prime Minister Margaret Thatcher vetoed the Channel Tunnel at her meeting with President Mitterand on 18 May*

(3.73) Hobbs's indexical representation:

$$Perfect(E_1) \wedge repeatedly(E_1) \wedge refuse'(E_1, Govt, E_2) \wedge$$
$$deny'(E_2, Govt, E_3) \wedge veto'(E_3, MT, CT) \wedge at'(E_4, E_3, E_5) \wedge$$
$$meet'(E_5, MT, PM) \wedge on'(E_5, 18May)$$

(3.74) EL representation

a. Indexical representation:

(The x: [[x meeting] $\wedge$ [x between MT PM] $\wedge$ [x on-time (date 5 18)]]
    (pres (perf ((adv-f repetitive)
        [Govt refuse (Ka (deny (That
            (past ((adv-e (at x)) [MT veto CT]))))))]))))

b. Deindexed representation (speech act neglected):

(The x: [[x meeting] $\wedge$ [x between MT PM] $\wedge$ [x on-time (date 5 18)]]
    ($\exists e_1$: [e_1 at-about u_1]
        [($\exists e_2$: [e_2 impinges-on e_1]
            [[[e_2 ((attr repetitive) multi-component-ep)] $\wedge$
             (mult [Govt refuse
                (K λa[[(fst a) deny (That
                    ($\exists e_3$: [e_3 before (rst a)]
                        [[[e_3 at x] $\wedge$ [MT veto CT]] ** e_3]))]
                  ** (rst a)])])]
            ** e_2])
        ** e_1]))

In (3.74b), u_1 is the utterance episode. Various predicates used in the above formula are explained in Chapters 8–9, but it should suffice here to say this: `impinges-on` is a predicate obtained by deindexing the `perf` operator and is to be particularized to `until` — thus, the above formula says that some repetitive event extends till e_1, i.e., till the utterance time. A multi-component episode (`multi-component-ep`) is an episode whose temporal projection is a multi-interval, thus, having two or more component episodes. The sentential operator mult indicates that every component of a "composite" episode characterized by (mult Φ) is uniformly of type Φ. In this case, each component of the multi-component episode e_2 is of type "Govt refusing some kind of actions," namely, refusing "denying the possible fact that MT vetoed CT." The λ-variable a denotes an action, an agent-episode pair; so, (fst a) denotes the (undefined) agent of the "denying" episode, and (rst a), the "denying" episode itself.

Note that Hobbs's logical form (3.73) is still indexical because of its use of *Perfect* as a predicate (presumably making implicit reference to a speech time and reference time). It would be hard to make inferences with such logical forms unless a fixed speech time and reference time is assumed for *all* the input sentences as well as for all the formulas in the knowledge base. Aside from this, it is not at all clear how to interpret a conjunction like $repeatedly(E_1) \wedge refuse'(E_1, Govt, E_2) \wedge deny'(E_2, Govt, E_3)$. The second and third conjuncts seem to say that E_1 is of type "Govt refuses E_2," where E_2 is a "denial of E_3 by Govt." *In the absence of the first conjunct*, we would presumably interpret E_1 as a *single* "refusal event," involving a single "denial event" E_2. But if that's what the second and third conjuncts say about E_1 and E_2 in isolation, they still say this when the first conjunct, $repeatedly(E_1)$ is added. At least, they still say this if the semantics for this logic is compositional, as one would suppose if this is ordinary FOL. Even if we interpret clauses like the second and third as potentially denoting multiple episodes of the specified type from the outset, it is not at all clear how a denial event of the sort specified in the third clause can get associated with *each* refusal event specified in the second clause, simply through assertion of the conjunction. In contrast, EL formulas are completely deindexed (via the deindexing mechanism developed in Part II of this thesis), ready for use in inference. More importantly, EL representations have a formal interpretation as will be seen in the next chapter.

In concluding this chapter, I should reiterate the rationale behind the rather rich "natural language-like" syntax of EL. The claim is certainly not that all of the syntactic features are *necessary* in either a general semantic representation for NLP or a general knowledge representation for commonsense reasoning. For instance, one can accommodate propositions in standard FOL, and similarly nonstandard quantifiers can be eliminated with the aid of some set-theoretic predicates and functions (see, e.g., [McCarthy, 1979]). Rather, the claim is that all features of EL are strongly motivated by corresponding expressive devices found in natural languages—i.e., generalized quantifiers, modifiers, nominalization, etc. By "mimicking" these devices in EL, we make possible a maximally simple transduction from surface form to meaning representation, and at the same time provide a high-level, intuitively comprehensible KR, which also permits intuitively obvious inferences to be modelled in a direct, straightforward way (as later chapters will indicate).

Chapter 4

The Logical Semantics of Episodic Logic

This chapter presents a partial formal semantics of EL. I start with the model structure of EL, reviewing the ontology discussed in Chapter 2. The focus of this chapter is on situations, however, especially their relationships to each other and to space and time. After that, I provide some semantic preliminaries, namely, the parameter mechanism, semantic types of atomic expressions, and definitions of classes of persistent functions. Then, I move on to actual semantic clauses for various EL constructs. This should be sufficient to interpret ELF expressions discussed in Chapter 3. (The discussion here refers to non-indexical ELFs.) I close the chapter with a discussion of entailment and anti-entailment in EL, which form the basis of EL inference rules discussed in Chapter 5, and a list of some valid EL schemas.

Given the commitment in this thesis, to comprehensive coverage from the outset with gradual refinement and deepening of the foundations, what is offered in this chapter is both incomplete and tentative. There have been some rather significant revisions within the last year or so. In particular, the semantics of connectives and quantification has been simplified and made persistent (so that the truth value of a quantified sentence persists under informational "enlargement" of a situation). The simplification in the semantics is the result of viewing bounded sentences as having outward persistent truth conditions. For instance, if "John mowed the lawn" is true in a certain one-hour situation, it is also true in longer situations (e.g., the day of the mowing). This view can be reconciled with the intuition that there *is* a certain fixed time that an event "occupies" (e.g., a one-hour period, rather than a day, in the case of the mowing) by making use of the distinction between '**' (characterization) and '*' (description) in EL. Characterizations pick out *minimal* situations (in both spatiotemporal extent and content), and are not outward persistent; but (bounded) descriptions *are* outward persistent.

Though some of the decisions presented in this chapter may yet need be revised, e.g., representation of space-time locations of situations, and there are still some remaining lacunas, e.g., semantics of questions and wh-formulas, future revisions to the semantics are expected to keep the basic framework more or less intact.

4.1 The Model Structure

The entities which make up a model structure for EL are rather numerous, as might be expected from the richness of the syntax and ontology. Moreover, this set of entities is somewhat open-ended, since it is not at all clear in EL where to draw the line between *logical* symbols—those whose meanings are not subject to interpretation and must be specified via separate elements of the model structure—and *nonlogical* ones subject to interpretation (and further specification through meaning postulates). Examples of EL symbols on the logical/nonlogical border are `subep-of`, `cause-of` (denoting binary relations over episodes), `clock-time-of` (denoting a function from episodes to real multi-intervals), `prog` (a sentence modifier corresponding to the English progressive), and `attr` (an operator which combines with a predicate to produce a predicate modifier, used in interpreting prenominal adjectives). In fact, even the basic episodic modal operators '**' and '*' are to some extent open to interpretation, though their interpretations are strongly constrained by the interpretations of other symbols, as well as by certain explicitly enumerated elements of a model structure (e.g., $\sqsubseteq$, $\preceq$, *Actual*, and *Nonactual*).

In view of this complexity and open-endedness, I forego the traditional "tuple" form of model structure, instead using a more informal format. I first briefly enumerate the most basic ingredients needed to build models, with their mathematical types (and sometimes object-language "counterparts"). Then I delve more deeply into the intuitive motivation for, and properties of, these ingredients. Here, as elsewhere in EL, the end result is not yet a finished project, but (it is hoped) a plausible initial draft, a concrete basis for some formal theorizing—e.g., about persistence, valid schemas, and inference—and point of departure for further refinement.

4.1.1 Some Essential Ingredients

Only the items marked with '%' are independently stipulated; the rest are defined in terms of these. Occasionally, an object-language counterpart of the item is given in `[]`.

% **Ontology** $\mathcal{O}$**:** a set of (possible) individuals $\mathcal{D}$ and subsets $\mathcal{S}$ (situations), $\mathcal{H}$ (exhaustive situations), $\mathcal{I}$ (times), $\mathcal{W}$ (worlds), $\mathcal{M}$ (moments), $\mathcal{P}$ (propositions), $\mathcal{F}$ (facts), $\mathcal{K}$ (kinds), $\mathcal{K}_{\mathrm{A}}$ (kinds of actions/attributes), $\mathcal{K}_{\mathrm{E}}$ (kinds of episodes/situations), $\mathbb{R}$ (the real numbers, along with $-\infty$ and $+\infty$), $\mathcal{R}$ (the real multi-intervals), $\mathcal{R}_2$, $\mathcal{R}_3$,

$\mathcal{R}_4$ (2, 3, and 4-D regions), $\mathcal{C}$ (collections), $\mathcal{V}$ (vectors/tuples), and various refinements and combinations of these. As shown in Figure 2.1, $\mathcal{S}$, $\mathcal{P}$, and $\mathcal{K}$ are pairwise disjoint, as are $\mathcal{K}_A$ and $\mathcal{K}_E$; and $\mathcal{W}$ and $\mathcal{M}$. Also, $(\mathcal{W} \cup \mathcal{M}) \subset \mathcal{I} \subset \mathcal{H} \subset \mathcal{S}$, $\mathcal{F} \subset \mathcal{P}$, and $(\mathcal{K}_A \cup \mathcal{K}_E) \subset \mathcal{K}$.

% **Coextensive part-of** $\preceq$ [coexten-subep-of] : a partial ordering on situations such that for any $s \in \mathcal{S}$, there is at least one maximal element $h \in \mathcal{H}$ such that $s \preceq h$. For $h, h' \in \mathcal{H}$, $h \preceq h'$ iff $h = h'$.

% **Part-of** $\sqsubseteq$ [subep-of] : a partial ordering on situations which extends the $\preceq$-ordering, such that for any $s \in \mathcal{S}$ there is at least one maximal element $w \in \mathcal{W}$ such that $s \sqsubseteq w$. If $s \in \mathcal{H}$, this maximum is unique. For $w, w' \in \mathcal{W}$, $w \sqsubseteq w'$ iff $w = w'$. Moreover, $\sqsubseteq$ forms a join semilattice[1] for each set $\{s \mid s \preceq w\}$, $w \in \mathcal{W}$, where any pair of situations occurring in more than one such set has the same *l.u.b.* in each.

Join $\sqcup$: the join operator for the $\sqsubseteq$-semilattices, i.e., for $s, s' \in \mathcal{S}$, where $s, s' \preceq w$ for some $w \in \mathcal{W}$, $s \sqcup s' = lub_{\sqsubseteq}(s, s')$.

% ***Actual*** [actual] : a relation $\subset \mathcal{D} \times \mathcal{S}$ extending $\sqsubseteq$, such that for $d \in \mathcal{D}$, $s, s' \in \mathcal{S}$,

$$Actual(d,s) \wedge s \sqsubseteq s' \supset Actual(d,s').$$

% ***Nonactual*** [nonactual] : a relation $\subset \mathcal{D} \times \mathcal{S}$ disjoint from *Actual*, such that for $d \in \mathcal{D}$, $s, s' \in \mathcal{S}$,

$$Nonactual(d,s) \wedge s \sqsubseteq s' \supset Nonactual(d,s').$$

World: the total function $\in \mathcal{H} \rightarrow \mathcal{W}$ such that for all $h \in \mathcal{H}$, $h \sqsubseteq World(h)$. I.e., $World(h)$ is the $\sqsubseteq$-maximum of h.

% ***Region***: a total function $\in \mathcal{S} \rightarrow \mathcal{R}_4$ such that for $s, s' \in \mathcal{S}$,

$$s \preceq s' \supset Region(s) = Region(s'), \text{ and}$$

$$Region(s \sqcup s') = Region(s) \cup Region(s').$$

For $h, h' \in \mathcal{H}$,

$$[Region(h) = Region(h') \wedge world(h) = world(h')] \supset h = h'.$$

For $w \in \mathcal{W}$,

$$Region(w) = \mathbb{R}^4.$$

For $i \in \mathcal{I}$,

$$Region(i) = r \times \mathbb{R}^3 \text{ for some } r \in \mathcal{R}.$$

By the discussion in Chapter 2, $Region(s)$ for any s is a closed subset of $\mathbb{R}^4$, and no finite straight line crosses the boundary of $Region(s)$ infinitely often.

[1]A poset is a join-semilattice if $lub(a, b)$ exists for any two elements. See [Grätzer, 1971] for lattice theory.

Clocktime [clock-time-of] : a total function $\in \mathcal{S} \rightarrow \mathcal{R}$ which is the projection of *Region* onto the first coordinate; i.e., for $s \in \mathcal{S}$,

$$Clocktime(s) = \{r_1 \mid \langle r_1, r_2, r_3, r_4 \rangle \in Region(s) \text{ for some } r_2, r_3, r_4\}.$$

By the discussion in Chapter 2, *Clocktime*(s) is a multi-interval.

Begin [begin-of] : the total function $\in \mathcal{S} \rightarrow \mathbb{R}$ giving the *g.l.b.* of *Clocktime*, i.e., the beginning of the clocktime multi-interval.

End [end-of] : the total function $\in \mathcal{S} \rightarrow \mathbb{R}$ giving the *l.u.b.* of *Clocktime*, i.e., the end of the clocktime multi-interval.

Inside : the partial ordering $\{\langle s, s' \rangle \mid Region(s) \subseteq Region(s');\ s, s' \in \mathcal{S}\}$.

During : the partial ordering $\{\langle s, s' \rangle \mid Clocktime(s) \subseteq Clocktime(s');\ s, s' \in \mathcal{S}\}$.

% ***Init*** : a function $\in \mathcal{S} \times \mathcal{W} \rightarrow \mathcal{M}$ defined whenever $Begin(s) \neq -\infty$, such that for all $\langle s, w \rangle \in domain(Init)$,

$$Clocktime(Init(s, w)) = \{Begin(Clocktime(s))\}, \text{ and } Init(s, w) \sqsubset w.$$

% ***Fin*** : a function $\in \mathcal{S} \times \mathcal{W} \rightarrow \mathcal{M}$ defined whenever $End(s) \neq +\infty$, such that for all $\langle s, w \rangle \in domain(Fin)$,

$$Clocktime(Fin(s, w)) = \{End(Clocktime(s, w))\}, \text{ and } Fin(s, w) \sqsubset w.$$

Also,

$$Init(s, w) = Fin(s, w) \text{ iff } Clocktime(Init(s, w)) = Clocktime(s) \text{ iff } Clocktime(s) = Clocktime(Fin(s, w)).$$

Time [time-of] : the total function $\in \mathcal{S} \times \mathcal{W} \rightarrow \mathcal{I}$ such that for all $s \in \mathcal{S}$, $w \in \mathcal{W}$,

$$Time(s, w) \sqsubseteq w, \text{ and}$$
$$Clocktime(Time(s, w)) = Clocktime(s).$$

(The remaining issues need further specification.)

% ***Prop*** : a boolean algebra $[\mathcal{P}, \dot{\vee}, \dot{\wedge}, \sim]$.

Fact : the subalgebra of *Prop* exclusive of the minima of *Prop*.

% ***Subkind*** : a partial ordering on $\mathcal{K}$ forming a semilattice.

% ***Part-of-coll*** : a partial ordering on collections in $\mathcal{C}$ forming a semilattice.

4.1.2 Spatio-Temporal Extents of Situations

In contrast with situation semantics, situations in EL are primitives rather than structured. They comprise a major category in the ontology of EL. As discussed in Chapter

2, the (possible) situations $\mathcal{S}$ include exhaustive situations $\mathcal{H}$ (situations with maximal propositional content), and these in turn include times $\mathcal{I}$ (viewed as spatially maximal situations with maximal propositional content, i.e., time slices or (multi-)intervals of the universe, supporting "everything that is true" in that multi-time slice), which in turn include moments of times $\mathcal{M}$ and possible worlds $\mathcal{W}$ (viewed as unbounded times, i.e., spatiotemporally maximal as well as maximal in propositional content). Moments of time $\mathcal{M}$ and possible worlds $\mathcal{W}$, usually taken as independent indices of possibility in possible-worlds semantics, are here temporally minimal and maximal time intervals respectively. Note that an exhaustive situation $h \in \mathcal{H}$ can be thought as *part(s) of time intervals*, i.e., spatially bounded chunks of times. Certain classes of sentences (atemporal/unlocated ones) have the same truth values in both, while other types of sentences may have different truth values in exhaustive situations than in the times of which those situations are a part. Here I should also comment on the term 'episode'. Strictly, this is synonymous with 'situation', but we prefer the term when referring to situations evoked by narratives, which are typically "episodic."

One of the fundamental properties of situations is their spatiotemporal location. These are the "when and where" of the events and circumstances described by EL sentences. It is important not to confuse the "when and where" of events and circumstances with that of their *participants*. Inferences about locations of participants are dependent on the particular predicates, terms and operators involved in a formula. For instance, if the event "John kissed Mary" occurred in the park, it is probably legitimate to conclude that both were in the park at the time. However, a similar conclusion does not follow for "As a result of the movie 'Amadeus', Mozart gained new admirers in the States."

At the beginning of this research, *time* was given a privileged role in relation to *space*, and only time was considered as an intrinsic property of situations. It appeared that if certain individuals participate in certain events or relations *at a certain time*, that is sufficient to "locate" the corresponding situation. The reason was that to the extent that "event locations" made sense to us, they seemed parasitic upon the locations of the participants. Also, the existence of tense inflections, but not locative inflections, in many languages seemed to indicate that the temporal dimension is privileged. Another concern was that spatial locations of events are often ill-defined, e.g., for sentences like the following.

(4.1) The star I am looking at in the telescope is slowly moving

(4.2) John saw Mary in the mirror

(4.3) The sun is shining through my window[2]

(4.4) The rumor spread as far as where John's grandparents live

(4.5) Bill has a friend in Montana

(4.6) Dodos became extinct

In the case of (4.6), one might argue that this event occurred wherever the last dodo died (e.g., in Mauritius), but this is incorrect since the extinction of a species, though *correlated* with the death of its last representative, *isn't* that event. A becoming-extinct is a becoming-extinct *everywhere*. In fact, the occasionally encountered difficulty in spatially locating events have been observed by many (see, e.g., [ter Meulen, 1986]).

However, sentences like

(4.7) Water is scarce in California, but abundant in Michigan

(4.8) George is popular in Montana but not in Ohio

(4.9) In many developed countries, more and more people quit smoking

(4.10) A dozen eggs cost fifty cents in Ohio, while they cost a dollar in New York

(4.11) The grass is greener near the bank

(4.12) The dollar rose slightly in Hong Kong yesterday, but dropped in Tokyo

clearly indicate that temporal location by itself is not always enough to locate situations. There are properties, e.g., *scarce* or *popular* above, which the same individual may possess in one place and lack in another simultaneously. These individuals may be ordinary ones as in (4.8) and (4.11), or they may be kinds as in (4.7) and (4.12), and the properties may be generic/habitual or otherwise. Thus, locating situations requires space be taken seriously.

Given that, the question is how to specify temporal and spatial locations of situations. For this, we first need to determine whether to separate spatial and temporal locations of situations or to amalgamate them somehow. Conceptually, it seems simplest to separate them, and for most sentences this might be unproblematic. For instance, "Mary bought a book at the Village Green" describes a situation whose temporal location is some time interval in the past, and the spatial one is some chunk of space inside the Village Green. However, as Cooper's [1985] well-known sentence,

(4.13) It didn't not snow on the trip from Madison to Chicago,

illustrates, separating spatial and temporal locations makes it hard to get an intuitively satisfactory interpretation. As Cooper [1985, p. 5] points out, (4.13) "could be true even if it had snowed during the trip on the road between Madison and Chicago and yet had not been snowing at any time at the place where the car was at the time." Thus, it seems

[2] This sentence is due to Alexander Nakhimovsky (personal communication).

unavoidable to "amalgamate" spatiotemporal locations. This is also a prevalent view of those researchers working within situation semantics frameworks; e.g., Barwise and Perry's [1983] "cylinders," Cooper's [1985] "sausages," and Hinrichs' [1985] "worms" —all these refer to such amalgamated spatiotemporal locations. However, they do not provide much guidance on how to specify them mathematically.

An intuitive requirement is that the temporal and spatial projections of event trajectories should not be too disjointed. The temporal projection (i.e., the times at which the event was in progress) may well consist of disconnected chunks (e.g., for repetitive or sporadically interrupted actions), but surely not of arbitrarily densely packed chunks. In other words, the temporal projection should consist of only finitely many subintervals over any finite stretch of time. This lies behind the use of multi-intervals in the temporal dimension. Analogously, the spatial projection of trajectories ought not to be arbitrarily convoluted. Requiring any finite straight line (in *4-D* trajectories) to have only finitely many boundary crossings on it appears to yield the desired properties. However, nothing in this thesis hinges on this assumption. It is made for specificity, but may be ultimately unsatisfactory.

As emphasized in Chapter 2, temporal projections of the *4-D* regions occupied by events, or "clock times," must be distinguished from times as certain kinds of exhaustive situations, i.e., situations supporting "everything that happened" in a particular world within a particular clock time.

4.1.3 Motivating the Algebraic Structures on Situations

Various algebraic structures have been assumed above, most importantly the orderings $\preceq$ and $\sqsubseteq$, the relation *Actual*, *Nonactual* and others. I now discuss these more fully.

The ordering $\preceq$ is intended to reflect basic intuitions about "cumulative information content" at some space-time location, constituting part of the state of affairs at that space-time location. For instance, one such piece of information may be that John hugged Mary at that place and time; this characterizes some situational part of the state of affairs there. He may also have concurrently kissed Mary, and that characterizes another situational part of that state of affairs. But this hugging and kissing intuitively join into a more comprehensive situation which is again part of that state of affairs, but is "informationally larger" than either, though spatiotemporally coextensive with them. It is this "information accumulation" which $\preceq$ is intended to capture. But one imagines that the information that can be added is ultimately exhausted (though of course not necessarily in any finite limiting process); we then have a locally complete state of affairs, an *exhaustive* situation $h \in \mathcal{H}$ that cannot be further enlarged.

Thus, *exhaustive situations* $\mathcal{H}$ can be identified as those situations (n.b., *situations*, not regions) maximal in the $\preceq$-ordering. In other words, exhaustive situations have "maximal content" relative to their *coextensive* parts. Exhaustive situations, in this sense, may be arbitrarily small, e.g., "here and now," or large, e.g., the history of universe or a complete *possible world*. As mentioned in Chapter 2, their factual content is "everything that happened, or was the case, during that time in that space in that world." Keep in mind exhaustive situations (times and worlds, as well) are episodically uncharacterizable because of the amount of information they have. For instance, an exhaustive situation whose location is "this room at this moment" would have not only information about "events of interest" at that time and place but *everything* that is the case there, such as what color the wallpaper is, what is written in the books on the shelf, and how air molecules are moving; in addition, it would support all the facts of the world it belongs to — who was talking with me five minutes ago, the chemical composition of the grains of sand at the seashore, who will win the next presidential election, what creatures populated the oceans two billion years ago, etc.

The $\sqsubseteq$ relation takes this accumulation of information a step further, by allowing for an expansion of the space-time region whose associated state of affairs is under consideration. With this new way for situations to "increase" (spatiotemporally as well as informationally), we should "eventually" get to the all-encompassing situation, i.e., a possible world $w \in \mathcal{W}$. Before getting there, though, if we maximize information content without maximizing space-time, we should get to some spatiotemporally bounded exhaustive situation, $h \in \mathcal{H}$. But once there, the world that this h is part of should be determinate (unique), since a complete state of affairs locally will also settle such issues as "Did John kiss Mary at such-and-such (other) places and times?" Thus, as suggested in Chapter 2, it carries in it the information of the entire world — past, present and future.

These remarks motivate the assumption that $\sqsubseteq$ forms a semilattice of situations with respect to each world, as well as the assumption that exhaustive situations uniquely determine the world to which they belong. The only additional assumption that was made about $\sqsubseteq$ in the previous subsection was that two situations do not have different *l.u.b.*'s (i.e., joins) with respect to different worlds. This assumption is made for simplicity, allowing a functional view of the *l.u.b.*, where it exists.

Then $\sqcup$, *situation join*, is the *l.u.b.* (or join) determined by $\sqsubseteq$. Note that this join is defined for any two situations belonging to at least one common world, but is not defined everywhere. In particular, the join of exhaustive situations (including times and worlds) is undefined if those situations belong to different worlds. This follows from the uniqueness of worlds determined by exhaustive situations and the maximality of worlds in the $\sqsubseteq$-ordering.

Two disjoint relations over $\mathcal{D}\times\mathcal{S}$, namely, *Actual* and *Nonactual*, determine what entities are actual and nonactual relative to a situation. Individuals can be *Actual* relative to any number of worlds, and must be *Actual* or *Nonactual* relative to any *given* world. Together, these are the *participants* in the situation. Note that since *Actual* extends $\sqsubseteq$,

$$s \sqsubseteq s' \text{ only if } Actual(s, s').$$

The "transitivity axiom" given earlier for *Actual* expresses the intuition that an actual entity participating in one situation is still actual in any larger situation, and also participates in that larger situation. Analogously, nonactuality persists through the $\sqsubseteq$-ordering.

In general, situations, like other individuals, can be *Actual* relative to any number of worlds; however, exhaustive situations are *Actual* relative to exactly one world, namely, the one of which they are a part, and *Nonactual* relative to all others. Similarly, times are *Actual* relative to exactly one world, of which they are a part. Intuitively, the reason for this assumption is that exhaustive situations, being factually maximal, i.e., *all* the "unlocated" facts of a world being actual to *all* exhaustive situations in that world, already encapsulate the history of their universe. That is, for any given exhaustive situation h, the fact that some earlier event occurred *is a fact of* h, and similarly for future events. Thus, the total function $World(h)$ supplies the unique world of a given time h. A situation-world pair then uniquely determines the exhaustive situation $\in \mathcal{H}$ as well as the time $\in \mathcal{I}$. Note the *times* $\mathcal{I}$ are those exhaustive situations which encompass all space, and among these *moments of time* $\mathcal{M}$ are the temporally minimal ones, i.e., momentary states of the universe. So, on our conception of exhaustive situations, the world component of a world-time index is redundant, and indeed in our semantics there is just one index of possibility, a situational one. If this index is a time, the world is implicitly determined as well.

The *Region* function expresses the fundamental intuition that events take place (and states of affairs) exist *some* place, *some* time, i.e., they have a space-time trajectory. Note the correspondence we have assumed between these space-time regions and exhaustive situations: a space-time region uniquely determines an exhaustive situation in any given world. Thus, as far as the semantic values of EL expressions at *exhaustive* situations are concerned, we could just as well have used regions plus possible worlds, rather than situations, as indices of evaluation. (In fact, we could have used clock-times and worlds, since a clock-time r determines a region $r \times \mathbb{R}^3$, and this is the region of an exhaustive situation, namely a time.)

The functions *Clocktime*, *Begin*, *End*, *Inside* and *During* are at this point self-explanatory. The functions *Init* and *Fin* yield moments from situation-world pairs such that $Init(s, w)$ and $Fin(s, w)$ are the initial moment and the final moment of the time i determined by $\langle s, w \rangle$ respectively. Moments are temporally minimal times with no duration, i.e., momentary states of the world as mentioned above.

4.2 Semantic Preliminaries

We are now ready to begin the discussion of interpretations and valuations. The first steps are the definition of parameters, the semantic types of atoms, and the classes of persistent functions used in these types. After this, I will outline the conditions satisfied by an interpretation and by its extension to a valuation function.

4.2.1 Parameters

First, we need to understand the quite unconventional notion of "preemptable" quantification and the notion of the "parameters" of a formula that are used in EL. These notions allow us to keep the phrase-by-phrase mapping from syntactic form to logical form uniform and simple, even when interpreting anaphora, including "donkey" anaphora. In effect, they make possible a DRT-like treatment of indefinites, allowing existential variables to be used out of their scope, in generic conditionals and elsewhere.

In EL, quantifiers are classed into preemptable ones (e.g., $\exists$ and The) and non-preemptable ones (e.g., $\forall$ and Most). Preemptable quantifiers, often called determiners, are "weak" or "preemptable" in the following sense. If an interpretation I already assigns a value d to a variable α in a quantificational formula ($\mathtt{Q}\,\alpha{:}\,\Phi\Psi$), $\mathtt{Q} \in \{\exists,$ The$\}$, then the quantifier is preempted; that is, instead of iterating over the entire domain $\mathcal{D}$, it "iterates" only over the singleton $\{d\}$. It has its usual quantificational force only if the variable it quantifies is *a priori* valueless; otherwise, it is "ignored" (though it still serves to mark its variable as preemptable). For instance, if the interpretation of x is well-defined under I, i.e., $I(x) = d$ for some $d \in \mathcal{D}$, then

$$(\exists x{:}\,[x \text{ man}]\,[x \text{ rich}])$$

means that a *certain* man, namely d, is rich. In other words, the formula is interpreted referentially as if it said simply $[[x \text{ man}] \wedge [x \text{ rich}]]$. $\exists$ iterates over the full domain $\mathcal{D}$ only if the quantified variable is undefined under the given interpretation. Thus, if $I(x)$ is undefined, $(\exists x{:}\,[x \text{ man}]\,[x \text{ rich}])$ has its ordinary meaning, "There exists a man who is rich." The definite quantifier The is treated similarly as preemptable, but with an added uniqueness condition (which is automatically satisfied in the referential case). An important consequence of this is that variable names matter in EL. For instance, given ($\exists$x: [x ball] ($\exists$y: [y window] [x hit y])), the continuations (The x: [x thing] [x break]) and (The y: [y thing] [y break]) have different meanings. (However, it is possible to convert to a normal form in which quantifiers play their customary role.[3])

---

[3]In particular, consider the following equivalence transformation:

(1) *Scope widening*: $\text{wffs}_1 \wedge (\exists\alpha : \Phi\Psi) \wedge \text{wffs}_2 \Longleftrightarrow (\exists\alpha : \Phi\ [\text{wffs}_1 \wedge \Psi \wedge \text{wffs}_2])$,
where the wffs are joined by conjunctions, and no free variables of Φ occur in wffs_1.

(2) *Copying domain constraints into consequent*: $(\exists\alpha : \Phi\Psi) \rightarrow \Upsilon \Longleftrightarrow (\exists\alpha : \Phi\Psi) \rightarrow (\exists\beta : \Phi_{\beta/\alpha}\Upsilon_{\beta/\alpha})$,
where Υ is a wff with free occurrences of α, and β is a new variable. Thus, the $\exists\alpha$ and $\exists\beta$ quantifiers will both have their conventional force. This rule also applies if instead of connective '$\rightarrow$', we have '$\rightarrow_{p,\overline{v}}$', as long as the controlled variables $\overline{v}$ do not include α.
Finally,

(3) *Dropping preempted quantifiers*: $(\exists\alpha_i : \Phi\Psi) \rightarrow_{p,\overline{\alpha}} \Upsilon \Longleftrightarrow [\Phi \wedge \Psi] \rightarrow_{p,\overline{\alpha}} \Upsilon$,
where α_i is one of the controlled variables $\overline{\alpha}$.

(1, 2) allows us to transform, e.g.,

Quantifiers like ∀, Most and Few are non-preemptable. These quantifiers iterate over the entire set of individuals that satisfy the restriction clause in the situation the quantified formula is describing. For instance, for the formula

$$[(\forall x\colon [x \text{ man}]\,[x \text{ rich}]) ** E],$$

$(\forall x\colon [x \text{ man}]\,[x \text{ rich}])$ iterates over the set of individuals that are men in situation E. That is, for every individual d, if he is a man in situation E, then he is rich. Note that the mere fact that a predication $[\tau\ \pi]$ or $[\tau\ \pi\ \eta]$ has value 1 in situation E does not necessarily entail that the things denoted by τ or η are physically at the location of E. Rather, such locative inferences depend very much on the predicates involved. For instance, if we say that in a certain situation "Mary wrote down the names of *all her friends*," those friends over which we are quantifying do not need to be where Mary is in that writing situation. They *do*, on the most natural reading, have to physically exist (and be her friends) at the *time* of writing, but not at the *place* of writing. So, the temporal and spatial implications of $[\tau\ \pi]$ or $[\tau\ \pi\ \eta]$ about *when* and *where* τ or η are located can be quite distinct.

Intuitively, the *parameters* of a formula are just the top-level existentially quantified variables of the formula, except that they may be embedded by "non-negative" operators such as ∃, The, ∧, ∨, *, and **.[4]

Formally, the *parameters* of formula Φ, written with an underscore function as $\underline{\Phi}$, the set of variables exported, are defined as in Table 4.1. Note that rule I covers, *inter alia*, expressions of form ((adv-a π) Π), ((adv-e π) Φ), ((adv-f π) Φ), (□ Φ), ((adv-p π) Φ), (That Φ), etc. Expressions which do not export parameters may nevertheless involve them internally. For instance, $\underline{[\Phi \rightarrow \Psi]} = \emptyset$, i.e., material conditional does not export parameters. But note that the parameters in the antecedent, i.e., $\underline{\Phi}$, are accessible to the consequent Ψ. It is just that they are not accessible outside the conditional. In other words, parameters have scopes (as in DRT).

A couple of further comments are in order. Since the general rule for $\underline{(\pi\ \alpha)}$ applies to modal sentences such as $\underline{\Box\Phi}$, $\underline{((\text{adv-p certain})\ \Phi)}$, $\underline{((\text{adv-p probable})\ \Phi)}$, etc., modally

$(\exists x\colon [x \text{ horse}]\,[\text{John own } x]) \rightarrow [\text{John riding } x]$

into the more conventional

$(\exists x\colon [x \text{ horse}]\,[\text{John own } x]) \rightarrow (\exists y\colon [y \text{ horse}]\,[\text{John riding } y])$.

(1, 2, 3) allow us to "conventionalize" probabilistic conditionals, eliminating all instances of an existential variable (either controlled or otherwise) occurring "anaphorically." (Note: controlled variables are considered *bound* by the conditional that controls them.)

[4]The motivation for allowing embedding by episodic operators lies in sentences like "If Jack sees a pretty girl, he follows her," or, in its rough episodic logical form,

(∃e [(∃x: [x ((attr pretty) girl)] [Jack see x]) ** e]) → (∃e1 [[Jack follow x] ** e1]).

Here, x needs to be exported out of the scope of its quantifier. As will be seen in Chapter 6, in the initial translation we would get the above formula instead of the following (in which "a pretty girl" has wider scope than '**'):

(∃e (∃x: [x ((attr pretty) girl)] [[Jack see x] ** e])) → (∃e1 [[Jack follow x] ** e1]).

A. $\underline{\alpha} = \emptyset$, for atomic expressions α

B. $\underline{(\exists\alpha\ \Phi)} = \underline{(\text{The}\,\alpha\ \Phi)} = \{\alpha\} \cup \underline{\Phi}$

C. $\underline{(\exists\alpha{:}\,\Phi\Psi)} = \underline{(\text{The}\,\alpha{:}\,\Phi\Psi)} = \{\alpha\} \cup \underline{\Phi} \cup \underline{\Psi}$

D. $\underline{(\lambda\alpha\ \Phi)} = \underline{\Phi}$

E. $\underline{[\Phi \wedge \Psi]} = \underline{[\Phi \vee \Psi]} = \underline{[\Phi \text{ because } \Psi]} = \underline{\Phi} \cup \underline{\Psi}$

F. $\underline{(\neg\Phi)} = \underline{[\Phi \rightarrow \Psi]} = \underline{[\Phi \rightarrow_{p,\alpha_1,\alpha_2,\ldots,\alpha_n} \Psi]} = \underline{(\mathrm{Q}\,\alpha\,\Phi)} = \underline{(\mathrm{Q}\,\alpha{:}\,\Phi\Psi)} = \emptyset$,
where Q is a non-preemptable quantifier such as $\forall$, `Most`, etc.

G. $\underline{[\Phi * \eta]} = \underline{[\Phi ** \eta]} = \underline{\Phi}$

H. $\underline{(\pi\ \Pi)} = \emptyset$, for $\pi \in \{$`K`, `Ka`, `Ke`, $\ldots\}$

I. $\underline{(\pi\ \alpha)} = \underline{\pi} \cup \underline{\alpha}$,
for *all other nonatomic* expressions of form $(\pi\ \alpha)$

Table 4.1: $\underline{\Phi}$, The Parameters of Formula Φ

embedded indefinites are accessible to pronominal reference. This is quite plausible for *factive* modals like *necessarily*, as illustrated by "There is necessarily *a smallest natural number*, and *it* is 1." It is somewhat less plausible for nonfactive modals, e.g., "There may be an *even number greater than 2 which is not the sum of two primes*. If anyone identifies *it*, he will become famous." Also, negative contexts are problematic, and the above decision to block variable exporting from negated contexts, while allowing export from *That*-contexts, is in need of refinement. For example, consider "I am not reading a novel. *It* is a biography," or "John has hardly/never dated a girl to the end of the evening. He always manages to offend *her* and makes *her* leave in the middle." Concerning *That*-contexts, contrast "John knows that Bill has a Ferrari. He wants to borrow *it*" with "John doubts that Bill has a Ferrari. But if he has **it*, he would like to borrow it."

4.2.2 Semantic Types of Atomic Expressions

Recall that the semantics of EL is based on function application as in Montague semantics. Thus, atoms in EL are typed. Table 4.2 shows the semantic types of EL atomic expressions. In the table, we use $2 = \{0, 1\}$ as truth values and write $A \rightarrow B$ for the set of partial functions from A to B. $A^n \rightarrow B$ will abbreviate $(A \rightarrow (A \rightarrow \ldots (A \rightarrow B) \ldots))$, with

Type of atom α	$I(\alpha)$ is an element of :
Individual constant	$\mathcal{D}$
Individual variable	$\mathcal{D}$, or is undefined
Function constant	$\mathcal{D}^n \rightarrow \mathcal{D}$
Sentence constant	$[\mathcal{S} \rightarrow 2]$
n-place predicate constant	$\lceil \mathcal{D}^n \rightarrow (\mathcal{S} \rightarrow 2) \rceil$
n-fold predicate modifier	$\lfloor \lceil \mathcal{D} \rightarrow (\mathcal{S} \rightarrow 2) \rceil \rightarrow \lceil \mathcal{D}^n \rightarrow (\mathcal{S} \rightarrow 2) \rceil \rfloor$
attr, adv-a, ...	$\mathcal{N} \rightarrow \lfloor \mathcal{N} \rightarrow \mathcal{N} \rfloor$, where $\mathcal{N} = \lceil \mathcal{D} \rightarrow (\mathcal{S} \rightarrow 2) \rceil$
Sentence modifier	$\lfloor \lceil \mathcal{S} \rightarrow 2 \rceil \rightarrow \lceil \mathcal{S} \rightarrow 2 \rceil \rfloor$
adv-e, adv-f, ...	$\lceil \mathcal{D} \rightarrow (\mathcal{S} \rightarrow 2) \rceil \rightarrow \lfloor \mathcal{N} \rightarrow \mathcal{N} \rfloor$, where $\mathcal{N} = \lceil \mathcal{S} \rightarrow 2 \rceil$
Predicate nominalization operator	$(\mathcal{D} \rightarrow (\mathcal{S} \rightarrow 2)) \rightarrow \mathcal{K}$
Sentence nominalization operator	$(\mathcal{S} \rightarrow 2) \rightarrow \mathcal{D}$
* , **	$\lceil \mathcal{S} \rightarrow 2 \rceil \rightarrow [\mathcal{S} \rightarrow (\mathcal{S} \rightarrow 2)]$

Table 4.2: Semantic Types of EL Atomic Expressions

n occurrences of '$A \rightarrow$'. $\lceil\ \rceil$, $[\ ]$, and $\lfloor\ \rfloor$ indicate classes of persistent functions discussed shortly.

It is important to understand that this "type table" provides only a partial specification of what counts as an interpretation I of EL. While EL makes available an unlimited number of "freely interpretable" constants (of each type listed), individual variables, and modifiers, which need not satisfy any constraints other than those in Table 4.2, there are also many "special" atoms on which additional constraints are imposed via semantic clauses and/or meaning postulates. For instance, certain individual constants—e.g., numerals—denote certain specific individuals; certain function constants—e.g., $+$, $-$, $|$, tuple, etc.—denote certain specific functions; the sentence constants $\top$, $\bot$ denote truth and falsity respectively; predicate constants such as *episode* and *time* are interpreted so as to correspond with categories of the ontology, etc.

4.2.3 Classes of Persistent Functions

I now define certain subclasses of semantic functions in relation to the $\preceq$ and $\sqsubseteq$ orderings. This is needed to explain what we mean by an interpretation and by its extension to a valuation function.

(1) **$\preceq$-persistent functions $\lceil \mathcal{G} \rceil$**. A partial function $f \in \mathcal{D}^n \rightarrow (\mathcal{S} \rightarrow 2)$, where $n \geq 0$, is $\preceq$*-persistent* ("upward" persistent, with '$\preceq$' understood) iff for all $d_1, \ldots, d_n \in \mathcal{D}$ and $s, s' \in \mathcal{S}$ such that $s \preceq s'$,

$$f(d_1)\cdots(d_n)(s) = f(d_1)\cdots(d_n)(s')$$

whenever the LHS is defined. For example, suppose that *kiss* is interpreted as an upward persistent function *kiss'* of 3 arguments, the last being a situation. Then, if $kiss'(John')(Mary')(s) = 1$, then $kiss'(John')(Mary')(s') = 1$ as well, for any situation s' coextensive with s and containing s as a subsituation. We write $\lceil \mathcal{G} \rceil$ for the subclass of $\preceq$-persistent functions in class $\mathcal{G}$.

(2) **$\sqsubseteq$-persistent functions $[\mathcal{G}]$**. A partial function $f \in \mathcal{D}^n \rightarrow (\mathcal{S} \rightarrow 2)$, where $n \geq 0$, is $\sqsubseteq$*-persistent* (both "outward" and "inward" persistent, with '$\sqsubseteq$' understood) iff for all $d_1, \ldots, d_n \in \mathcal{D}$ and $s, s' \in \mathcal{S}$ such that $s \sqsubseteq s'$,

$$f(d_1)\cdots(d_n)(s) = f(d_1)\cdots(d_n)(s')$$

whenever the LHS is defined. For example, suppose "ended in 1945" is interpreted as a persistent function *end'45* of 2 arguments, the second being a situation. Suppose further that *end'45* (*WW2*) $(s) = 1$ for some situation s — say, the situation at the conclusion of World War II. Then by outward persistence, *end'45* (*WW2*) $(s') = 1$, where s' is the (exhaustive situation corresponding to) the 20th century. And hence, by inward persistence, *end'45* (*WW2*) $(s'') = 1$, where s'' is (the exhaustive situation corresponding to) the year 1992. By the same token, *end'45* (*WW2*) $(s') = 1$, for *any* s' whatsoever, if the LHS does have a value at all, and *end'45* (*WW2*) $(s) = 1$ for *some* s. In other words, *end'45* (*WW2*) is "eternally true," if true at all. We write $[\mathcal{G}]$ for the subclass of $\sqsubseteq$-persistent functions in class $\mathcal{G}$. (Note: If we instead had said "$\cdots$ whenever the LHS = 1," or "$\cdots$ whenever the LHS = 0," we would have obtained the "outward persistent" or "inward persistent" functions, respectively.)

(3) **$\sqsubseteq$-persistence preserving functions $|\ \mathcal{G}\ |$**. A partial function $f \in (\mathcal{D}^m \rightarrow (\mathcal{S} \rightarrow 2)) \rightarrow (\mathcal{D}^n \rightarrow (\mathcal{S} \rightarrow 2))$, where $m, n \geq 0$, is $\sqsubseteq$*-persistence preserving* iff

for all $g \in [\mathcal{D}^m \rightarrow (\mathcal{S} \rightarrow 2)]$,
if $f(g)$ is defined, then $f(g) \in [\mathcal{D}^n \rightarrow (\mathcal{S} \rightarrow 2)]$.

(We use m=1, and arbitrary n in n-fold predicate modifiers, and m=n=0 in sentence modifiers.) We write $|\mathcal{G}|$ for the subclass of $\sqsubseteq$-persistence preserving functions in class $\mathcal{G}$.

As discussed in Chapter 2, the bounded/unbounded distinction plays an important role in persistence of information, i.e., in evaluating $[\![\Phi]\!]^s$. That is, outward persistence and inward persistence (modulo granularity) correspond to the distinction between bounded and unbounded predicates. More specifically, for Φ bounded (and *not* negatively polarized): if $[\![\Phi]\!]^s = 1$, then $[\![\Phi]\!]^{s'} = 1$ for any s' such that $s \sqsubseteq s'$; and if $[\![\Phi]\!]^s = 0$, then $[\![\Phi]\!]^{s'} = 0$ for any s' such that $s' \sqsubseteq s$. However, for Φ unbounded: if $[\![\Phi]\!]^s = 1$, then $[\![\Phi]\!]^{s'} = 1$ for any s' such that $s' \sqsubseteq s$; and if $[\![\Phi]\!]^s = 0$, then $[\![\Phi]\!]^{s'} = 0$ for any s' such that $s \sqsubseteq s'$.

The bounded/unbounded distinction is presumed to be a recursively computable property of EL formulas. However, note that the bounded/unbounded classification is not taken to be exhaustive; there could be sentences which are neither bounded nor unbounded, e.g., a conjunction of a bounded and an unbounded sentences.

4.3 Formal Semantics

4.3.1 Semantic Clauses: Extending an Interpretation

Besides the type constraints on interpretation I stated earlier, there are more specific constraints on I and on its extensions to a valuation $[\![\,]\!]_I$, which follow below. If $d \in \mathcal{D}$, I is an interpretation of the atomic symbols of the logic, and α is a variable, $I\,(\alpha{:}d)$ denotes the interpretation identical with I except that it interprets α as d (regardless of whether or not α already had a value under I). Also, if $\underline{d}$ is a tuple of n elements of $\mathcal{D}$, and $\underline{\Phi}$ consists of n variables, then $I(\underline{\Phi}{:}\underline{d})$ denotes the interpretation obtained from I by setting the denotations of those variables in $\underline{\Phi}$ *which have no prior values* to the corresponding individuals in $\underline{d}$ (e.g., make the assignments in lexicographic order of the variables). In other words, if some of the variables in $\underline{\Phi}$ have no values under I and some do, then $I(\underline{\Phi}{:}\underline{d})$ changes only the interpretations of the variables without prior values to the corresponding elements of $\underline{d}$, leaving denotations of variables with prior values unchanged. The following semantic clauses state constraints both on an interpretation I and on its extension to a valuation function $[\![\,]\!]_I$. Whenever $[\![\,]\!]$ occurs unsubscripted, it is an abbreviation for $[\![\,]\!]_I$. Also recall that only one index of possibility, $s \in \mathcal{S}$, is used in semantic clauses in EL, as opposed to the world-time indices used in possible-world semantics.

1. *Valuation of Atomic Expressions*

   If α is an atomic expression, then $[\![\alpha]\!] = I(\alpha)$.

   Remarks. See Table 4.2 for the semantic type of $I(\alpha)$.

2. *Valuation of Functional Expressions*

If π, α are expressions of type $[\![\pi]\!] \in A \to B$ and $[\![\alpha]\!] \in A$ for sets A, B (derived from the types in Table 4.2), then $[\![(\pi\ \alpha)]\!] = [\![\pi]\!]^{[\![\alpha]\!]}$, i.e., $[\![(\pi\ \alpha)]\!] = [\![\pi]\!]([\![\alpha]\!])$.

Remarks. More precisely, functions ought to be typed syntactically, in parallel with the set-theoretic types given in Table 4.2, and the present rule need be expressed in terms of the syntactic types of π and α, instead of their (in principle ambiguous) set memberships. But the intention should be clear.

Note that the above covers the sentence modifiers, $\Box$, (adv-p *necessary*), etc., and operators $*$ and $**$; however, additional constraints are imposed indirectly on their interpretations by the clauses that follow (see clauses *3, 4, 16*).

3. *Valuation of '*'*

For $s \in \mathcal{S}$,

(a) $[\![\Phi * \eta]\!]^s = 1$ only if $Actual([\![\eta]\!], s)$ and $[\![\Phi]\!]^{[\![\eta]\!]} = 1$
$\qquad = 0$ only if $Nonactual([\![\eta]\!], s)$ or $[\![\Phi]\!]^{[\![\eta]\!]} \neq 1$; and

(b) for the special case that $s \in \mathcal{H}$ (i.e., s is an exhaustive situation), these two conditionals ('only if's) become biconditionals ('iff's).

Remarks. Thus '*' denotes truth at an (actual) situation. As shown in Table 4.2, the "input" formula for '*' will always be $\preceq$-persistent. The previously given set membership requirement for $I(*)$ in Table 4.2 ensures that the intension of a sentence of form $[\Phi * \eta]$ (with the intension of Φ upward persistent) is $\sqsubseteq$-persistent, i.e., unlocated. This is consistent with the additional requirements on $[\![\Phi * \eta]\!]$ inasmuch as the two conditionals can be consistently strengthened to biconditionals, and in that case, $\sqsubseteq$-persistence is a *consequence* of these biconditionals. This follows from the basic axioms for *Actual* and *Nonactual*, namely,

$$Actual(d, s)\ \&\ s \sqsubseteq s' \supset Actual(d, s'),$$
$$Nonactual(d, s)\ \&\ s \sqsubseteq s' \supset Nonactual(d, s').$$

4. *Valuation of '**'*

For $s \in \mathcal{S}$,

(a) $[\![\Phi ** \eta]\!]^s = 1$ iff $[\![\Phi * \eta]\!]^s = 1$, and
there is no $r \sqsubset [\![\eta]\!]$ such that $[\![\Phi]\!]^r = 1$;
$\qquad = 0$ iff $[\![\Phi * \eta]\!]^s = 0$, or
for some $r \sqsubset [\![\eta]\!]$, $[\![\Phi]\!]^r = 1$; and

(b) for the special case that $s \in \mathcal{H}$ (s is an exhaustive situation), these two conditionals ('only if's) become biconditionals ('iff's).

Remarks. Again, the $\sqsubseteq$-persistence constraint expressed by the set membership requirement is consistent with the further requirements, in that it becomes deducible if the two conditionals are made biconditionals.

5. *Interpretation of Other Special Functors (i.e., Logical Words)*

We state only three examples.
For all $d, e \in \mathcal{D}$ and $h \in \mathcal{H}$,

(1) $I(\text{episode})(d)(h) = 1$ iff $d \in \mathcal{S}$ and $Actual(d,h)$;
$= 0$ iff $d \notin \mathcal{S}$ or $Nonactual(d,h)$.

(2) $I(\text{during})(d)(e)(h) = 1$ iff $d, e \in \mathcal{S}$, $Actual(d,h)$, $Actual(e,h)$, and $Clocktime(d) \subseteq Clocktime(e)$;
$= 0$ iff $d \notin \mathcal{S}$, or $e \notin \mathcal{S}$, or $Nonactual(d,h)$, or $Nonactual(e,h)$, or $Clocktime(d) \not\subseteq Clocktime(e)$,

(3) $I(\text{time})(d)(h) = 1$ iff $d \in \mathcal{I}$ and $Actual(d,h)$;
$= 0$ otherwise,

where

$I(\text{episode}) \in \lceil \mathcal{D} \rightarrow (\mathcal{S} \rightarrow 2) \rceil$,
$I(\text{during}) \in \lceil \mathcal{D}^2 \rightarrow (\mathcal{S} \rightarrow 2) \rceil$, and
$I(\text{time}) \in \lceil \mathcal{D} \rightarrow (\mathcal{S} \rightarrow 2) \rceil$.

Remarks. Note in (3) that only *actual* times satisfy the predicate. The above "functors" are all predicative (ultimately yielding sentence intensions) rather than "functional" in the sense of mapping individuals to individuals. We use few functional symbols, in this sense, in EL. One we do use is the pairing function (e.g., in interpreting actions, as seen in Chapter 3). Its semantics is $[\![[\alpha \mid \beta]]\!] = \langle [\![\alpha]\!], b_1, \ldots, b_m \rangle$ where $\langle b_1, \ldots, b_m \rangle = [\![\beta]\!]$. This is an $(m+1)$ tuple that is defined whenever both arguments are defined.[5] Also, arithmetic functions have been mentioned in Chapter 3, which we find convenient, but they have little bearing on the concerns in this thesis.

Other special functors include `subep-of`, `coexten-subep-of`, `clock-time-of`, `subset-of`, `actual`, $\top$, $\bot$, `K`, `Ka`, `Ke`, `That`, and others. The detailed algebraic specifications of kinds, propositions, etc., remain to be completed. An important

[5]Note that $\langle d \rangle$ is not distinct from d. Thus, the "theory of tuples" is of the sort used in most of mathematics, rather than like the theory of lists in lisp or prolog.

point about the last four functors listed (i.e., K, Ka, Ke, and That) is that, like * and **, they produce time-independent values.

In the remaining clauses, Φ and Ψ are formulas, α is a variable, η is a term, and s *is a situation* $\in \mathcal{S}$. Also, it will be convenient in the semantics of '$\exists$' and 'The' to let $\mathcal{D}_{I(\alpha)}$ denote $\mathcal{D}$ if $I(\alpha)$ is undefined, and singleton set $\{I(\alpha)\}$ otherwise.

6. *Valuation of Negation*

$[\![\neg\Phi]\!]^s = 1$ iff $[\![\Phi]\!]^s = 0$;
$= 0$ iff $[\![\Phi]\!]^s = 1$.

Remarks. This is as in any standard logic (apart from truth-value gaps).

7. *Valuation of Conjunction*

Using $\underline{\alpha}$ for the parameter set $\underline{\Phi \wedge \Psi}$, and
$\mathcal{D}^{|\underline{\alpha}|}$ for $\mathcal{D}\times\mathcal{D}\times\cdots\times\mathcal{D}$ (n times), where $n = |\underline{\alpha}|$, the cardinality of set $\underline{\alpha}$:
$[\![\Phi \wedge \Psi]\!]^s = 1$ iff for some $\underline{d}\in\mathcal{D}^{|\underline{\alpha}|}$, $[\![\Phi]\!]^s_{I(\underline{\alpha}:\underline{d})} = [\![\Psi]\!]^s_{I(\underline{\alpha}:\underline{d})} = 1$;
$= 0$ iff for all $\underline{d}\in\mathcal{D}^{|\underline{\alpha}|}$, either $[\![\Phi]\!]^s_{I(\underline{\alpha}:\underline{d})} = 0$, or $[\![\Psi]\!]^s_{I(\underline{\alpha}:\underline{d})} = 0$.

Remarks. This semantics is perfectly consistent with the intuition that "This soup is hot and salty" means that the soup is hot and salty simultaneously, while "John packed his suitcase and left" does not necessarily mean that John's packing and leaving occurred at the same time. The intuitive contrast here is accounted for if we assume that "hot" and "salty" are stative, while "packing" and "leaving" are not, and that only statives are of necessity inward persistent. Also, note that formulas like $[[[(+\ 2\ 2){=}4] \wedge (\neg[(+\ 2\ 2){=}4])] * E_1]$ are 0.

The above allows for forward and backward anaphora, through the "parameter" mechanism. In effect, they are evaluated as if existential quantifiers at the highest level in Φ and Ψ had wide scope over the entire conjunction. (So, "Some man x is ill, and x is coughing" is evaluated as "For some man x, x is ill and x is coughing.") Note again that according to these truth conditions, names of existentially quantified variables matter. For example, $[(\exists x{:}\Phi\Psi) \wedge (\exists x{:}\Phi\,\neg\Psi)]$ will be logically false, since there is only one parameter, x, which is varied "simultaneously" in both conjuncts. On the other hand, $[(\exists x{:}\Phi\,\neg\Psi) \wedge (\exists y{:}\Phi\,\neg\Psi)]$ has the usual truth conditions (unless x, y have "prior" values, for instance, as a result of embedding within a wider-scope Qx or Qy quantification, or as a result of being a part of a larger conjunction or conditional which contains other occurrences of $\exists x$ or $\exists y$.

8. *Valuation of Disjunction*

$$[\![\Phi \vee \Psi]\!]^s = 1 \text{ iff } [\![\Phi]\!]^s = 1 \text{ or } [\![\Psi]\!]^s = 1;$$
$$= 0 \text{ iff } [\![\Phi]\!]^s = 0 \text{ and } [\![\Psi]\!]^s = 0.$$

Remarks. This is as in most versions of situation semantics, e.g., [Fenstad *et al.*, 1987] or [Devlin, 1991].

9. *Valuation of Existential Formulas*

$$[\![(\exists\alpha{:}\,\Phi\Psi)]\!]^s = 1 \text{ iff for some } d \in \mathcal{D}_{I(\alpha)}, [\![\Phi \wedge \Psi]\!]^s_{I(\alpha:d)} = 1;$$
$$= 0 \text{ iff for all } d \in \mathcal{D}_{I(\alpha)}, [\![\Phi \wedge \Psi]\!]^s_{I(\alpha:d)} = 0.$$

Remarks. This allows for "referential" occurrences of indefinites, when $I(\alpha)$ is defined [Fodor and Sag, 1982]. Intuitively, indefinites seem to require truth of the restriction predicate throughout the episode described, except possibly at the end (when the main clause describes a "culmination"). Examples are "A child grew up," "A bubble burst," "An actor retired," "An ice cube melted," etc.[6] This observation is compatible with the semantics of nominal predicates in EL. These are inward persistent (stative), but this will cause no problems, *provided* we do *not* insist that inward persistence also entails truth at the end-*points* of the interval in question. Such a view of statives seems reasonable in general. For example, "John was asleep for the past hour" seems not to require that he was asleep at the end-*points* of the 1-hour episode, only that he was asleep during all sub*intervals* of that episode. (And we might have a notion of "grain size," dependent on predicates used that says that the inward persistence holds only for subintervals significantly larger than the grain size.) With such a view of inward persistence, "a child grew up" does not entail being a child at the endpoint.

Concerning examples like "A planet formed," "He became an adult," "Mary baked a cake," or "The institution graduated five accountants this year," we feel that these require an intensional account, i.e., the verbs are essentially predicate operators, operating on predicative rather than quantificational interpretation of their NP operands. Evidence for the nonquantificational interpretations of these NP operands can be found in the contrasting valid and invalid inferences from the progressive forms: A child was growing up, therefore there was a child; A planet was forming, *therefore there was a planet; etc.

[6]There are apparent counterexamples, e.g., "The institution graduated *five accountants* this year," "Some *actors* did not go to high school," etc. In the former sentence, they became accountants at the end of graduation; in the latter sentence, they were not actors at the beginning of the implicit reference time. This will be discussed shortly.

For sentences like "A Messiah was/will be born in Bethlehem," "A little boy (who used to live here) is now a man," "A certain violinist was once a child prodigy," and "Some actors did not go to high school" (*cf.* [Enç, 1981]), it appears there are two additional phenomena involved: (i) A nominal (predicate) operator *sometime*, where [α (*sometime* π)] means that "α is at some time a π," i.e., it will be true at *all* times, if it is ever true (so the result is *atemporal*, i.e., eternal). This appears to be involved in the first sentence, and in most sentences with *event*-nominals, such as "A supernova was observed at Palomar." (ii) A scope phenomenon, in which a quantifier escapes from the scope of the tense operator, thereby getting its temporal reference from the utterance event; this seems to be involved in one reading of *a certain violinist* in the third sentence, *viz.*, the reading where this has *present* reference. (In another reading, (*sometime violinist*) is involved.) However, this is tentative in view of the discussion in Chapter 2 and Section 3.3.4. We allow for the possibility of implicit predicate operators; not just *sometime*, but also tense-like ones— *present*, *past/former/erstwhile*, *future/prospective* — or, maybe even the sentential tense operators with use of orienting relations, *and* scope phenomena.

10. <u>*Valuation of* The-*formulas*</u>

$$\begin{aligned}[\![(\text{The } \alpha{:}\, \Phi\Psi)]\!]^s = 1 \ \text{ iff } & [\![\Phi]\!]^s_{I(\alpha:d)} \text{ is defined for all } d \in \mathcal{D}_{I(\alpha)}, \\ & \text{there is a unique } d \in \mathcal{D}_{I(\alpha)} \text{ satisfying } [\![\Phi]\!]^s_{I(\alpha:d)} = 1, \text{ and} \\ & [\![\Phi \wedge \Psi]\!]^s_{I(\alpha:d)} = 1\,; \\ = 0 \ \text{ iff } & [\![\Phi]\!]^s_{I(\alpha:d)} \text{ is defined for all } d \in \mathcal{D}_{I(\alpha)}, \\ & \text{there is a unique } d \in \mathcal{D}_{I(\alpha)} \text{ satisfying } [\![\Phi]\!]^s_{I(\alpha:d)} = 1, \text{ and} \\ & [\![\Phi \wedge \Psi]\!]^s_{I(\alpha:d)} = 0\,. \end{aligned}$$

Remarks. This leads to a "referential" reading if α has a prior value, and a situation-dependent "Russellian" one otherwise. If there is no unique individual satisfying the restriction, the formula is truth-valueless.[7] We might have chosen falsity at least for the cases where there are no such individuals at all, by writing the RHS of the falsity condition as

iff for all $d \in \mathcal{D}_{I(\alpha)}$, $[\![\Phi \wedge \Psi]\!]^s_{I(\alpha:d)} = 0$.

Note that for referential occurrences of The and $\exists$, there is only a slight semantic difference: if the (predetermined) referent fails to satisfy the restriction, the $\exists$-formula is false, while the The-formula is truth-valueless.

[7] In that sense, we might say that the above truth conditions are only half-Russellian, i.e., they conform with Russell's conditions for truth, but not falsity.

11. Valuation of Universal Formulas

$$[\![(\forall\alpha\colon\Phi\Psi)]\!]^s = 1 \text{ iff for all } d\in\mathcal{D},\ [\![\Phi]\!]^s_{I(\alpha:d)} = 0 \text{ or } [\![\Phi\wedge\Psi]\!]^s_{I(\alpha:d)} = 1$$
$$= 0 \text{ iff for some } d\in\mathcal{D},\ [\![\Phi]\!]^s_{I(\alpha:d)} = 1 \text{ and } [\![\Phi\wedge\Psi]\!]^s_{I(\alpha:d)} = 0.$$

Remarks. Note that if Ψ is outward persistent, the conditions for truth allow the Ψ-subepisodes corresponding to various Φ-instances to be arbitrarily dispersed in time, as in "Every fisherman caught a fish." (Here, Ψ is "individual x catches a fish," an outward persistent achievement sentence.) However, the restriction sentence Φ is stative (if derived from an English nominal), so the restriction will select those individuals for which the restriction applies uniformly *throughout* s (at all subepisodes). Thus, the "fishermen" in the example have that property throughout the comprehensive episode (wherein each of them caught a fish).[8] By the same token, if the matrix sentence Ψ is stative (as in "Every fisherman remained luckless"), then universal truth requires Ψ to be true for all Φ-instances *throughout* s. (Again, *sometime* and scoping phenomena can lead to apparent violations of these observations.)

Note that the conditions for truth in effect require the truth value of the restriction Φ to be determinate for *all* individuals in the domain of discourse (as in the semantics of The). They also deny that a universal statement $(\forall\alpha{:}\Phi\Psi)$ can be true in a situation in which Φ is true and Ψ has no value, or for which Ψ is false and Φ has no value. If these possibilities are not denied, a non-persistent semantics results. For instance, one might then judge "Every child is asleep" true in a certain limited situation (e.g., a certain house at a certain time) merely because that situation determines definite truth values for the predicates *child* and *asleep* for very few actual children. In a more comprehensive situation, the same sentence might then be false. Given the above commitment to a persistent $\forall$-semantics, we would not regard

(pres ($\forall x$: [x child] [x asleep]))

as correctly expressing the meaning of "Every child is asleep," where the intent is to quantify only over individuals in a limited, contextually salient locale. Rather, the correct translation would make this locale explicit, e.g., as in

(pres ($\forall x$: [[x child] $\wedge$ (pres[x in-loc House])] [x asleep])).

[8] Apparent counterexamples such as "Last year every newborn baby at the Misericordia Hospital was put in an oxygen tent" are tentatively viewed as involving implicit *sometime* operators operating on the nominal.

12. *Valuation of* Most-*formulas*

$$[\![(\text{Most}\,\alpha\!:\Phi\Psi)]\!]^s = 1 \text{ iff for more than half of those } d \in \mathcal{D} \text{ such that } [\![\Phi]\!]^s_{I(\alpha:d)} \neq 0,\ [\![\Phi \wedge \Psi]\!]^s_{I(\alpha:d)} = 1;$$
$$= 0 \text{ iff for at least half of those } d \in \mathcal{D} \text{ such that } [\![\Phi]\!]^s_{I(\alpha:d)} \neq 0,\ [\![\Phi]\!]^s_{I(\alpha:d)} = 1 \text{ and } [\![\Phi \wedge \Psi]\!]^s_{I(\alpha:d)} = 0.$$

Remarks. Again, this semantics is designed to be persistent. It avoids making the LHS true merely because s only "knows about" very few of the existing Φ's, and those it "knows about" are mostly Ψ's. Conditions similar to clauses *11* & *12* can be formulated for other *monotone increasing* quantifiers, such as *many, at least one hundred, etc.* These are quantifiers Q such that for unlocated Φ, Ψ, and Υ, $(Q\alpha\!: \Phi\Psi)$ and $(\forall\alpha\!: \Psi\Upsilon)$ together entail $(Q\alpha\!: \Phi\Upsilon)$ (see [Barwise and Cooper, 1981]). We assume that non-monotone increasing quantifiers can be recast in terms of monotone increasing ones and negation. Thus, $(\text{Few}\,\alpha\!: \Phi\Psi)$ becomes $\neg\,(\text{Many}\,\alpha\!: \Phi\Psi)$, and $(\text{Exactly-}n\,\alpha\!: \Phi\Psi)$ becomes $(\text{At-least-}n\,\alpha\!: \Phi\Psi) \wedge \neg(\text{More-than-}n\,\alpha\!: \Phi\Psi)$, etc. (In the current version of EL syntax, numbers are not treated as quantifiers. But derived quantifiers such as these cardinal ones are under consideration for possible inclusion in EL.) The above truth conditions are rough-and-ready inasmuch as they do not have a mathematically clear meaning when infinitely many individuals satisfy the restriction clause. Ultimately a measure-theoretic approach should be used.

13. *Valuation of Material Conditionals*

$$[\![\Phi \rightarrow \Psi]\!]^s = 1 \text{ iff } [\![\Phi]\!]^s = 0 \text{ or } [\![\Phi \wedge \Psi]\!]^s = 1;$$
$$= 0 \text{ iff } [\![\Phi]\!]^s = 1 \text{ and } [\![\Phi \wedge \Psi]\!]^s = 0.$$

Remarks. Through the clause for conjunction, this plausibly handles many non-generic donkey sentences, such as "If Pedro owns a donkey, he will ride it to town tomorrow" (*cf.,* [Schubert and Pelletier, 1989]). This will *not* require Pedro to ride *all* his donkeys to town. Rather it only requires that there be *some* donkey which he both owns *and* rides to town, if he has any donkeys. The iteration over candidate donkeys is done implicitly through the iteration over parameter values in clause 7, where these values preempt the existential quantifier for the donkey.

14. *Valuation of λ-expressions*

For π a formula or n-place predicative expression ($n \geq 1$),
$[\![\lambda\alpha\,\pi]\!] = \{<d, [\![\pi]\!]_{I(\alpha:d)}> \mid d \in \mathcal{D},\ [\![\pi]\!]_{I(\alpha:d)} \text{ defined}\}$.

Remarks. Thus, for instance, $[\![\lambda x[x \text{ tornado}]]\!] = [\![\text{tornado}]\!]$. (Note, however, that $[\![\lambda e\ [[\text{John kiss Mary}] * e]]\!] \neq [\![[\text{John kiss Mary}]]\!]$. The LHS is of type $\mathcal{D}\rightarrow(\mathcal{S}\rightarrow 2)$, while the RHS is of type $\mathcal{S}\rightarrow 2$.)

15. *Valuation of Equalities*

For τ_1, τ_2 terms,

$$\begin{aligned}[\![\tau_1 = \tau_2]\!]^s &= 1 \text{ iff } [\![\tau_1]\!] = [\![\tau_2]\!], \text{ with both defined and } \textit{actual} \text{ or } \textit{nonactual} \text{ relative to } s\,;\\ &= 0 \text{ iff } [\![\tau_1]\!] \neq [\![\tau_2]\!], \text{ with both defined and } \textit{actual} \text{ or } \textit{nonactual} \text{ relative to } s.\end{aligned}$$

Remarks. Note that these conditions allow an equality to be undefined in some situations, even when both of the terms equated have determinate and equal denotations—namely, those situations in which the entities denoted simply play no role (are non-participants).

16. *Valuation of Necessity Formulas*

$$\begin{aligned}[\![\Box\Phi]\!]^s &= 1 \text{ only if for all exhaustive situations } h \in \mathcal{H},\ [\![\Phi]\!]^h = 1\,;\\ &= 0 \text{ only if for some exhaustive situation } h \in \mathcal{H},\ [\![\Phi]\!]^h = 0.\end{aligned}$$

Remarks. Note that since $\Box$ is a sentence modifier, it is also constrained to be of type $|\lceil\mathcal{S}\rightarrow 2\rceil \rightarrow \lceil\mathcal{S}\rightarrow 2\rceil|$, so that $\Box$ is persistence preserving.

$\Box\,(\Phi \rightarrow \Psi)$ guarantees that if Φ is true in a situation s, then there is a situation s' such that $s \preceq s'$ and Ψ is true in s'. That's all we need to make the inferences we want from meaning postulates.

17. *Valuation of Probabilistic (Generic) Conditionals*

$$\begin{aligned}[\![\Phi \rightarrow_{p,\alpha_1,\ldots,\alpha_n} \Psi]\!]^s &= 1 \text{ iff for "at least a proportion } p\text{" of elements}\\ &\qquad \underline{d} \in \mathcal{D}^n \text{ such that } [\![\Phi]\!]^s_{I(\underline{\alpha}:\underline{d})} \neq 0,\\ &\qquad [\![\Phi \wedge \Psi]\!]^s_{I(\underline{\alpha}:\underline{d})} = 1\,;\\ &= 0 \text{ iff for "more than a proportion } (1-p)\text{" of elements}\\ &\qquad \underline{d} \in \mathcal{D}^n \text{ such that } [\![\Phi]\!]^s_{I(\underline{\alpha}:\underline{d})} \neq 0,\\ &\qquad [\![\Phi]\!]^s_{I(\underline{\alpha}:\underline{d})} = 1 \text{ and } [\![\Phi \wedge \Psi]\!]^s_{I(\underline{\alpha}:\underline{d})} = 0\,;\\ &\text{where } \underline{\alpha} = \alpha_1, \ldots, \alpha_n\,.\end{aligned}$$

Remarks. Note the similarity to the semantics of `Most` in clause *12*. This is only a rough approximation to what is required. First, the talk of "proportions" needs to be replaced by a notion of measure, based on a distribution over individuals (as in the case of `Most`). And second, instead of using purely "extensional statistics" (proportions of Ψ-instances relative to Φ-instances at s), the truth conditions should in general be modalized to reflect the *nomic* character of many generic sentences;

this could be done by evaluating the antecedent and consequent not only at s, but also at "nearby" *nonactual* situations extending over the same space-time region as s.

The use of "controlled variables" $\alpha_1, \ldots, \alpha_n$ in this construct can solve the "proportion problem" mentioned in Chapter 3. For instance, "If a farmer owns a donkey, he is usually rich" could be expressed as in the following:

$$\begin{array}{l}(\exists x{:}\,[x \text{ farmer}] \\ \quad (\exists e{:}\,[e \text{ episode}]\ [(\exists y{:}\,[y \text{ donkey}]\ [x \text{ own } y]) \ ** \ e])) \\ \rightarrow_{.8,e} [[x \text{ rich}] * (\text{time-of } e)]\end{array}$$

Note that the $\exists$-quantifier of the y-variable (varying over donkeys) is inside the scope of '**'. Note also that y is not a controlled variable here. If the $\exists y$ quantifier were outside the scope of '**', and y were controlled, along with e, we would have obtained an intuitively implausible reading of the English sentence. For instance, the sentence would have been judged true in a situation where 99 farmers who each owns a donkey are poor and 1 farmer who owns 500 donkeys is rich.

One might wonder what the effect of controlling x (the farmer variable), as well as e, might be. Here, there would be no effect, since for every choice of e, the value of x is already determinate. This is because e is *minimal* with respect to its characterization, so that it supports the truth of that characterization for only one choice of x-value. If x were the *only* controlled variable, without e, we would get a bizarre reading to the effect "If a farmer owns a donkey at some time, then he is usually a farmer who owns a donkey and is rich at some time." For example, if most donkey owning farmers were briefly rich at some time, but poor most of their lives, we would not judge the original English sentence true, but the x controlled translation (with e not controlled) would render it true.

We ignore the problem about the duration of episodes, i.e., the problem of how to count stative episodes, for now. For example, suppose Pedro owned a donkey 5 times each for about a month, and he was rich during those periods. Next, he owned a donkey for 12 years, while being poor. So, there are 6 donkey owning episodes; during 5 of them he was rich, during one of them he was poor. The question is, can we say "When Pedro owned a donkey, was he usually rich"? Here, the answer should be "No," but we can't give the correct answer unless we incorporate measurement of durations of episodes. We should consider the *durations of those 5 1-month owning episodes* versus the *duration of that one 12-year episode*; i.e., we should consider 5 months versus 12 years, rather than 5 episodes versus 1 episode. So, it seems, for stative e, we need to consider durations, in the future.

Despite the neglect of infinite sets of cases, and the nomic character of generic sentences, probabilistic conditionals provide a useful approximation to many generic sentences. In particular, they allow probabilistic conclusions to be drawn based on the much-discussed inductive principle of *direct inference* (e.g., [Kyburg, 1983; Bacchus, 1988; Bacchus, 1990]). Given that Φ holds for particular values of the controlled variables, $\Phi \rightarrow_{p,\overline{\alpha}} \Psi$ allows us to conclude Ψ for those values, with degree of certainty $\geq p$ (at least in the absence of other information).

This completes our enumeration of truth conditions. Primary omissions are the semantics of many quantifiers (such as 'Few' which can be "filled in" more or less analogously as mentioned under clause *12*), some sentential connectives (especially, *because*), and various functors (e.g., the specific properties of various sentence modifiers and various nominalization operators and their inverses). Note that there are as well additional operators in EL relating to questions and *wh*-nominals which we have only mentioned and not further addressed, for lack of semantic details. Next we state two persistence theorems, and then discuss entailment and anti-entailment in EL.

4.3.2 Persistence Theorems

We distinguish unlocated predicate constants such as =, *cause-of, before, number,* etc., from located ones such as *walk, girl, kiss, popular, dead,* etc. Intuitively, the unlocated predicates are those that either hold for given arguments everywhere, at all times (in any given world), or nowhere, at no time *j*, whereas the truth of located predicates for given arguments is place and time dependent, at least for some arguments, in some worlds. Atomic unlocated predicates are so marked in the logical lexicon.

An *unlocated expression* is defined as an EL expression such that (1) any located predicate occurring in it lies within the scope of one or more of {*, **, K, Ka, Ke} (unlocated predicates such as =, *cause-of,* and *before* can occur anywhere in an unlocated formula) and (2) an expression of form $[\Phi * \eta]$ is unlocated provided that no subterm of η occurs in Φ (η assumed to be nonanaphoric). Here are the two fundamental persistence theorems. Given that a formula describes a certain situation, they allow us to infer that it also describes any "enlarged" situations.

1. *Upward persistence of formulas.* If Φ is a formula, and s, s' are situations such that $s \preceq s'$, then $[\![\Phi]\!]^s = [\![\Phi]\!]^{s'}$ if the LHS is defined.

2. *Persistence of unlocated formulas.* If Φ is an unlocated formula, and s, s' are situations such that $s \sqsubseteq s'$, then $[\![\Phi]\!]^s = [\![\Phi]\!]^{s'}$ if the LHS is defined.

4.3.3 Entailment and Anti-entailment

Entailment (truth preservation) in EL is defined for a set of premises $\Phi_1, \ldots, \Phi_n$ and a conclusion Ψ as follows.

$$\Phi_1, \ldots, \Phi_n \models \Psi \ , \quad \text{for } \Phi_1, \ldots, \Phi_n, \Psi \ \text{ unlocated,}$$

iff $[\![\Psi]\!]_I^w = 1$, whenever (i) I is an interpretation which assigns denotations to all constants and parameters of all Φ_i and to all constants and free variables of Ψ; (ii) $[\![\,]\!]_I$ is a valuation function extending I; and (iii) w is a world $\in \mathcal{W}$ such that $[\![\Phi_i]\!]_I^w = 1$, for $i = 1, \ldots, n$.

This is more general than needed for the inference rules that will be proposed in Chapter 5, covering anaphora, e.g.,

$$\begin{aligned}
&(\exists x{:}\,[x \text{ girl}]\ (\forall y{:}\,[y \text{ boy}]\ [y \text{ love } x])),\ [\text{John boy}] \\
&\quad \models [\text{John love } x]\,, \\
&\quad \models (\exists z\ [\text{John love } z])\,, \\
&\quad \models [[y \text{ boy}] \rightarrow [y \text{ love } x]]\,, \ \text{ etc.}
\end{aligned}$$

Note that by defining entailment for worlds, it essentially applies only to unlocated formulas (since located ones are generally not true in world, i.e., for all time, but rather for only a limited time).

Anti-entailment (falsity preservation), $\Phi \mathrel{=\!\!\mid} \Psi$, is defined as follows, much like entailment.

$$\Phi \mathrel{=\!\!\mid} \Psi \ , \quad \text{for } \Phi, \Psi \ \text{ unlocated,}$$

iff $[\![\Psi]\!]_I^w = 0$, whenever (i) I is an interpretation which assigns denotations to all constants and parameters of Φ and to all constants and free variables of Ψ; (ii) $[\![\,]\!]_I$ is a valuation function extending I; and (iii) w is a world $\in \mathcal{W}$ such that $[\![\Phi]\!]_I^w = 0$.

Thus, for example,

$$(\exists x{:}\,[x \text{ boy}]\ [x \text{ love Mary}]) \ \mathrel{=\!\!\mid}\ [[y \text{ boy}] \rightarrow [y \text{ love Mary}]]\,.$$

That is, if it is false that "There is a boy who loves Mary," we infer it is false that "If y is a boy, he loves Mary."

We will now prove the following entailment, as an illustration of how the truth conditions for times interact with the persistence theorems *1* & *2* above:

Claim (Inward persistence of negated outward persistent formula).

For φ an outward persistent formula and η a term,
$[(\neg\varphi) * \eta] \models (\forall t\colon [[t \text{ time}] \wedge [t \text{ during } \eta]] (\neg [\varphi * t]))$.

Proof. For any world $w \in \mathcal{W}$, $[\![(\neg\varphi) * \eta]\!]^w = 1$ implies $[\![\eta]\!] \sqsubseteq w$ and $[\![\neg\varphi]\!]^{[\![\eta]\!]} = 1$ (by semantics of '*', clause *3*). Hence if i is the time concurrent with $[\![\eta]\!]$ in w, $[\![\neg\varphi]\!]^i = 1$ (by upward persistence). Hence, for all times $j \sqsubseteq i$, $[\![\varphi]\!]^j = 0$ (by outward persistence of φ, and the semantics of negation, clause *6*). Hence for all times $j \sqsubseteq i$, $[\![\varphi * t]\!]^j_{I(t:j)} = 0$ (by clause *3*). Hence for all times $j \sqsubseteq i$, $[\![\neg[\varphi * t]]\!]^i_{I(t:j)} = 1$ (by *6*). Hence for all times $j \sqsubseteq i$, $[\![\neg[\varphi * t]]\!]^w_{I(t:j)} = 1$ (because '*' yields a result in $[\mathcal{S}\rightarrow(\mathcal{S}\rightarrow 2)]$, and clause 6 preserves the $\sqsubseteq$-persistence property). Then by the constraints placed in clause 5 on $I(\textit{time})$ and $I(\textit{during})$, and the truth conditions for '$\forall$' (clause *11*) and '$\wedge$' (clause *7*), $[\![(\forall t\colon[[t \text{ time}] \wedge [t \text{ during } \eta]] (\neg[\varphi * t]))]\!]^w = 1$. □

This result justifies a meaning postulate,

$$\Box\, [(\neg\varphi) * \eta] \models (\forall t\colon [[t \text{ time}] \wedge [t \text{ during } \eta]] (\neg [\varphi * t])),$$

which is quite useful. For example, it allows us to infer that if John did not leave in an episode spanning yesterday, then there is no time *during* yesterday at which John left.

We now close this section with some more valid EL schemas without proof.

Some Valid EL *Schemas*

1. a. ($\forall$e: [e episode] [e coexten-subep-of e])
   b. ($\forall$e ($\forall$e′ [[e coexten-subep-of e′] → [e subep-of e′]]))
   c. ($\forall$e ($\forall$e′: [e′ episode] [[e coexten-subep-of e′]
   ↔ [(region-of e) subset-of (region-of e′)]]))

2. [Φ * η] → ($\exists$e: [e coexten-subep-of η] [Φ ** e]), for Φ stative
   → ($\exists$e: [e subep-of η] [Φ ** e])

3. ($\forall$e ($\forall$e′: [e coexten-subep-of e′] [[Φ * e] → [Φ * e′]]))

4. For Φ telic or factual,
   ($\forall$e ($\forall$e′: [e subep-of e′] [[Φ * e] → [Φ * e′]]))

5. [Φ ** η] ↔ [[Φ * η] $\wedge$ ¬ ($\exists$e: [e proper–subep-of η] [Φ * e])]

6. ($\exists$x [Φ * η]) ↔ [($\exists$x Φ) * η]

7. [($\exists$x Φ) ** η] → ($\exists$e: [η coexten-subep-of e] ($\exists$x [Φ ** e]))

8. $[(\neg\Phi) * \eta] \rightarrow \neg[\Phi * \eta]$

9. $[(\neg\Phi) * \eta] \rightarrow \neg(\exists e{:}\,[[e\ \text{time}] \wedge [e\ \text{during}\ \eta]]\,[\Phi * e])$

10. $[[\Phi \wedge \Psi] * \eta] \rightarrow [\Phi * \eta]$, with no cataphor
$\rightarrow [\Psi * \eta]$, with no anaphor

11. For Φ unlocated,
    a. $[\Phi * \eta] \rightarrow \Phi$
    b. $\Phi \rightarrow (\forall t{:}\,[t\ \text{time}]\,[\Phi * t])$

4.4 Summary and Remarks

I have shown the model structures and the semantics of EL. One of the major categories in EL ontology is the set of possible situations. Depending on the nature of spatiotemporal locations they occupy and the amount of information they support, situations may be classified into exhaustive situations $\mathcal{H}$ (situations with maximal informational content), times $\mathcal{I}$ (situations spatially unbounded, with maximal informational content), and worlds $\mathcal{W}$ (situations both spatially and temporally unbounded, with maximal informational content), where $\mathcal{W} \subset \mathcal{I} \subset \mathcal{H}$. Two partial orderings on situations based on regions and informational content are $\sqsubseteq$ (general part-of relation between a pair of situations) and $\preceq$ (informational part-of relation between a pair of coextensive situations). $\sqsubseteq$ forms a join semilattice with join operator $\sqcup$, with respect to each set $\{s \mid s \sqsubseteq w\}$, for $w \in \mathcal{W}$.

In the semantics, some important features of EL that facilitate interpretation and valuation of ELFs are: a DRT-like parameter mechanism, *typed* EL expressions, and classes of persistent functions.

It should be noted, though, that the semantics of EL is still under development, and there are still some uncertainties and gaps in the logical semantics (as in any situation theory currently being developed). However, EL at least subsumes classical logic, provides tentative extensions in several major directions, and is sufficiently carefully formalized to allow future systematic analysis and revision. Moreover, the semantic ideas were not conceived in isolation, but with an eye on the mapping from surface structure to logical form (which will be discussed in Part II) and on facilitating the inferences that support narrative understanding (see Chapter 5). In these respects, the development of EL is at least a step in the right direction.

Chapter 5

Rules of Inference in Episodic Logic

I have discussed at length the NL-like *expressiveness* of EL. This also provides a basis for concise, easily understood inferences. The main inference rules in EL are RI (Rule Instantiation) and its dual GC (Goal Chaining). These are generalizations of what are commonly referred to as "forward inference" (or "input-driven inference" or "spontaneous inference") and "backward chaining" (or "goal-driven inference") in AI terminology. In addition, natural deduction is used in goal-driven, i.e., "backward," inference. In this chapter, I show RI and GC, illustrating them with examples.

5.1 Input-driven and Goal-driven Inferences

In the course of understanding natural language or answering questions, humans automatically make obvious inferences. Inference rules in EL, RI and GC, are purported to generate such inferences so that the system can understand actual story fragments and answer questions. Since the final ELF is nonindexical, it can be used in concert with facts in a knowledge base to work out immediate consequences of new inputs and to answer questions. RI is heavily used in input-driven inference. It automatically generates "obvious" inferences people would spontaneously make given input. GC, the dual of RI, similarly dominates goal-driven inference, while trying to answer questions. These rules were formulated by looking at examples of what seems to follow in a single step (intuitively), given the "background facts" for the story. For instance, suppose we have the following pieces of knowledge.

(5.1) Every large carnivore is dangerous when it is hungry

(5.2) If a small animal yelps, it is either hungry or sad or sick

Suppose now Fenris the wolf is hungry. Then, we infer he is dangerous by (5.1). Next if we see Pluto the dog yelp, then we infer he is either hungry or sad or sick by (5.2). (Let us assume that wolves are large carnivores and that dogs are small animals.) This kind of inferences are input-driven or forward inference. On the other hand, suppose we want to know if Fenris is dangerous at the moment (if so, we wouldn't want to let Alice play with him!). Then we may be interested to know if he is hungry now. Similarly, if we want to find out if Pluto is sad, then one way of getting an answer is to check if he is neither hungry nor sick and yet is yelping.

The rules of inference in EL are classical. That is, all the nonclassical constructs, e.g., modals like beliefs, are handled by meaning postulates. This is done by the very general inference rules RI and GC. These resemble resolution except that they allow arbitrarily embedded quantifiers. Before stating the rules, however, we need to first fix some terminology.

A quantifier or a subformula occurs "positively" if it lies within an even number of negations, where conditional antecedents and $\forall$-quantifier restrictions count as negation.[1] Similarly, a quantifier or a subformula occurs "negatively" if it lies within an odd number of negations. This is illustrated in the following ELF representations of (5.1) and (5.2). '+' and '−' signs below indicate positive and negative occurrence of the embedded quantifiers or formulas, respectively.

(5.1) Every large carnivore is dangerous when it is hungry

a. ($\underline{\forall}^+$ x: [x ((attr large) carnivore)]$^-$
[($\underline{\exists}^-$ e1 [[x hungry] ** e1]$^-$)
→ ($\underline{\exists}^+$ e2 [[e1 same-time e2]$^+$ ∧ [[x dangerous] ** e2]$^+$])]])

(5.2) If a small animal yelps, it is either hungry or sad or sick

a. ($\underline{\exists}^-$ x: [x ((attr small) animal)]$^-$
($\underline{\exists}^-$ e1 [[x yelp] * e1]$^-$))
→ ($\underline{\exists}^+$ e2: [e1 same-time e2]$^+$
[[[x hungry] ** e2]$^+$ ∨ [[x sad] ** e2]$^+$ ∨ [[x sick] * e2]$^+$])

Note that the above informal definition of positive and negative occurrences of subformulas is intended to apply to "classically" embedded subformulas only. That is, the

[1]This +/− system is apparently due to Peirce (see [Roberts, 1973]), and has come up repeatedly in the automated theorem proving literature, e.g., [Bibel, 1979; Andrews, 1981; Traugott, 1986].

embedding operators are just the truth-functional connectives $\neg, \wedge, \vee$ and $\rightarrow$ and quantifiers $\exists$ and $\forall$. If a subformula is within the scope of any operators other than these, it occurs neither positively nor negatively. (One exception made later is that the consequent of a generic conditional is considered to occur positively and the antecedent negatively in that conditional. In this respect, $\Phi \rightarrow_{p,x_1,\ldots,x_n} \Psi$ is treated just like $\Phi \rightarrow \Psi$.) I now show basic inference rules RI and GC.

5.2 RI: Rule Instantiation

RI (Rule Instantiation), which is heavily used in input-driven inference, allows arbitrarily many minor premises to be matched against arbitrarily deeply embedded subformulas of a rule. It subsumes *modus ponens* and *modus tollens*, but can also instantiate probabilistic conditionals. I first state the rules formally, and then show how they work via examples.

Singly Instantiating the Rule

For comprehensibility, we first show rules with just *one* minor premise ("fact") with unit probabilities.

RI: Single Instantiation Cases
— Unit Probability versions —

$$\text{(I)}\quad \frac{R^-(\Phi),\ F^+(\Psi)}{R^-_\sigma(\neg\,(F^+_\sigma(\bot)))} \qquad\qquad \text{(II)}\quad \frac{R^-(\Phi),\ F^+(\Psi)}{F^+_\sigma(R^-_\sigma(\top))}$$

'R' stands for *Rule*, and 'F' for *Fact*. $R(\Phi)$ and $F(\Psi)$ are formulas with bound variables standardized apart. The '+' and '−' signs are intended to indicate positive and negative occurrence of the embedded Φ, Ψ formulas being unified. Thus "rule" and "fact" terminology is used only to suggest the typical use of these inference rules—there is no formal distinction between rules and facts. What matters is the "sign" of the embedding of Φ and Ψ. (The signs are marked in the rules above only as a reminder though; they have no separate mathematical roles.) At the bottom of both rules, σ is a substitution that unifies Φ and Ψ, and $\top$ and $\bot$ are truth and falsity respectively. (As discussed in Chapter 3, $\top$ and $\bot$ are formulas which uniformly *denote* truth values 1 and 0, respectively.) Unification of Φ with Ψ is defined in a way that allows substitution of arbitrary terms for explicitly quantified, "matchable" variables which occur free (unbound) in Φ or Ψ, but are bound in R or F as a whole.[2] A variable in a rule or fact is "matchable" if it

[2]In addition, for variables bound within Φ or Ψ, unification allows for variable renaming. This applies

is bound by a positively occurring universal quantifier or negatively occurring existential quantifier. When a term is substituted for a quantified variable, the quantifier is deleted. For instance, substitution of w for x in a positively embedded subformula ($\forall x$:[x P][x Q]) yields [[w P] $\rightarrow$ [w Q]], and the same substitution in a negatively embedded subformula ($\exists x$:[x P][x Q]) yields [[w P] $\wedge$ [w Q]].

Thus, rule (I) says: if a positively occurring subformula of fact F, Ψ, is unifiable with a negatively occurring subformula of rule R, Φ, with substitution σ, then do the substitution throughout R and F, replace the Ψ portion in rule F with falsity, negate the resultant F, and replace the Φ portion in rule R with this resultant F. Rule (II) says similarly.[3]

Rule (I) is *sound* if Ψ contains no unmatchable free variables which are bound in F as a whole. Rule (II) is *sound* if Φ contains no unmatchable free variables which are bound in R as a whole. So, in particular, rule (I) is sound if F contains only constants and top-level universal, hence matchable, variables. (The soundness proof can be found in [Schubert, in preparation].)

I now illustrate RI with some examples.

Examples (RI-I,II) : "Moby Dick is not a fish; Wanda is not a whale"

Consider rule

(5.3) a. No whale is a fish
b. ($\forall x$ [[x whale] $\rightarrow \neg$[x fish]])

Here x is a matchable variable since it is quantified by a positively occurring $\forall$-quantifier. [x whale] and [x fish], called, say, Φ_1 and Φ_2, occur negatively, while $\neg$[x fish] occurs positively. Suppose now we have the assertion

[Moby-Dick whale]

in the knowledge base. Then, by RI (I), with [x whale] in the role of Φ and [Moby-Dick whale] in the role of Ψ, the substitution $\langle$Moby-Dick/$x\rangle$ unifies Φ with Ψ, with result

($\neg\top$) $\rightarrow$ $\neg$[Moby-Dick fish],

which is simplified to $\neg$[Moby-Dick fish], i.e., "Moby Dick is not a fish." We would have obtained the same result with RI (II). (The results of (I) and (II) are easily seen to

to *all* variable-binding operators, including nonclassical ones like Many and λ. Such renaming must in principle preserve anaphoric connectives, but this complication can be avoided by converting formulas to a standard form not involving anaphora.

[3] Jon Traugott [1986] proposes essentially the same rule as RI. However, he assumes all quantifiers have been skolemized away.

be equivalent whenever $F^+(\Psi) = \Psi$, i.e., the matched "subformula" of F is the entire formula F.) Conversely, if we have

[Wanda fish]

in the knowledge base, unification with [x fish], with substitution ⟨Wanda/x⟩, yields the inference

[Wanda whale] $\rightarrow$ $\perp$,

which simplifies to ¬ [Wanda whale], i.e., "Wanda is not a whale."

Example (RI-I,II): "Oops, Fenris *does* and does *not* live in something??"

Let us now look at a slightly more complicated case. Consider the following examples (with their simplified ELFs).

(5.4) a. Fenris does not live in anything
b. (∀w: [w thing] (¬ [Fenris live-in w]))

(5.5) a. Every wolf lives in a lair
b. (∀x: [x wolf] (∃y [[y lair] ∧ [x live-in y]]))

Here, [Fenris live-in w] in (5.4 b) and [x live-in y] in (5.5 b) are candidates for matching. Note that [Fenris live-in w] is negatively occurring and that [x live-in y] is positively occurring, and hence that (5.4) should be taken as a rule and (5.5) as a fact to use RI. Next, notice that w in (5.4 b) and x in (5.5 b) are matchable variables, but y in (5.5 b) is not. Thus, we need to take rule (II) to generate a sound result, which yields (with substitution ⟨Fenris/x, y/w⟩):

[Fenris wolf] → (∃y [[y lair] ∧ (¬ [y thing])]).

This would simplify further to

(∃y [[y lair] ∧ (¬ [y thing])]),

by an application of rule (I), if we already have [Fenris wolf] in the knowledge base. So, what we would get is that there is some lair that is not a thing, which would be a contradiction if a lair is known to be a thing; i.e., we would conclude that (5.4) and (5.5) are together inconsistent with the knowledge base.

If we had used rule (I), with the same substitution, we would have gotten

$$[[\text{y thing}] \rightarrow [[\text{Fenris wolf}] \rightarrow (\exists \text{y} \; [[\text{y lair}] \wedge \bot])]], \text{ or,}$$

$$[[\text{y thing}] \rightarrow [[\text{Fenris wolf}] \rightarrow \bot]], \text{ or,}$$

$$[[\text{y thing}] \rightarrow \neg [\text{Fenris wolf}]],$$

which is not a closed formula, and leads to the unsound conclusion $\neg [y \text{ thing}]$, if it is known that [Fenris wolf].

Example (RI-I,II): "Every map is 4-colorable"

RI also applies to *axiom schemas*. The following axiom schema is a meaning postulate about modal predicate "know."

(5.6) For τ, η terms and Φ a sentence:
$[[\tau \text{ know (That } \Phi)] ** \eta] \rightarrow \Phi$

Now, suppose we are told that

(5.7) a. John knows that every map is 4-colorable
b. [[John know (That (∀x: [x map] [x 4-colorable]))] ** E7],
with some simplification

Instantiating the above axiom schema with this formula, we get the following, with unification ⟨John/x, (∀x: [x map] [x 4-colorable])/Φ, E7/η⟩:

(∀x: [x map] [x 4-colorable]).

Multiply Instantiating the Rule

Let us now further consider the special case of rules (I) and (II), with $F(\Psi) = \Psi$, i.e., Ψ is the same as the fact F as a whole. (This case was already illustrated in the "Moby Dick" example.) Then both rules are simplified to

$$\frac{R^{-}(\Phi),\ \Psi}{R^{-}_{\sigma}(\top)} \;.$$

If we substitute $\neg\Phi'$ for Φ and $\neg\Psi'$ for Ψ in this special case of rules (I) and (II), then the embeddings of Φ' and Ψ' are the "opposite" of that of Φ and Ψ (as in the previous "Wanda" example), and we get the following variant:

$$\frac{R^+(\Phi'),\ \neg\Psi'}{R^+_\sigma(\bot)}\ .$$

Suppose now we iterate the above two cases an arbitrary number of times. Then we get a rule for multiple instantiation that could be expressed schematically as follows.

`RI: A Multiple Instantiation Case`
— Unit Probability Version —

$$\text{(III)}\qquad \frac{R(\Phi_1,\ldots,\Phi_m,\Phi'_1,\ldots,\Phi'_n),\quad \Psi_1,\ldots,\Psi_m,\neg\Psi'_1,\ldots,\neg\Psi'_n}{R_\sigma(\top,\ldots,\top,\bot,\ldots,\bot)}$$

where $R(\Phi_1,\ldots,\Phi_m,\Phi'_1,\ldots,\Phi'_n)$, $\Psi_1,\ldots,\ \Psi_m,\ \neg\Psi'_1,\ldots,\ \neg\Psi'_n$, are formulas with bound variables standardized apart; all Φ_i's occur negatively in $R(\Phi_1,\ldots,\Phi_m,\ \Phi'_1,\ldots,\Phi'_n)$, all Φ'_i's occur positively in it; and substitution σ unifies the Φ_i with corresponding Ψ_i and Φ'_i with corresponding Ψ'_i. As before, the substitution σ may replace only matchable variables, i.e., those that are $\forall$-quantified by a positively occurring quantifier, or $\exists$-quantified by a negatively occurring quantifier, in $R(\Phi_1,\ldots,\Phi_m,\Phi'_1,\ldots,\Phi'_n)$ or one of the Ψ_i. Again, computing $R_\sigma(\top,\ldots,\top,\bot,\ldots,\bot)$ involves elimination of quantifiers of variables replaced by σ, e.g., if b is substituted for x, then $(\forall x{:}\Phi\Psi)_{b/x}$ becomes $[\Phi_{b/x} \rightarrow \Psi_{b/x}]$. $R_\sigma(\top,\ldots,\top,\bot,\ldots,\bot)$ is then simplified to eliminate the "truth values," $\top$ and $\bot$.

That is, RI (III) allows *arbitrarily many* minor premises to be matched against arbitrarily deeply embedded subformulas of a rule. (Apart from its avoidance of skolemization, and its restriction to matching of parts of *distinct* formulas, it resembles Andrews' [1981] *general matings* and Bibel's [1979] *connections*.) A rule instantiation typically instantiates the complete antecedent of a conditional and infers the particularized consequent, though it may in principle match only part of the antecedent, or match part or all of the consequent, giving a "contrapositive" inference. (The previously shown deduction of ¬ [Wanda whale] from fact [Wanda fish] is of this type.)

Example (RI-III) : "Fenris is dangerous"

As an example for RI (III) with m=2, let us consider the rule (5.1) again.

(5.1) Every large carnivore is dangerous when it is hungry

a. (∀x: [x ((attr large) carnivore)]
 [(∃e1 [[x hungry] ** e1])
 → (∃e2 [[e1 same-time e2] ∧ [[x dangerous] ** e2]])]])

Here x and $e1$ are matchable, x being quantified by a positively occurring $\forall$-quantifier, and $e1$ by a negatively occurring $\exists$-quantifier. Moreover, [x ((attr large) carnivore)] and [[x hungry] ** $e1$] occur negatively, so that they can play the roles of Φ_1 and Φ_2 in (III). Suppose now we get input

[Fenris wolf] and
[[Fenris hungry] ** E7]

Suppose also that by using additional facts from the knowledge base, including hierarchical knowledge implicit in a (type) specialist, we can infer

[Fenris ((attr large) carnivore)]

from [Fenris wolf]. Then, substitution ⟨Fenris/x, E7/e⟩ unifies [Fenris ((attr large) carnivore)] and [[Fenris hungry] ** E7] with Φ_1 and Φ_2, with result

$$\top \rightarrow [\top \rightarrow (\exists e2\, [[\text{E7 same-time e2}] \wedge [[\text{Fenris dangerous}] ** e2]])],$$

which simplifies to

$$(\exists e2\, [[\text{E7 same-time e2}] \wedge [[\text{Fenris dangerous}] ** \text{e2}]]).$$

This process amounts to making the inference "Fenris is dangerous at the time of E7 (the episode of his being hungry)."

Example (RI-III) **: "Fenris is gray"**

As an example with m=n=1, consider rule

(5.8) a. Every wolf is either gray or black
b. (∀x [[x wolf] → [[x gray] ∨ [x black]]]).

Here x is a matchable variable since it is quantified by a positively occurring $\forall$-quantifier. Since [x wolf] occurs negatively, it can play the role of Φ_1. Also, [x gray] and [x black] occur positively and so in particular [x black] can play the role of Φ'_1. Suppose now we obtain as input or inferences

[Fenris wolf] and ¬[Fenris black].

Then, substitution ⟨Fenris/x⟩ unifies [Fenris wolf] and [Fenris black] with Φ_1 and Φ'_1 respectively, with result

$$\top \rightarrow [[\text{Fenris gray}] \vee \bot]$$

which simplifies to [Fenris gray]. This process amounts to making the inference "Fenris is gray," given "Every wolf is either gray or black" and "Fenris the wolf is not black."

Also, we can make the following type of inference, with either RI (I) or RI (II):

$$(\forall x{:}[x\ P]\ [x\ Q]), (\forall y{:}[y\ R]\ [y\ P]) \models (\forall z{:}[z\ R]\ [z\ Q]).$$

In practice, RI is implemented roughly as follows. A newly inferred conclusion, corresponding to one of the Ψ_i or $\neg\Psi'_i$, is used to index to the rule R. An initial determination is then made whether the instantiation is likely to succeed and yield a useful result. If the decision is to instantiate, then the attempt to do so is performed by a recursive algorithm applied to R, which actively seeks to find appropriate Ψ_i and $\neg\Psi'_i$ instances in the knowledge base to unify with negatively and positively occurring subformulas of R. Actually, as was indicated in the *Fenris is dangerous* example, the Ψ_i and $\neg\Psi'_i$, (that is, [Fenris ((attr large) carnivore)] in that example), need not even occur explicitly in the knowledge base. They may be inferred by specialists for type taxonomies, temporal relations, or other special classes of relations, or by a limited amount of Prolog-like backchaining.

Probabilistically Instantiating the Rule

Most of human general knowledge (i.e., working knowledge) seems to be in the form of unreliable generalizations. Even the "facts" we start with may be uncertain. If we assume such uncertain rules and facts come from statistical knowledge, they may be best cast in the form of probabilistic conditionals shown in previous chapters. Using probabilistic conditionals, one can express the frequency with which certain consequences follow in certain cases. Then, by the principle of "direct inference" which treats propositions as if their participants were randomly chosen (apart from satisfying the knowledge we have about them), one can get the probability of a consequence proposition. We can use this probability—called epistemic probability—as a certainty on the conclusion obtained. Given a probabilistic conditional with statistical frequency, probabilistic versions of RI and GC derive conclusions with epistemic probabilities. For example, a rule which asserts that "A person who has recently eaten a meal is unlikely to be hungry" could be instantiated with the fact that a certain individual is hungry, leading to the conclusion that he probably has not eaten for some time.

A probabilistic version of RI results when the Ψ_i or $\neg\Psi'_i$ are allowed to have non-unit lower epistemic probabilities and/or R is a generic conditional. The generalization of RI for probabilistic inference is essentially the same as the unit probability version, but with the controlled variables of a generic conditional counting as matchable when the conditional occurs in a positive environment. Thus, for a positively occurring generic conditional, the controlled variables, existential variables in the antecedent, and universal variables in the consequent are all matchable (depending on whether the conditional as a whole occurs positively or negatively). Additional rules of simplification are also needed, including:

$\top^q \rightarrow_p \Phi$ becomes Φ^{pq}

$\bot^q \rightarrow_p \Phi$ becomes $\top$

$\Phi \rightarrow_p \top^q$ becomes $\top$

$\Phi \rightarrow_p \bot^q$ becomes $(\neg\Phi)^{pq}$.

Here, p, q, $(0 \leq p, q \leq 1)$, are probabilities. They denote the lower bound on the statistical probability when they appear as subscript on the conditional, and the lower bound on the epistemic probability when they appear as superscript of a formula. Statistical probabilities appearing as subscript of conditionals are part of the EL object language, but epistemic probabilities are not. They just indicate the lower bound on the degree of belief on the formula in a particular situation. (These rules are justifiable in terms of probability theory for $q = 1$ or 0, but for intermediate q they merely reflect our intuitions at this point.) As an illustration, if the two conditionals in (5.1 a) in the "large carnivore" example above had weights p and q, the previous conclusion would be obtained with lower probability pq.

Let us now consider examples with non-unit probabilities.

Example (RI, probabilistic): "Perhaps Fenris is hungry"

First, let us slightly modify rule (5.1) as follows.

(5.9) a. For every carnivore, if it is not hungry, it probably is not dangerous

b. $(\forall x\, [[x \text{ carnivore}]$
$\rightarrow [(\exists e1\, [(\neg[x \text{ hungry}]) * e1])$
$\rightarrow_{.6,e1} (\neg(\exists e2\, [[e2 \text{ same-time } e1] \wedge [[x \text{ dangerous}] ** e2]]))]])$,

By the semantics of such conditionals, the conditional "iterates" over the possible values of the controlled variable $e1$, and this preempts the existential quantifier. Thus, (5.9 b) is equivalent to

(5.9) c. $(\forall x\, [[x \text{ carnivore}]$
$\rightarrow [[(\neg[x \text{ hungry}]) * e1]$
$\rightarrow_{.6,e1} (\neg(\exists e2\, [[e2 \text{ same-time } e1] \wedge [[x \text{ dangerous}] ** e2]]))]])$.

Here all three variables are matchable: x being quantified by a positively occurring $\forall$-quantifier, $e1$ being generically controlled, and $e2$ being quantified by negatively occurring $\exists$-quantifier. Also, note that [x carnivore], [$e2$ same-time $e1$] and [[x dangerous] ** $e2$], which will play the roles of Φ_1, Φ_2 and Φ_3 respectively, all occur negatively.

Suppose now we have input

> [Fenris wolf]
> [E7 same-time Now6]
> [[Fenris dangerous] ** E7].

As before, using additional facts from the knowledge base, we infer [Fenris carnivore] from [Fenris wolf]. Then, substitution ⟨Fenris/x, E7/$e2$, Now6/$e1$⟩ unifies [Fenris carnivore], [E7 same-time Now6] and [[Fenris dangerous] ** E7] with Φ_1, Φ_2 and Φ_3, with result

$$\top \rightarrow [[(\neg[\text{Fenris hungry}]) * \text{Now6}] \rightarrow^{.6} (\neg[\top \wedge \top])], \quad \text{i.e.,}$$

$$[[(\neg[\text{Fenris hungry}]) * \text{Now6}] \rightarrow^{.6} \bot],$$

which simplifies to $(\neg[(\neg[\text{Fenris hungry}]) * \text{Now6}])^{.6}$, i.e., "It's probably not the case that Fenris is not hungry" or "It's likely that Fenris is hungry." (Note the superscript .6 which is the lower bound on the epistemic probability.) This process amounts to making the inference "Perhaps Fenris is hungry," given "A carnivore may not be dangerous unless it is hungry" and "Fenris is dangerous."

Example (`RI, probabilistic`): "The crack is not likely due to corrosion"

I illustrate probabilistic inference with one further example, drawn from a more realistic domain. Consider rule:

(5.10) a. If an aircraft that is less than 3 years old has a crack, usually the crack is not due to corrosion

b. $(\exists x{:}[[x \text{ aircraft}] \wedge [(\text{age } x \text{ year}) < 3]]$
$(\exists y{:}[y \text{ crack}][y \text{ located-on } x]))$
$\rightarrow_{.8,x,y} (\neg[y \text{ due-to (K corrosion)}])$

Suppose now we have the following fact:

(5.11) a. The two year old aircraft VB12 has a crack

b. $[[\text{VB12 aircraft}] \wedge [(\text{age VB12 year}) = 2] \wedge$
$[\text{C4 crack}] \wedge [\text{C4 located-on VB12}]]$

Then, RI matches these formulas against the antecedent of the conditional, unifying VB12/x and C4/y, and derives[4]

$$(\neg\,[\text{C4 due-to (K corrosion)}])^{.8}.$$

However, if the system is given only

$$[[\text{VB12 aircraft}] \wedge [(\text{age VB12 year}) = 2] \wedge [\text{C4 located-on VB12}]]$$

instead of (5.11), it will derive

$$[(\text{age VB12 year}) < 3] \rightarrow^{.8} (\neg\,[\text{C4 due-to (K corrosion)}]),$$

which says "If VB12 is less than 3 years old, the crack C4 is not likely due to corrosion."

This kind of inference, based on unreliable generalizations, allows evidence for explanations or predictions to be weighted, much as is done in expert systems. Note that this is in contrast with frame-based approaches which rely on default values (e.g., Krypton) [Brachman *et al.*, 1983]. They can distinguish default conclusions and reliable conclusions, but not *degrees of uncertainty*. In many applications, however, it is natural and often important to assess the degree of uncertainty of conclusions.

Before turning to goal-directed inference rules, we should mention λ-conversion and substitution of equals for equals as further deductive rules available in EL. For example, in the following, (c)-formulas are obtained by applying λ-conversion to (b)-formulas.

(5.12) a. The wolf who ate LRRH was Old Father Wolf[5]

b. (The x: [x (λy[[y wolf] $\wedge$ [y eat LRRH]])]
[x = Old-Father-Wolf])

c. (The x: [[x wolf] $\wedge$ [x eat LRRH]]
[x = Old-Father-Wolf])

(5.13) a. Old Father Wolf was big and wicked

b. [Old-Father-Wolf (λz[[z big] $\wedge$ [z wicked]])]

c. [[Old-Father-Wolf big] $\wedge$ [Old-Father-Wolf wicked]]

An application of skolemization $\langle$W1/$x\rangle$ and splitting to (5.12c) gives us

[4] In EPILOG, the number specialist successfully matches [(age VB12 year) = 2] against [(age x year) < 3].

[5] Note that for simplicity I use *Mary, John, LRRH, Old-Father-Wolf*, etc., as constants denoting individuals with those names throughout this thesis. To be accurate, however, it should be translated into $\langle$The λx [x named "Mary"]$\rangle$, $\langle$The λx [x named "Little Red Riding Hood"]$\rangle$, etc.

(5.12) d1. [W1 wolf]
d2. [W1 eat LRRH]
d3. [W1 = Old-Father-Wolf]

Next by applying substitution of equals for equals ⟨W1/Old-Father-Wolf⟩ to formula (5.13c), we get

(5.13) d. [[W1 big] ∧ [W1 wicked]]

In substitution of equals for equals, we assume a control strategy that prohibits repeatedly applying the substitution to the same formula (unless given a new equality).

5.3 GC: Goal Chaining

GC (Goal chaining), which dominates goal-driven inference (in response to questions), is a pair of very general chaining rules. Chaining from rule consequents to antecedents is a special case.

Singly Instantiating the Rule

As with RI, I first show two unit probability cases, with just one "goal" ("goals" play much the same role here as "facts" in RI):

GC: Goal Chaining
— Unit Probability versions —

$$\text{(I)}\quad \frac{R^{+}(\Phi),\ ?G^{+}(\Psi)}{?\neg(R^{+}_{\sigma'}(\neg(G^{+}_{\sigma'}(\top))))} \qquad\qquad \text{(II)}\quad \frac{R^{+}(\Phi),\ ?G^{+}(\Psi)}{?G^{+}_{\sigma'}(\neg(R^{+}_{\sigma'}(\bot)))}$$

Here, the formulas below the line are inferred subgoals (rather than conclusions). R stands for "Rule", and G for "Goal." σ' "antiunifies" Φ, Ψ (i.e., with positive existentials and negative universals in G regarded as matchable). Rule (I) is *sound* if Ψ contains no unmatchable free variables which are bound in R as a whole; rule (II) is *sound* if Φ contains no unmatchable variables which are bound in G as a whole.[6]

Examples (GC-I,II) : "Is there a forest dweller? a gray individual?"

For example, consider the following rule and goal:

[6]In this case "sound" means *falsity*-preserving; i.e., if the given goal is false, so is the inferred subgoal. (Thus, if the subgoal is true, so is the given goal.)

(5.14) a. Every wolf is a meat eater and a forest dweller
b. R: (∀x: [x wolf] [[x meat-eater] ∧ [x forest-dweller]])

(5.15) a. Is there a forest dweller?
b. G: ? (∃y [y forest-dweller])

In R, subformulas [x meat-eater] and [x forest-dweller] occur positively; in G, y is matchable as it is quantified by a positively occurring ∃-quantifier. Thus, via substitution $\langle y/x \rangle$, we can unify [y forest-dweller] with [x forest-dweller], and using GC (II), get the following new goal

? (∃y (¬ [[y wolf] → [[y meat-eater] ∧ ⊥]])), i.e.,
? (∃y (¬ [[y wolf] → ⊥])), i.e.,
? (∃y [y wolf]).

This process amounts to reducing the question "Is there a forest-dweller?" to "Is there a wolf?", using knowledge "A wolf is a meat eater and a forest dweller." Note that we would have got the same result had we used GC (I).

As another example, consider the following rule and goal:

(5.16) A wolf is either gray or black
R: (∃x [x wolf]) → [[x black] ∨ [x gray]]

(5.17) Is there a gray individual?
G: ? (∃y [y gray])

In R, subformulas [x gray] and [x black] occur positively; in G, y is matchable much as before. Thus, via substitution $\langle y/x \rangle$, we can unify [y gray] with [x gray], and, either by (I) or (II), get

? ¬ [(∃y [y wolf]) → [[y black] ∨ ⊥]],

which is equivalent to ? ¬ [(∃y [y wolf]) → [x black]], i.e., ? (∃y [[y wolf] ∧ ¬ [y black]]). This process amounts to reducing the question "Is there a gray thing?" to "Is there a wolf that is not black?", using knowledge "A wolf is either gray or black."

Example (GC-I,II): "Does Pluto have a toy?"

Next, let us consider the following rule and goal:

(5.18) a. Every dog has a bone
b. R: (∀x: [x dog] (∃y: [y bone] [x own y]))

(5.19) a. Does Pluto has a toy?
b. G: ? (∃z: [z toy] [Pluto own z])

Note that subformulas [x own y] in R and [z toy] and [Pluto own z] in G occur positively. Also, in G, z is matchable. Thus, by substitution ⟨Pluto/x, z/y⟩, we can unify [Pluto own z] in G with [x own y] in R, and using GC (I), get the following new subgoal

? ¬ [[Pluto dog] → (∃y: [y bone] (¬ [[y toy] ∧ ⊤]))], i.e.,
? ¬ [[Pluto dog] → (∃y: [y bone] (¬ [y toy]))], i.e.
? ¬ [(¬ [Pluto dog]) ∨ (∃y: [y bone] (¬ [y toy]))], i.e.,
? [Pluto dog] ∧ ¬ (∃y: [y bone] (¬ [y toy]))], i.e.,
? [Pluto dog] ∧ (∀y: [y bone] [y toy]))].

This process amounts to breaking down the goal "Does Pluto has a toy?" into two subgoals "Is Pluto a dog?" and "Is every bone a toy?".

<u>GC *in Combination with* RI</u>

The general version of GC, like the general version of RI, allows arbitrarily many subsidiary knowledge base facts (or presuppositions) to be invoked in the process of chaining from the given goal to a subgoal. This could be expressed schematically as follows.

GC, Combining RI
— Unit Probability Versions —

$$\text{(III)}\quad \frac{R(\Phi_0, \Phi_1, \ldots, \Phi_m, \Phi'_1, \ldots, \Phi'_n), \quad ?G(\Psi_0), \quad \Psi_1, \ldots, \Psi_m, \neg\Psi'_1, \ldots, \neg\Psi'_n}{?\neg(R_{\sigma'}(\neg(G_{\sigma'}(\top), \bot, \ldots, \bot, \top, \ldots, \top)))}$$

$$\text{(IV)}\quad \frac{R(\Phi_0, \Phi_1, \ldots, \Phi_m, \Phi'_1, \ldots, \Phi'_n), \quad ?G(\Psi_0), \quad \Psi_1, \ldots, \Psi_m, \neg\Psi'_1, \ldots, \neg\Psi'_n}{?G_{\sigma'}(\neg(R_{\sigma'}(\bot, \top, \ldots, \top, \bot, \ldots, \bot)))}$$

(III) is derived from GC (I) and RI (III), and (IV) is derived from GC (II) and RI (III). σ' differs from σ in RI in that it treats variables of Ψ_0 with positively occurring ∃-quantifiers or negatively occurring ∀-quantifiers as matchable. These are very general chaining rule, allowing not only chaining from rule consequents to antecedents, but from *any* positively occurring subformula to the rest of $R(\Phi)$ (negated and suitably instantiated).

Also, probabilities are handled much as in RI. (A subgoal Φ^p is interpreted as meaning that if Φ can be proved with probability q, then the original goal is established with probability pq.)

Example (GC-III) : "Does Fenris the wolf live in something?"

We will consider a similar question with the last one (Pluto and a bone). Let us (5.5) again, which is repeated here as (5.20).

(5.20) a. Every wolf lives in a lair
b. R: $(\forall x\colon [\text{x wolf}]\ (\exists y\ [[\text{y lair}] \wedge [\text{x live-in y}]]))$

(5.21) a. Does Fenris the wolf live in something?
b. G: $?\,(\exists z\colon\ [\text{z thing}]\,[\text{Fenris live-in z}])$
F: [Fenris wolf], *taking* [Fenris wolf] *as presupposition*

This is similar to Pluto's bone question, except that it also uses the fact [Fenris wolf] to match part of the rule.

Note that subformulas [x live-in y] in R and [z thing] and [Fenris live-in z] in G occur positively; whereas [x wolf] in R occurs negatively, while [Fenris wolf] in F occurs positively. As before, z in G is matchable. Thus, with substitution $\langle \text{Fenris}/x,\ z/y \rangle$, we can simultaneously unify [Fenris live-in z]$^+$ in G with [x live-in y]$^+$ in R, and [Fenris wolf]$^+$ in F with [x wolf]$^-$ in R, using a combination of GC(I) and RI(I), and get the following new subgoal

$$?\,\neg\,[(\neg\bot) \rightarrow (\exists y\colon [\text{y lair}]\ (\neg\,[[\text{y thing}] \wedge \top]))],\quad \text{i.e.,}$$
$$?\,\neg\,[\top \rightarrow (\exists y\colon [\text{y lair}]\ (\neg\,[\text{y thing}]))],\quad \text{i.e.}$$
$$?\,\neg\,(\exists y\colon [\text{y lair}]\ (\neg\,[\text{y thing}])),\quad \text{i.e.,}$$
$$?\,(\forall y\colon [\text{y lair}]\,[\text{y thing}]))].$$

Thus, we are left with a reduced subgoal, "Is every lair a thing?".

For instance, in the example above, presupposition [Fenris wolf] led to immediate success (i.e., subgoal $(\neg\bot)$, which is $\top$).

The second class of goal-directed methods consists of standard natural deduction rules such as proving a conditional by assuming the antecedent and deriving the consequent; or proving a negative formula by assuming the positive and deriving a contradiction; or proving a universal by proving an "arbitrary instance" of it. Such rules are needed for completeness, since goal chaining cannot prove valid formulas such as $\Phi \rightarrow \Phi$. An interesting future possibility, in the case of proofs involving assumption-making, is to activate input-driven inferencing (primarily, RI) once an assumption has been made, so that its important consequences will be worked out, making it easier to complete the goal-directed proof.

5.4 Summary and Future Work

I have shown that even though EL is very expressive, inference is straightforward with RI and GC. Yet, there are a couple of remaining problems. First, RI should also be generalized so that it can be applied "internally" to a single formula. For instance,

$$(\forall x\ (\exists y\ [[x \text{ P } y] \land \neg [x \text{ P } y]]))$$

should give an immediate contradiction by "internal" RI.

Second, it turns out that RI and GC are not as general as one would like for certain cases that can be thought of as involving multiple resolutions, but intuitively should be single-step deductions. For instance, one may construct examples in which there are three resolution pairs, where the above rules cannot be applied without violating constraints. An example might be derive a contradiction from

$$\begin{aligned} F\text{: } & (\forall w\,[[w \text{ wolf}] \\ & \quad \rightarrow (\exists u\,[[u \text{ forest}] \land \\ & \qquad (\forall v\,[[v \text{ part-of } u] \rightarrow [w \text{ roams-in } v]])])]) \\ R\text{: } & (\forall x\,[[x \text{ forest}] \\ & \quad \rightarrow (\exists y\ [[y \text{ part-of } x] \land \neg [\text{W1 roams-in } y]])]) \end{aligned}$$

(where W1 is (separately) known to be a wolf), or, prove

$$G\text{: } ?(\exists x\,[[x \text{ forest}] \land (\forall y\,[[y \text{ part-of } x] \rightarrow [\text{W1 roams-in } y]])])$$

from

$$\begin{aligned} R\text{: } & (\forall w\,[[w \text{ wolf}] \\ & \quad \rightarrow (\exists u\,[[u \text{ forest}] \land (\forall v\,[[v \text{ part-of } u] \rightarrow [w \text{ roams-in } v]])])]). \end{aligned}$$

The latter is not provable using GC although it is provable using the natural deduction techniques. It seems desirable to develop not only "internally applicable" versions of RI and GC, but also ones where one can proceed step-by-step to "resolve" one pair of literals at a time if necessary, not necessarily substituting one *entire* formula into the other, but possibly just conjunctive parts thereof.

Despite these remaining problems, many interesting inferences are already made. RI (III) and GC (II), together with natural deduction rules, have been implemented in the EPILOG system [Schaeffer *et al.*, 1991], a hybrid reasoning system combining efficient storage and access mechanism, forward and backward chaining, agenda-driven control structure, and multiple "specialists" for taxonomies, temporal reasoning, etc., which carries out efficiently the kinds of inferences described in this chapter. An extended sample run from EPILOG will be shown in Chapter 10, illustrating RI and GC at work.

Part II

FROM ENGLISH TO EPISODIC LOGIC

Introduction to Part II

In Part I of this thesis, I presented the syntax and the semantics of Episodic Logic, EL, and its inference rules. I also pointed out its NL-like *expressiveness* as one of the most important features of EL. The main concern in Part II is how to obtain such expressive logical forms from English surface syntax easily and systematically. As one might imagine, the NL-like expressiveness of EL indeed makes it easy to derive LF-translations from English sentences. The translation from phrase structure to the preliminary, indexical logical form (LF) is accomplished with simple GPSG-like syntactic and semantic rules, while the final nonindexical episodic logical form (ELF) is obtained by applying simple recursive equations called "deindexing" rules to the preliminary LF and a component of context structures called "tense trees."

Through Chapters 6–9, I will discuss the process of deriving ELF from surface structure, with particular emphasis on temporal deindexing. In Chapter 6, I briefly discuss the transduction process from English to preliminary logical form. In Chapter 7, I review some of the well-known works in English tense-aspect interpretation, motivating the tense tree structure to be developed. In Chapter 8, I describe tense trees and the basic tense-aspect deindexing rules. Possible deindexing rules for complex expressions are provided in Chapter 9, together with some pointers for future extension of the mechanism developed.

Chapter 6

Computing Episodic Logical Form

The computation process for obtaining the episodic logical form, ELF, from surface structure may conceptually be viewed as consisting of several stages— from the initial parsing stage to the final stage involving pragmatic and semantic inference. The early stages, especially parsing and to some extent the computation of a superficial indexical LF, are relatively straightforward to specify, thanks to the efforts of our predecessors. Certain stages, e.g., deindexing of tense and aspect, however, have not seen much progress, despite extensive investigation by many researchers both in AI and in the linguistics community.

In this chapter, I will first provide our view of the conceptual stages of the natural language understanding process. I then focus on the stage that computes preliminary logical form from surface structures, illustrating the process with a simple grammar fragment. I will subsequently show a more extensive set of lexical and phrase structure rules, in particular, rules for VPs and adverbials. Finally, I will provide a preview of what is involved in the transduction from preliminary logical from to completely deindexed episodic logical form. The next three chapters will be devoted on this transduction.

6.1 Conceptual View of the NL Understanding Process

Figure 6.1 depicts our current view of the stages of the understanding process, at a theoretical level. The first three stages in this view are fairly conventional, though the details are eclectic, incorporating ideas from GPSG, HPSG, DRT, and research on mapping English into logic, in particular, [Schubert and Pelletier, 1982, 1989]. At the procedural level, these stages are intended to be interleaved, with on-line disambiguation based on syntactic, semantic and pragmatic principles and preferences.

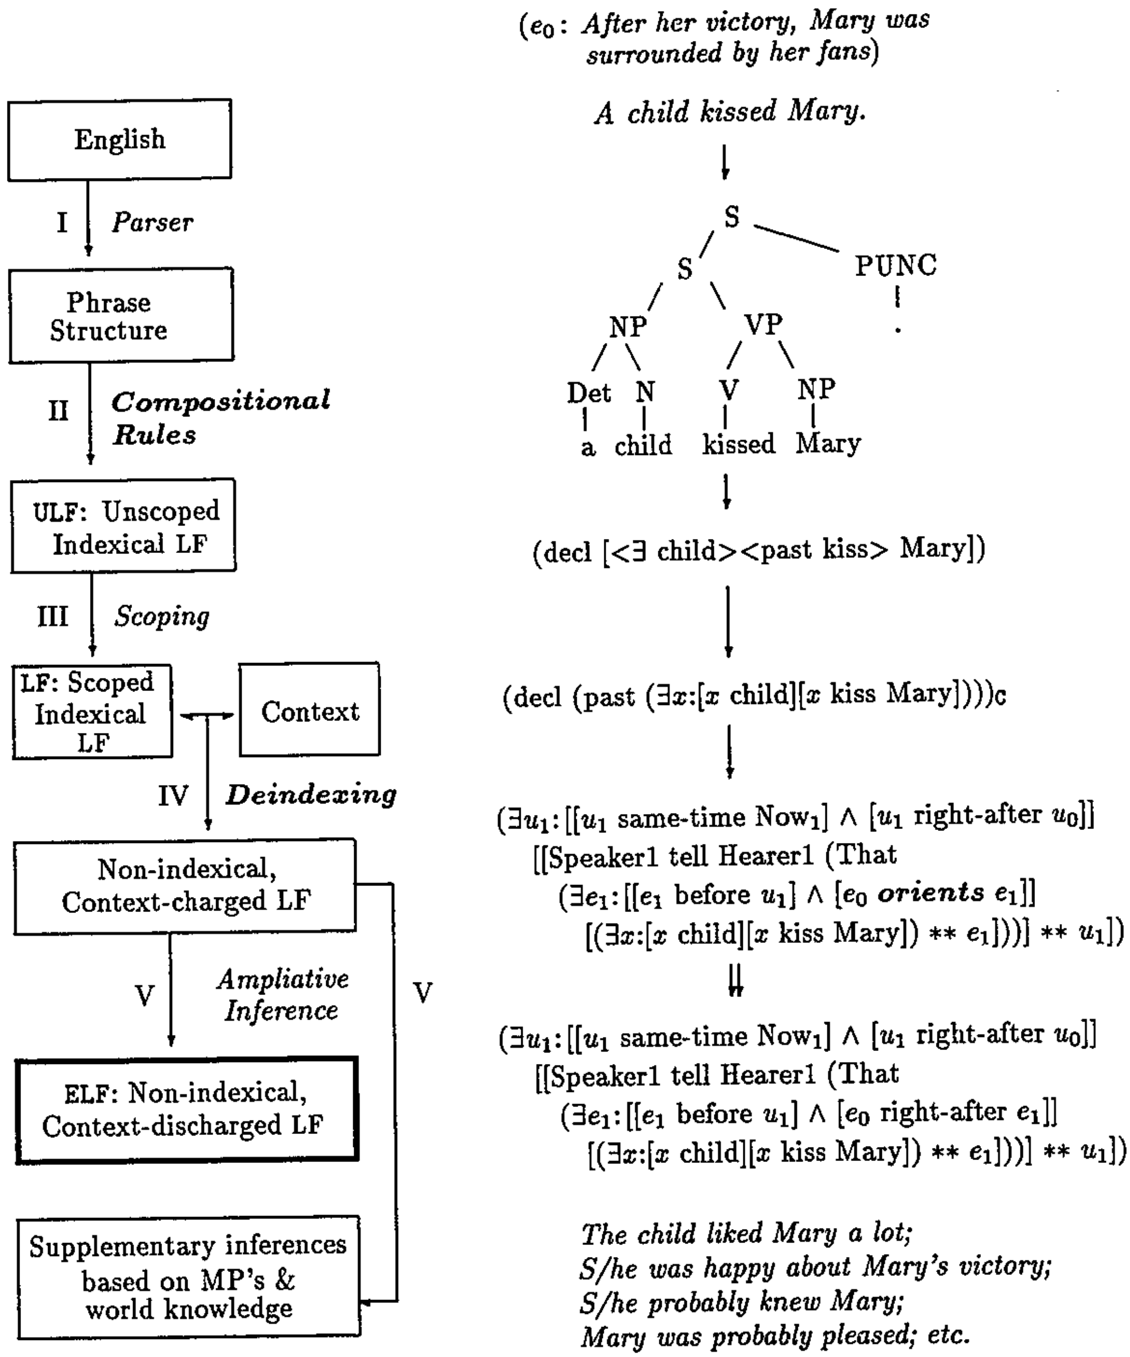

Figure 6.1: The Conceptual Stages of NL Understanding

Let us now consider each of the stages shown in Figure 6.1. Suppose we have the following short passage.

(6.1) After her victory, Mary was surrounded by her fans.

(6.2) A child kissed her {Mary}.

On the RHS of the figure, I illustrate the stages specified on the LHS with a trace of LF-computation for (6.2). In stage I, we obtain parse trees from English, i.e., initial phrase structure translations, using a GPSG-like parser. See the sample parse tree on the RHS of the figure. I will not discuss this stage any further. See [Allen and Schubert, 1991] for some relevant points.

From this initial translation of phrase structure, we get the preliminary, unscoped indexical logical form, ULF, in stage II. This is accomplished with simple semantic rules paired with the phrase structure rules used in stage I. This preliminary ULF is in general ambiguous — e.g., with respect to the scopes of quantifiers and other operators — and context-dependent — e.g., involving indexical operators like *past*, whose interpretation depends on the utterance time. Again see the sample ULF on the RHS. As discussed in Chapter 3, predicate infixing is used for readability. In the logical form, `decl` indicates sentence mood "declarative." This will later be converted to a speech act. Angle brackets indicate unscoped operators that are to be "raised" to some sentence-level position. The sample ULF involves two operators that need be scoped: $\exists$ and `past`. The various processing stages are aimed at removing ambiguity and context-dependence.

Scoping quantifiers in stage III involves an introduction of a variable, i.e., x in this case, and conversion of the restriction predicate to a restriction formula, i.e., *child* becomes $[x$ child$]$. Also scoped at this stage are tense operators and coordinators. `past` and `pres` are considered as sentence-level operators despite syntactic appearances.[1] In general,

[1]Tense and aspect operators are sometimes taken to operate on predicates or even only over verb meanings. One of the arguments in favor of the latter approach may be found in [Enç, 1981]. Enç disputes the validity of the usual argument for sentential scope. This argument is based on apparent ambiguities such as in "All scholarship recipients will attend an award ceremony." If tense is a sentence operator, then it can have narrower or wider scope than *all the scholarship recipients*; in the former case, the NP is evaluated at the speech time (so that the sentence is about *present* scholarship recipients), while in the latter, the NP is evaluated at some future time (so that it may include prospective scholarship recipients). This argument applies to Prior-like tense logics, which were what she was considering. But if nouns have their *own* time reference, this argument falls by the wayside.

Hinrichs adopts Enç's approach and assigns tense scope only over the main verb of the sentence [1988, p. 9]. However, at the level of his proposed logical form, it is not clear whether tense scope is an issue at all, since he replaces indexical operators like PAST or PRES by existentially quantified time variables and their relationships. He takes a sort of Davidsonian approach, introducing a time variable as an "extra" argument of the verb, i.e., he would translate "John kissed every girl" into

$$\forall x[\exists t[\text{girl}(x)(t) \,\&\, \text{R}(x)(t)] \rightarrow \exists t'[t' < ts \,\&\, t' \subseteq tr \,\&\, \text{kiss}(x)(\text{john})(t')]],$$

where R is a context dependent predicate. To me, however, this seems just as readily derivable from a view of tense as sentence operators.

tense has a strong, though not absolute, wide-scoping tendency (right below the sentence mood indicator `decl`); like quantifiers, however, it is "trapped" by scope islands, such as embedded clauses. See the ∃-quantifier and `past` operator in the scoped LF in the example, which is repeated below.

(6.3) (decl (past ($\exists x$: [x child] [x kiss Mary])))$_C$

The subscripted C indicates the explicit context structure with respect to which the scoped LF is to be interpreted. Among other things, it consists of a "tense tree," which serves the purpose of context-dependent tense-aspect interpretation, a "clock" which generates a succession of *Now*-points for speech times, and hearer and speaker parameters. (This will be discussed again in Chapter 8.)

The scoped, indexical translation is to be deindexed with respect to this context C in the next stage. Such a transformation into a nonindexical LF is essential because, to be useful for *inference*, a situational logic must be nonindexical. In stage IV, the computation of the nonindexical ELF from the LF is obtained by a simple, recursive deindexing mechanism that makes use of the context structure C, whose main component for the purpose of this thesis is the tense tree. This handles tense, aspect, and many temporal adverbials and their interaction, and brings the context information into the logical form, removing context dependency. In particular, tense and aspect operators are replaced by relationships among episodes, and explicit *episodic variables* are introduced into the formula on the RHS, which is repeated below.

(6.4) ($\exists u_1$: [[u_1 same-time *Now1*] $\wedge$ [u_1 right-after u_0]]
[[Speaker1 tell Hearer1 (That
($\exists u_1$: [[u_1 before *Now1*] $\wedge$ [e_0 *orients* e_1]]
[($\exists x$: [x child] [x kiss Mary]) ** e_1]))] ** u_1])

Here, u_0 is the utterance episode of the previous sentence, i.e, (6.1), and e_0 is the episode introduced by it, i.e., that of Mary's being surrounded by her fans. *Now1* is the speech time of sentence (6.2). While producing this deindexed formula, the deindexing process also modifies the tense tree component of the context by adding branches and episode tokens as a "side effect."

The `orients` relationship in (6.4) is intended to echo Leech's [1987, p. 41] notion of a point of orientation. This relation is considered as *context-charged*. The idea is that their

In EL, nouns are taken to be untensed "by default" (i.e., only in special cases they are modified by implicit temporal modifiers like `former`, `erstwhile` or `sometime`), and this allows us to account for many cases of ambiguous temporal reference of nominals in the traditional way. Also, Richard and Heny [1982] had an argument that for *untensed* adverbially modified sentences and *tensed* ones to be handled uniformly, tense has to have wide scope over adverbials. As well, the majority on this issue, e.g., Reichenbach, Prior, Dowty, etc. (their work will be discussed in Chapter 7), took tense as sentence operators.

meaning is to be "discharged" using *uncertain* (probabilistic) inference. Though these uncertain inferences will no longer depend on explicit context structure C, the conclusions they lead to will depend on what is "coherent with" the already interpreted prior discourse. In that sense, the conclusions still depend on context. For instance, the fact that e_0 *orients* e_1, i.e., e_0 serves as point of orientation for e_1, "suggests" among other possibilities that e_1 immediately follows e_0 (in e_0's "consequent" or "result" phase, in the terminology of Moens and Steedman [1988]). Given the telic nature of e_0 (Mary being surrounded by her fans) and e_1 (a child kissing Mary) and the circumstances described, this is a very plausible inference, but in other cases the most plausible conclusion from the *orients* relation may be a subepisode relation, an explanatory relation, or any of the discourse relations that have been discussed in the literature (more on this in Chapter 8).

In general, stage IV is also envisaged as performing other kinds of deindexing, most importantly the explicit augmentation of anaphoric expressions with context-derived predications, supplying superficially preferred antecedents. For instance, if the sample sentence were

(6.5) a. <u>The child</u> kissed Mary,
or in LF,
b. <u>(The x: [x child]</u>
$(\exists e_1: [[e_1 \text{ before Now3}] \wedge [e_0 \textit{ orients } e_1]]\ [[x \text{ kiss Mary}] ** e_1]))$,

the underlined anaphoric expression could be augmented in the deindexing stage as shown in (6.6).

(6.6) $(\text{The } x: [[x \text{ child}] \wedge [x \textit{ has-preferred-antecedents } (\text{tuple}\, \tau_1\ \tau_2\ \ldots \tau_m)]]$
$(\exists e_1 : [[e_1 \text{ before Now3}] \wedge [e_0 \textit{ orients } e_1]]\ [[x \text{ kiss Mary}] ** e_1]))$.

Like the earlier `orients` relation, `has-preferred-antecedents` is also a *context-charged* relation, triggering probabilistic inferences whose ultimate conclusion depends on overall coherence. But this will not be pursued further here.

The deindexing stage is followed by an ampliative inference stage (V) which brings to bear MPs and world knowledge, driving a coherence-seeking plausible inference process. This is thought of as leading simultaneously to supplementary inferences about the discourse situation (including the discourse subject matter) and to discharging of context-charged relations in the LF at that point. Thus, (6.6) might become

(6.7) $(\exists u_1 : [[u_1 \text{ same-time } \textit{Now1}] \wedge [u_1 \text{ right-after } u_0]]$
$[[\text{Speaker1 tell Hearer1 (That}$
$(\exists u_1 : [[u_1 \text{ before } \textit{Now1}] \wedge [e_1 \textit{ right-after } e_0]]$
$[(\exists x : [x \text{ child}][x \text{ kiss Mary}]) ** e_1]))] ** u_1])$.

Note that $[e_0$ *orients* $e_1]$ has been particularized into $[e_1$ *right-after* $e_0]$.

At the same time, unique referents for referring expressions, predictions, and explanations are computed which ultimately give a causally coherent elaboration of what has been said. Although I assumed at the beginning that the referent of pronoun "her" in (6.2) had been resolved to Mary, in actuality it would be resolved at this stage. Note the sample inferences indicated on the RHS in Figure 6.1. Finally, it should again be emphasized that though the stages have been described as if they ran sequentially, and the implemented stages are in fact sequenced (implemented partly in the TRAINS system and partly in EPILOG), they are intended to be interleaved eventually. The sequencing is feasible only as long as structural disambiguation, scoping, referent determination and "discharging" of context-charged relations can be adequately "guessed," based only on syntactic preferences and crude semantic checks.

6.2 From English to Preliminary Logical Form ULF

It is one thing to posit a representation, but quite another to actually obtain such a representation from English input. An important advantage of EL is that it can be directly and uniformly computed from syntactic analyses of input sentences. This is possible because episodic logical form is close to English surface form, allowing the logical form to be computed in simple rule-by-rule fashion.

As mentioned, a GPSG-style grammar is used to compute indexical translations with ambiguously scoped quantifiers, connectives and tense operators. In this section, I will focus on the computation of such translations in stage II. We do not have a complete grammar, but the following tentative examples are sufficient to convey the flavor of the grammar-building task. I will first show a step-by-step translation of a very simple sentence, using a small grammar fragment. Then I will add many more grammar rules, mainly for tense, aspect and adverbials. NPs are not treated in any detail. For a substantial NP grammar, the interested reader is referred to [Hwang and Schubert, 1992b].

6.2.1 A Simple Example: Illustration of Translation Process

I will illustrate the derivation of a logical form for the sentence

(6.8) John realized that Pluto was tired.

Table 6.1 is a GPSG fragment adequate for the above sentence, where each lexical or phrase structure rule is paired with a corresponding semantic rule. In the rules, the arrows are in "reverse" direction to indicate that the rules express node admissibility conditions. Note that ID and LP rules have been combined into traditional PSRs for simplicity. Certain feature principles are assumed here—namely, certain versions of the head feature principle, the control agreement principle, and the subcategorization principle (*cf.*, [Gazdar *et al.*,

A.	NP ← *Pluto*; Pluto	
B.	NP ← *John*; John	
C.	A[pred] ← *tired*; tired	
D.	V[be, past, 3per, sing] ← *was*; λP<past P>	
E.	V[_S[that, tensed], past] ← *realized*; <past realize>	
F.	AP ← A[pred]; A′	
G.	VP ← V[be] AP[pred]; (V′ AP′)	
H.	VP ← V[_S[that]] S[that]; (V′ S′)	
I.	S ← NP VP; [NP′ VP′]	
J.	S[that, tensed] ← COMPL[that] S[tensed]; (That S′)	
K.	PUNC[tell] ← . ;	
L.	S[tell] ← S[full-decl] PUNC[tell]; (decl S′)	

Table 6.1: GPSG Fragment I — A simple sentence

1985; Pollard and Sag., 1987]). The subcategorization principle obviates the need for explicit rules like H, but I show the rule for greater clarity.

Rule H is for embedded sentences that are complements of verbs like *believe, infer, prove, know, discover, etc.*, i.e., verbs that take propositions (or possible facts) as their complements. As discussed in Part I of this thesis, propositions are all represented by expressions headed by the nominalization operator `That`.

Figure 6.2 shows the feature system used in Table 6.1. It is a slightly modified version of the one developed by Schubert for the TRAINS project [Allen and Schubert, 1991]. In this system, features are treated as trees rather than as attribute-value functions as in standard GPSG. For instance, feature `tell` in rules K and L is a daughter of feature `utt`. `utt` is the root of the feature tree for "utterance type," or, in view of how punctuation is used to distinguish utterance types, "punctuation type," and has daughters `tell`, `ask`, `instruct`, and `interject` as shown in Figure 6.2. A tree node is considered compatible with any of its ancestors or descendants (the unifier being the "lower," more particular node), and incompatible with the remaining nodes.

The second feature tree is for "sentence mood," with `mood` as root. `mood` has daughters `decl`, `ques`, `imper` and `excl`, and `decl` in turn has daughters `YN-decl` (for sentences "Yes" or "No") and `full-decl` (for declarative sentences other than "Yes" and "No"), and `ques` has `YN-ques` and `wh-ques` as daughters. Feature trees for "imperatives" and

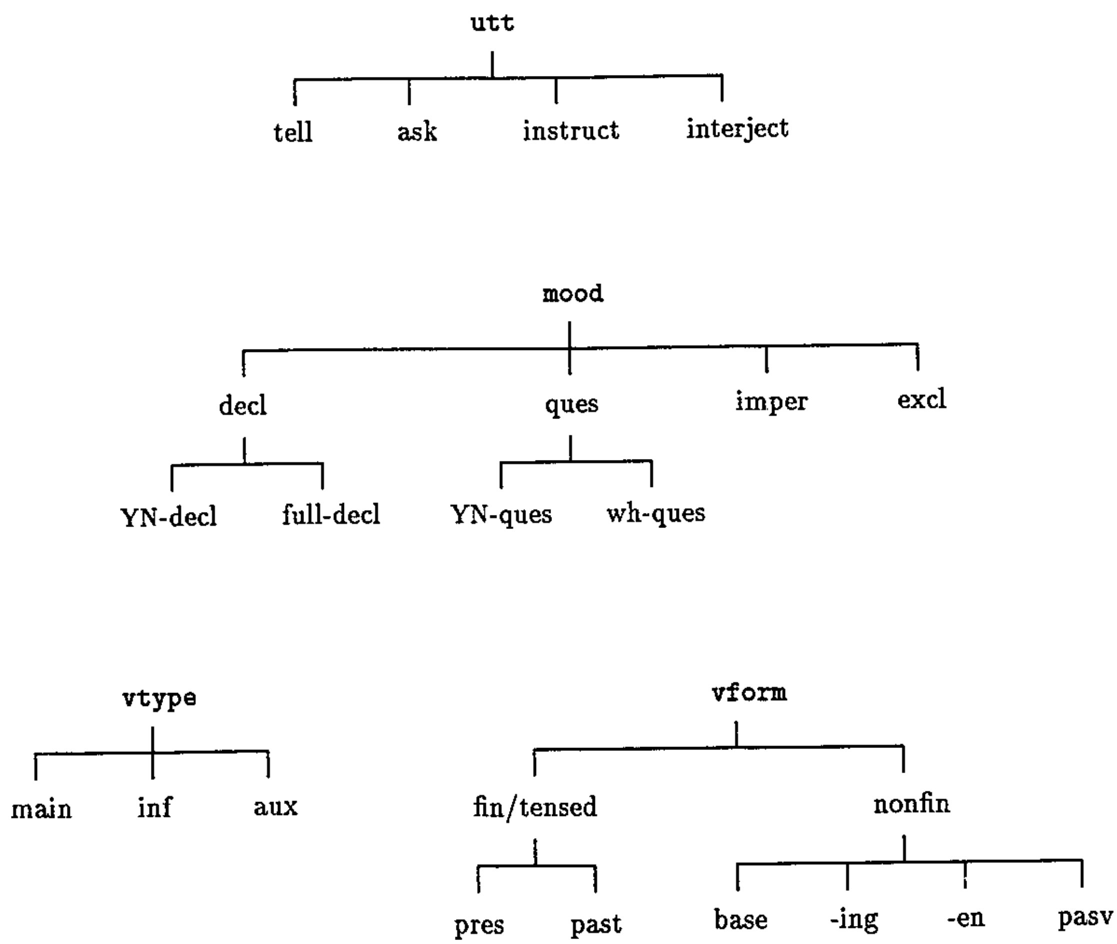

Figure 6.2: `utt`, `mood`, `vtype` and `vform` Feature Hierarchies

"exclamations" have not been fully developed yet. The next feature tree is for "verb types," rooted at `vtype`. `vtype` has daughters `main` (for "main verbs"), `inf` (for the infinitival particle "to"), and `aux` (for "auxiliary verbs"). The feature tree for "verb forms" is headed by `vform`. `vform` has daughters `fin` (for "finite" verb form) and `nonfin` (for "nonfinite" verb form). `fin` is also called `tensed`, and has daughters `pres` and `past`. `nonfin` has daughters `base` ("root" form like *take*), `-ing` (present participles like *taking*), `-en` (past participles like *eaten* or *finished* with *perfect* inflection), and `pasv` (past participles like *eaten* or *finished*, but with *passive* inflection).

I also assume feature trees rooted at `pers` (person) and `numb` (number); the former with daughters `1per`, `2per` and `3per`, the latter with daughters `sing` and `plur`. Also assumed is a feature tree with root `compl` with daughters `that`, `whether`, etc. I omit discussing feature trees for APs, except for noting that `pred` (for "predicative") in rule C in Table 6.1 is from the `atype` feature hierarchy. (The interested reader is refer to [Hwang and Schubert, 1992b].) Feature trees for PPs are discussed later in Section 6.2.3.

Let us now go back to sample sentence (6.8), and trace phrase-formation. To trace it in bottom-up order, I will start with *tired*, then proceed to *was tired*, *Pluto was tired*, etc. I take for granted the upward propagation of certain head features, such as `sing`, `pred`, and `fin` (or `tensed`).

a. AP[*tired*]′ = tired ;
   by applying PS rule F to the translation given by lexical rule C

b. VP[*was tired*]′ = <past tired> ;
   by applying PS rule G to the translation given by D and the result a

c. S[*Pluto was tired*]′ = [Pluto <past tired>] ;
   by applying PS rule I to the translation given by A and the result *b*

d. S[*that Pluto was tired*]′ = (That [Pluto <past tired>]) ;
   by applying PS rule J to the result in *c*

e. VP[*realized that Pluto was tired*]′
   = (<past realize> (That [Pluto <past tired>])) ;
   by applying PS rule H to the translation given by E and the result *d*

f. S[*John realized that Pluto was tired*]′
   = [John (<past realize> (That [Pluto <past tired>]))] ;
   by applying PS rule I to the translation given by B and the result *e*

g. S[*John realized that Pluto was tired.*]′
   = (decl [John (<past realize> (That [Pluto <past tired>]))]) ;
   by applying PS rule L to the result of *f*

Thus, we have the following ULF as a initial translation for sentence (6.8):

(6.9) (decl [John <past realize> (That [Pluto <past tired>])]).

The second phase consists of "raising" the occurrence of past to a permissible sentential level (stage III in Figure 6.1). Following Richard and Heny [1982] and Schubert and Pelletier [1989], I treat tense as sentence operators with wide scope over adverbials (see also Footnote 1). Details of the scoping stage are not covered in this thesis. We now have the following unique result for (6.9), since decl and That act as "scope traps":

(6.10) (decl (past [John realize (That (past [Pluto tired]))])).

This is still indexical in that it is past *relative* to some implicit utterance time, and perhaps also involves implicit relations between the described event and previously described events. Thus, the next step is to combine the indexical translation with a context structure, especially tense trees, and then apply equivalence transformations to the combination, which recursively eliminate the dependence on context, ultimately giving the desired nonindexical (context-independent) translation, as well as introducing episodic variables.[2] This deindexing stage will be discussed extensively in Chapters 8 and 9.

6.2.2 Lexical/PS Rules for VPs

We have seen how a simple English sentence is translated into a preliminary logical form. I now discuss translation of tense and aspect, negation and to-infinitives in some detail.

Table 6.2 shows sample lexical rules for verbs and auxiliaries and PS rules for VPs. Syntactically, perfect and progressive aspects are handled straightforwardly through lexical rules and the auxiliary VP rule shown in the table. In the auxiliary VP rule, XP means any phrase, e.g., NP, AP, VP, PP, etc. "minus"-superscripts indicate optional constituents; where such constituents are not present, the correct semantic rule is obtained by replacing their translations by the identity operator, λPP. The rule is intended to cover all auxiliaries preceding the main verb, as well as copula *be*, i.e., V[be, _XP[pred]] (so that the earlier rule G in Table 6.1 now becomes redundant). Feature constraints such as that if V[aux] is the perfect *have*, then XP must be a VP with feature -en, are assumed to be enforced through subcategorization features on the V, such as V[aux, _VP[-en]]. This is schematically indicated in the rule, with the subcategorization feature _XP. These subcategorization constraints, together with the available forms of auxiliaries, are sufficient to limit auxiliary verb patterns to the usual ones. For example, **would will leave*, **is having left*, and **will do leave*, are ruled out since there is no untensed (base) form of *will* (contrary to the subcategorization requirements of *would*), no progressive (-ing) form of perfect *have* (contrary to the subcategorization requirements of *be*), and no untensed (base) form

[2]In [Schubert and Hwang, 1989], episodic variables and the episodic operators, '*' and '**', were introduced immediately into the logical form by the semantic rules of the GPSG grammar, so that tense, aspect and adverbials can be expressed in terms of relations between episodes. However, this ran into difficulties with the interaction between tense, perfect aspect, negation and time adverbials. The unnatural *t*, *e*, and *h* features in our previous fragment are symptoms of some of these difficulties.

V[aux, pres, 3per, sing, _VP[main, base]] ← *does*; λP<pres P>
V[aux, pres, pers, numb, _VP[base]] ← *will*; λPλx <pres (futr [x P])>†
V[aux, past, pers, numb, _VP[base]] ← *would*; λPλx <past (futr [x P])>
V[aux, pres, 3per, sing, _VP[-en]] ← *has*; λPλx <pres (perf [x P])>
V[aux, base, pers, numb, _VP[-en]] ← *have*; λPλx(perf [x P])
V[aux, past, pers, sing, _VP[-ing]] ← *was*; λPλx <past (prog [x P])>†
V[aux, -en, pers, numb, _VP[-ing]] ← *been*; λPλx(prog [x P])

V[main, pres, 3per, sing] ← *leaves*; <pres leave>
V[main, -en, pers, numb] ← *left*; leave
V[_S[that, base], past] ← *suggested*; <past suggest>
V[main, pres, 3per, sing, _XP[pred]] ← *seems*; <pres seem>
V[main, pres, 3per, sing, _VP[inf]] ← *seems*; λPλx[(That (pres [x P])) <pres seems-true>]
V[inf, _VP[base]] ← *to*; λPP

ADV[neg] ← *not*; ¬

VP ← V[aux, _XP] ADVL[pre-VP]$^-$ ADV[neg]$^-$ XP;
(ADVL′ λx(ADV′ [x (V′ XP′)]))

VP ← ADV[neg]$^-$ V[inf, _VP[base]] VP[base]; λx(ADV′ [x VP′])
VP ← V[_XP] NP; (V′ XP′)
VP ← V[_NP, _NP[ninf]] NP NP[ninf]; (V′ NP′ NP[ninf]′)†
NP[ninf, 3per, sing]‡ ← VP[inf]; (Ka VP′)

PP ← P[for] NP; NP′
S[for-to] ← PP[for] VP[inf]; [PP′ VP′]
NP[for-to-clause, 3per, sing]‡ ← S[for-to]; (Ke S′)
NP[base-that-clause, 3per, sing]‡ ← S[that, base]; (Ke S′)
S[that] ← COMPL[that] S[base]; (Ke S′)

Table 6.2: GPSG Fragment II — VPs

† *Will* and *would* are also used in subjunctives or counter-factuals. Here, we are concerned with futural modality only. prog is an operator indicating "progressive" aspect. ninf is a feature that indicates a n*ominal* derived from an inf*initival* VP.

‡ Such VP-derived and S-derived nominals must not be admitted in all NP positions, for instance, to avoid PPs such as **in to leave* and VPs such as **saw for Mary to leave*. So, strictly, "normal" noun phrases ought to be marked as such. For instance, the lexical rule for *kiss* might be V[base, main, _NP[norm]]. Schubert [1992b] has proposed an ntype feature hierarchy to deal with this.

of auxiliary *do* (contrary to the subcategorization requirements of *will*). (There *is*, of course, a progressive form of *have* as a main verb, and an untensed form of *do* as a main verb.)

I now show some sample sentences and their preliminary translations obtained with the rules in Table 6.2. I show both ULFs (unscoped) and LFs (scoped). In the ULFs, I suppress lambda conversion at several places to help readers get the better grasp of how the translation has been obtained.

(6.11) a. Mary *did not leave*.
b. (decl [Mary λ*x*(¬[*x* ((λP<past P>) leave)])])
c. (decl (past (¬ [Mary leave])))

(6.12) a. John *will realize* that Mary *has left*.
b. (decl [John ((λPλ*x*<pres (futr [*x* P])>)
(realize (That [Mary ((λP′λ*x*′<pres (perf [*x*′ P′])>) leave)])))])
c. (decl (pres (futr [John realize (That (pres (perf [Mary leave])))])))

(6.13) a. John *realized* that Mary *was leaving*.
b. (decl [John <past realize>
(That [Mary ((λPλ*x*<past (prog [*x* P])>) leave)])])
c. (decl (past [John realize (That (past (prog [Mary leave])))]))

(6.14) a. John *thought* that Mary *would not leave*.
b. (decl [John <past think>
(That [Mary λ*x*(¬[*x* ((λPλ*x*′<past (futr [*x*′ P])>) leave)])])])
c. (decl (past [John think (That (past (¬ (futr [Mary leave]))))]))

I show further examples below, this time, focusing mostly on kinds of events and attributes. Note that verbs that take reified sentence intensions, i.e., kinds of episodes/events/situations, as complements are distinguished from verbs that take propositions as complements by their lexical subcategorization. Compare the lexical rule for *suggested* and the PS rule for S[that] in Table 6.2 with rules E and J in Table 6.1. The reason for translating tensed *that*-complement clauses using That, but untensed ones using Ke, is that the former denotes propositions or facts, while the latter denotes a kind of episode/event/situation. For instance, in (6.18) below, what Mary suggests may be an event of the type JOHN LEAVE, *not* a fact of John's leaving. In contrast, in "Mary thinks that John left," what Mary thinks is a proposition, a fact or a non-fact, that reifies a possible past situation in which John leaves.

(6.15) a. Mary likes *to swim*.
b. (decl (pres [Mary like (Ka swim)]))

(6.16) a. *For Mary to have left* is unfortunate.
b. (decl (pres [(Ke (perf [Mary leave])) unfortunate]))

(6.17) a. *For Mary not to be cheerful* is rare.
b. (decl (pres [(Ke ($\neg$[Mary cheerful])) rare]))

(6.18) a. Mary had suggested that *John leave.*
b. (decl (past (perf [Mary suggest (Ke [John leave])])))

(6.19) a. Mary seems *happy*.
b. (decl (pres [Mary (seem happy)]))

(6.20) a. Mary seems *to have inherited Jane's smile.*
b. (decl (The y: [[y smile] $\wedge$ [y of-genitive Jane]]
(pres [(That (pres (perf [Mary inherit y]))) seems-true])))

Note in (6.20 b) how the "subject-intensional" verb *seems* is translated (see the lexical rule for *seems* in Table 6.2). With this kind of translation, sentences like "A unicorn seems to be approaching" is successfully handled in EL. Note also that in (6.20 b), rules for determiners like the following have been used.

Det $\leftarrow$ a{n} ; $\exists$
Det $\leftarrow$ The ; The
Det $\leftarrow$ NP 's; λP<The λx[[x P] $\wedge$ [x of-genitive NP′]]>
NP $\leftarrow$ Det $\overline{\text{N}}$; <Det′ $\overline{\text{N}}'$>

6.2.3 Lexical/PS Rules for Adverbs and PPs

Adverbs and PPs may be syntactically classified into three groups: (i) modifiers (modifying adjectives, adverbs, PPs or NPs), (ii) predicates (i.e., predicative PPs), and (iii) adverbials. I will briefly discuss adverbs and PPs that are used as modifiers or predicates, and then focus on adverbials.

Adverbs/PPs as Modifiers or Predicates

The examples (6.21)–(6.29) below show adverbs or PPs as modifiers, and (6.29)–(6.31) show those used as predicates.

(6.21) Mary is *slightly* obnoxious.

(6.22) John is *very* obnoxious.

(6.23) John runs *very* fast.

(6.24) Mary is *completely* innocent.

(6.25) Jack is *extremely* smart.

(6.26) John is *quite* smart.

(6.27) Mary is *quite* a pianist.

(6.28) The girl *with a purple purse* is Jane.

(6.29) John's comment was {*right*} {*to the point*}.

(6.30) John is *under the table*.

(6.31) Mary is *in love*.

Note that many of the adverbs that modify AP, ADVP, NP, PP, etc., are intensifiers, e.g., *very, extremely, quite,* etc., or downtoners, e.g., *slightly* and *barely*.[3] The lexical and PS rules (with a minimal NP grammar) that are adequate to translate the above sentences are shown in Table 6.3. In the table, `adv-m` (for manner adverbs/adverbials) and `adv-q` (for quality adverbs/adverbials) are functions that map predicates into predicate modifiers. (It may be that they are kinds of action/attribute modifier. That is, (`adv-m` π) may be rewritten as (`adv-a` (in-manner π)), while (`adv-q` π) may be rewritten as (`adv-a` (in-quality π)) or (`adv-a` (in-degree π)), where `adv-a` (for actions and attributes) is a function that uniformly maps predicates over actions and attributes to predicate modifiers.) `adv-m-inv` is the inverse of function `adv-m` (which "strips off the *-ly*," semantically speaking).

I show below translations for some of the sentences (6.21)–(6.31). They are straightforwardly obtained with the rules just discussed (except (6.28′) and (6.31′) that involve PPs, which will be discussed shortly). In (6.28′), 'with-accomp' is a relational predicate indicating accompaniment.

(6.21′) (decl (pres [Mary ((adv-q slight) obnoxious)])), or equivalently
(decl (pres [Mary ((adv-a (in-degree slight)) obnoxious)]))

(6.22′) (decl (pres [John (very obnoxious)]))

(6.23′) (decl (pres [John ((adv-m (very fast)) run)])), or equivalently
(decl (pres [John ((adv-a (in-manner (very fast))) run)]))

(6.25′) (decl (pres [Jack ((adv-q extreme) smart)]))

(6.27′) (decl (pres [Mary (quite pianist)]))

[3] Intensifiers or downtoners may be interchangeably used with degree indicating predicate modifiers, perhaps in the form of (degree *p*), where *p* is a number. For example, on a scale of 0–5, (degree 4) and (degree 5) may be used as intensifiers; (degree 0) and (degree 1), downtoners.

ADV[_AP] ← *very* ; very
ADV[_AP] ← *extremely* ; (adv-q extreme)
ADV[_XP[pred]] ← *quite* ; quite
ADV[_XP[pred]] ← *barely* ; barely
ADV[_XP[pred]] ← *slightly* ; (adv-q slight)

AP ← ADV AP ; (ADV′ AP′)
ADV[manner] ← ADV ADV[manner] ; (adv-m (ADV′ (adv-m-inv ADV[manner]′)))
XP ← ADV[_XP[pred]] XP[pred] ; (ADV′ XP′)

NP[pred] ← Det[a{n}] $\overline{\text{N}}$[count] ; $\overline{\text{N}}'$
NP[pred] ← $\overline{\text{N}}$[plur] ; $\overline{\text{N}}'$
NP[pred] ← NP ; $\lambda x[x = \text{NP}']$
$\overline{\text{N}}$ ← AP[attr] $\overline{\text{N}}$; (AP′ $\overline{\text{N}}'$)

$\overline{\text{N}}$ ← $\overline{\text{N}}$ PP ; $\lambda x[[x\ \text{P}] \wedge [x\ \text{PP}']]$
NP ← $\overline{\text{N}}$[mass] ; (K $\overline{\text{N}}'$)
NP ← $\overline{\text{N}}$[plur] ; (K $\overline{\text{N}}'$)
NP ← Det[a{n}] $\overline{\text{N}}$[count] ; (K1 $\overline{\text{N}}'$)
NP ← $\overline{\text{N}}$[num] ; (K1 $\overline{\text{N}}'$)

Table 6.3: GPSG Fragment III — Adverbs

(6.28′) (decl (The z: [z λx[[x girl] ∧ (∃y: [y ((attr purple) purse)] [x with-accomp y])]]
(pres [z = Jane])))

(6.30′) (decl (The x: [x table] (pres [John under-loc x])))

(6.31′) (decl (pres [Mary in-love]))

Adverb/PP-Adverbials

I should mention first that I confine myself in this thesis to non-clausal adverbials. Syntactically, I take all adverbials as combining with VPs. Semantically, however, they are classed into ones that operate on *sentences* (i.e., formulas) and ones that operate on *predicates*. The examples (6.32)–(6.38) below are adverbials that operate on sentences, and those in (6.39)–(6.46) are ones that operate on predicates.

(6.32) John finished the project *yesterday.*

(6.33) John kissed Mary *behind the curtain.*

(6.34) Mary was *temporarily* out of job.

(6.35) John *regularly* dated Mary.

(6.36) John *repeatedly* called Mary.

(6.37) John called Mary *twice.*

(6.38) *Luckily*, the police arrived in time.

(6.39) Bill has grown up *considerably.*

(6.40) Mary *politely* declined the invitation.

(6.41) John baked the cake *with Mary.*

(6.42) Olga called *from Russia.*

(6.43) John hit the ball *against the wall.*

(6.44) The temperature dropped *below the freezing point.*

(6.45) John walked *along the shore.*

(6.46) Mary has freckles *on her cheeks.*

Adverbials that are sentence operators map sentence intensions to sentence intensions, i.e., functions of type $(\mathcal{S}\rightarrow 2)\rightarrow(\mathcal{S}\rightarrow 2)$. The most frequently encountered ones are episodic adverbials, frequency adverbials, cardinal adverbials, and propositional adverbials. *Episodic adverbials* specify the properties of episodes, especially, their temporal or spatial nature, e.g., *yesterday* in (6.32), *behind the curtain* in (6.33), and *temporarily* in (6.34). *Frequency adverbials* indicate repetition of a certain kind of episode, e.g., *regularly* in (6.35) and *repeatedly* in (6.36). *Cardinal adverbials* like *twice* in (6.37) indicate numeric frequencies. And *propositional adverbials* are mostly attitude operators like *luckily* in (6.38).

Episodic adverbials are translated into (adv-e π), where π is a predicate over episodes, and adv-e is a function that transforms predicates over episodes into sentence modifiers. Recall that in EL, sentence intensions are partial functions from situations/episodes to truth values. Note that the episodic adverbial *temporarily* in (6.34) transforms unbounded-stative sentences into bounded-stative ones. Thus, "Mary was out of job" characterizes an unbounded episode, but "Mary was temporarily out of job" characterizes a bounded one.[4]

[4]Since "temporarily" applies to stative sentences only, sentence "John left the job temporarily" needs to be interpreted as talking about the state resulting from his leaving the job. "John left the job for three months" is interpreted similarly. This kind of shift in aspectual class will be discussed later.

Frequency adverbials like 'frequently', 'regularly', 'repeatedly', etc., are translated into the form (adv-f π), where adv-f maps monadic predicates over sequences[5] into sentence modifiers. Cardinal adverbials like 'twice', 'five times', etc., are translated into (adv-n π), where adv-n maps monadic predicates over collections into sentence modifiers. Note that frequency adverbials transform bounded sentences into unbounded-stative ones, and cardinal adverbials apply to bounded formulas only. Proposition-modifying adverbials are translated into (adv-p π), where adv-p is an operator that maps monadic modal predicates into proposition modifiers, e.g., modal operators. adv-e, adv-f, adv-n and adv-p are functions of type $(\mathcal{S}\to(\mathcal{S}\to 2))\to((\mathcal{S}\to 2)\to(\mathcal{S}\to 2))$, as discussed in Chapter 4.

There are other kinds of sentential level adverbials as well; e.g., *alternatively* in "Alternatively, John can see a movie," or *consecutively* or *alternately* as in "Mary consecutively went shopping, mowed the lawn, and cleaned the house" and "John alternately laughed and cried." Note that *alternatively* is used only as sentence adverbial — sentence premodifying operator like "but," "therefore" or "or" — and it seems to implicitly *connect* the modified proposition with a previous one. This kind of adverbials is often called *conjunctive* ones, and may be translated using adv-p, as follows.

(pres ((adv-p (alternative-to <The thing>)) (can ($\exists x$:[x movie] [John see x]))))

This is still very tentative though.

The latter kind of adverbials (i.e., *consecutively* or *alternately*) may be represented using a meta-operator such as adv-meta, the details of which have not been thought out. Note that the meaning of the VP modified by such adverbials depends on both the meaning of the conjuncts and their *syntactic order* (cf. "the former; the latter"). So, Mary went shopping, mowed, and cleaned, in the order specified by surface structure, i.e., "in the order in which the phrase occur."[6] Similarly, John did laughing and crying repeatedly, but neither laughing nor crying twice in a row. All this requires further thought, and I will not dwell on them in this thesis.

Next is a large class of adverbials that are functions transforming predicates into predicate modifiers of various types. *Considerably* in (6.39) is a quality/degree/quantity modifying adverbial; *politely* in (6.40) is a manner adverbial; *with Mary* in (6.41) is an action/attribute modifying adverbial, indicating "accompaniment," in particular; *from Russia* in (6.42), *against the wall* in (6.43), *below the freezing point* in (6.44), and *along the shore* in (6.45) are path modifying adverbials — "origin," "target," "destination" and "trajectory," respectively. *On her cheeks* in (6.46) may be considered either as an episodic adverbial specifying a spatial location or an attribute modifying adverbial specifying the

[5] *Multi-component episodes*, to be precise. An episode is a multi-component episode if its temporal projection is a multi-interval.

[6] *Consecutively* is not always a meta operator; cf., *consecutively* in "The axioms are numbered consecutively" is a manner (or attribute modifying) adverbial.

"focal" location at which the property is manifested. I tentatively take the latter view. Note that the examples discussed so far are not exhaustive.

As mentioned earlier, adv-m and adv-q transform predicates over manner and qualities/degrees/quantities into predicate modifiers, respectively. Examples are (adv-m polite) for "politely" and (adv-q considerable) for "considerably." Most "-ly" adverbs may be used as quality adverbials, with a few exceptions; cf., in general, propositional adverbials cannot presumably because of the difference in semantic types. Action or attribute modifying adverbials are translated into (adv-a π); e.g., (adv-a (with-accomp Mary)), for "with Mary."

I now show lexical rules for some adverbs. Table 6.4 lists sample lexical entries for

ADV[pre-VP] ← *certainly*; $\lambda P \lambda x$((adv-p certain) [x P])
ADV[mod-VP] ← *slowly*; (adv-a (in-manner slow))

ADV[pre-VP] ← *kindly*; (adv-a kind)
ADV[mod-VP] ← *kindly*; (adv-a (in-manner kind))
ADV[pre-VP] ← *foolishly*; (adv-a foolish)
ADV[mod-VP] ← *foolishly*; (adv-a (in-manner foolish))

ADV[pre-VP] ← *strangely*; $\lambda P \lambda x$((adv-p strange) [x P])
ADV[mod-VP] ← *strangely*; (adv-a (in-manner strange))
ADV[pre-VP] ← *clearly*; $\lambda P \lambda x$((adv-p clear) [x P])
ADV[mod-VP] ← *clearly*; (adv-a (in-manner clear))

ADV[mod-VP] ← *briefly*; $\lambda P \lambda x$((adv-e brief) [x P])
ADV[mod-VP] ← *temporarily*; $\lambda P \lambda x$((adv-e temporary) [x P])

ADV[mod-VP] ← *frequently*; $\lambda P \lambda x$((adv-f frequent) [x P])
ADV[mod-VP] ← *regularly*; $\lambda P \lambda x$((adv-f regular) [x P])
ADV[mod-VP] ← *repeatedly*; $\lambda P \lambda x$((adv-f repetitive) [x P]) [x P])

ADV[mod-VP, cardinal] ← *twice*;
$\lambda P \lambda x$((adv-n ((num 2) (plur episode))) [x P])

Table 6.4: Lexical Rules for Sample Adverbs

adverbs — mostly "-ly" adverbs. As shown in the table, many "-ly" adverbs are ambiguous between proposition modifying adverbials, action/attribute modifying ones, and manner describing ones. The distinction between them is intuitively significant. For instance, "He stupidly grinned at Mary" is ambiguous between "His *action* of grinning at Mary was stupid" (i.e., it was stupid of him to grin at Mary) versus "The *manner* of his grinning at Mary was stupid." Below are some additional examples that involve ambiguous "-ly" adverbs.[7]

Mary *kindly* offered me a ride.	*action* (preferred reading)
Mary talked to the child *kindly*.	*manner*
John *foolishly* did not sign the paper.	*attribute*
John *foolishly* stepped to the front.	*action* (preferred reading)
John talks *foolishly*.	*manner*
Strangely, there was no one in the room.	*propositional*
Mary *strangely* looked at John.	*manner* (preferred reading)

For instance, "Mary *kindly* offered me a ride" above could mean "It was kind of Mary to offer me a ride" (action modifying) or "The way Mary offered me a ride was kind" (manner describing). "John *foolishly* did not sign the paper" has only one interpretation though, i.e., "It was foolish of John not to sign the paper" (cf., "*The way John *performed* the action of *not* signing the paper was foolish" does not make sense). However, as the next couple of sentences illustrate, *foolishly* may be attribute- or manner-modifying as well. Finally, *strangely* may be propositional (the first example above meaning "The fact that there was no one in the room was strange") or manner describing (the second example, "Mary looked at John in a strange way").

Figure 6.3 shows a tentative `ptype` ("preposition type") feature hierarchy. In this hierarchy, the feature `arg` indicates those prepositions that form PP-arguments of verbs (e.g., "Mary is always complaining *about* food"). `ppred` indicates ones that form PP-predicates. Such predicates could be either episode modifying, i.e., `e-mod`, or action/attribute modifying, i.e., `a-mod`. `ppropos` indicates ones that head propositional adverbials, i.e., those with the feature `p-mod`. As indicated in the figure, `ptype` features have these features by a co-occurrence restriction.

I should mention that the `ptype` feature hierarchy has been developed for pragmatic purposes, and is not necessarily an optimal one. This hierarchy provides us with translations of prepositions that are convenient for meaning postulates to be applied to. For instance, preposition 'at' is translated into *at-time*, *at-loc*, etc., depending on its argument (e.g., whether it is a temporal nominal or a locative one), which then allows meaning postulates about time, location, etc., to be applied straightforwardly. However, making such distinctions in translation based on syntax only may not be a very general approach; for instance, metaphors may not be handled. A better method would translate preposition 'at' uniformly as *at*, and later disambiguate it based on semantic grounds.

The grammar for adverbials is shown in Table 6.5. But I need first to reiterate that adverbials are taken to be uniformly *VP-adverbials* at the level of syntax (both are VP adjuncts). Initial adverbials in sentences like "Yesterday John left" are treated as topi-

[7]These examples are due to Len Schubert and Phil Harrison (personal communication).

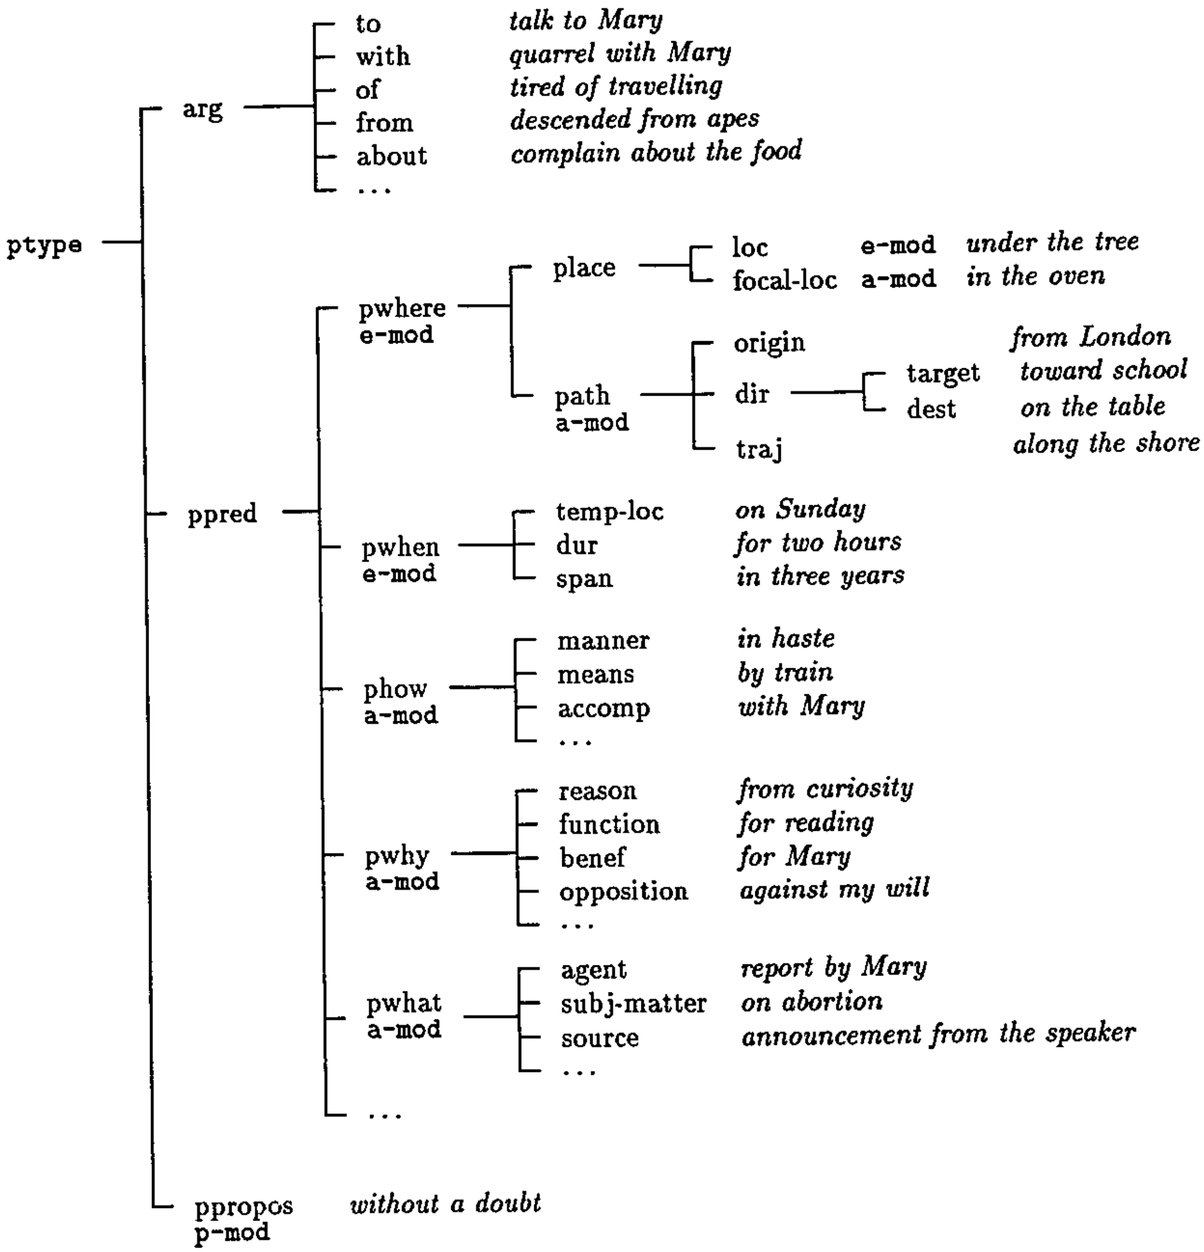

Figure 6.3: The ptype Feature Hierarchy

P[to] ← *to*; λPP
V[_P[to]] ← *talk*; talk-to
NP[time] ← *yesterday*; *yesterday*
ADVL[post-VP] ← NP[time]; λPλx((adv-e (during NP′))[x P])
P[loc] ← *behind*; behind-loc
P[dur] ← *for*; lasts
P[benef] ← *for*; for-benef
PP ← P NP; (P′ NP′)
ADVL[post-VP] ← PP[e-mod]; λPλx((adv-e PP′) [x P])
ADVL[post-VP] ← PP[a-mod]; (adv-a PP′)

VP ← ADVL[pre-VP] VP; (ADVL′ VP′)
VP ← VP ADVL[post-VP]; (ADVL′ VP′)

Table 6.5: GPSG Fragment IV — Adverbials

calized. Semantically, i.e., at the level of LF, however, temporal adverbials and locative adverbials are sentence modifiers while action modifying adverbials (including manner and quality adverbials) are predicate modifiers as mentioned earlier. The grammar fragment in Table 6.5 handles various adverbs and NP/PP-adverbials. Some relevant lexical rules are also included. Note that adverbials are divided into pre-VP and post-VP adverbials, and those that can be either, with feature `mod-VP`. (A feature hierarchy with `mod-vp` as root, and `pre-vp` and `post-vp` as daughters, is assumed.) With these rules, sentences (6.32)–(6.46) are translated in a straightforward way. I show some of their translations below. Because predicates are interpreted formally as "curried" functions (i.e., applicable to one argument at a time) and *during, before, below-loc, in-loc,* etc., are 2-place episode predicates, (during *Yesterday*), (below-loc x), etc., below are monadic predicates (with an indexical constant *Yesterday* and a variable x).

(6.32′) (decl (past (The x: [x project] ((adv-e (during *Yesterday*)) [John finish x]))))

(6.33′) (decl (past (The x: [x curtain] ((adv-e (behind-loc x)) [John kiss Mary]))))

(6.34′) (decl (past ((adv-e temporary) [Mary (out-of job)])))

(6.35′) (decl (past ((adv-f regular) [John date Mary])))

(6.37′) (decl (past ((adv-n ((num 2) (plur episode))) [John call Mary])))

(6.38′) (decl ((adv-p lucky)
(past (The x: [x (plur police)] [x ((adv-a in-time) (plur arrive))]))))

(6.39′) (decl (pres (perf [Bill ((adv-q considerable) grow-up)])))

(6.40′) (decl (past (The x: [x invitation]
[Mary ((adv-a (in-manner polite)) (decline x))])))

(6.41′) (decl (past (The x: [x cake] [John ((adv-a (with-accomp Mary)) (bake x))])))

(6.45′) (decl (past (The x: [x shore] [John ((adv-a (along-path x)) walk)])))

(6.46′) (decl (pres (The x: [[x (plur cheek)] $\wedge$ [x of-genitive Mary]]
($\exists y$: [y (plur freckle)] [Mary ((adv-a (on-focus x)) (have y))]))))

Note that some of the above translations are still tentative. As mentioned, it is not completely clear how to translate sentence (6.46).[8] "Mary's having freckles" is true on her cheeks, but not necessarily elsewhere. So, "on her cheeks" might be considered as property of an episode. Similar examples are "The lawn has lots of dandelions on it near the bank," "The lawn is greener near the bank," etc. We may eventually want to exploit the location-dependence of sentence truth for these examples.

Meaning postulates later apply the predicates which are the arguments of these operators to episodes and actions, respectively. This will be discussed in Chapter 9.

6.2.4 Aspectual Class Shifts

Before closing this section, I need to mention shifts in aspectual classes. Figure 6.4 shows the tentative aspectual class feature hierarchies, `stativeness` and `boundedness` (similar to Passonneau's [1988] system — see Section 7.2.3).[9]

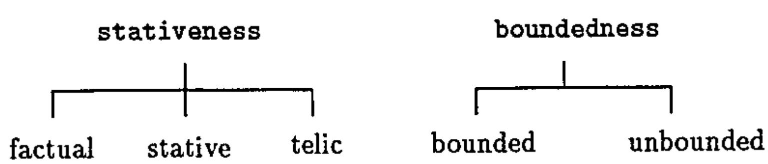

Figure 6.4: `stativeness` & `boundedness` Feature Hierarchies

[8]The translation (6.46′) is due to Phil Harrison (personal communication).

[9]We also need to consider *polarity* (see Footnote 16 in Chapter 2) in combination with aspectual classes, but I will neglect it here.

factual unbounded	stative unbounded	stative bounded	telic bounded
2 and 2 make 4	*Mary cheerful*	*Mary briefly happy*	*Mary blink*
WW2 after WW1	*Mary write letters*	*John (bounded sick)*	*Mary write a book*
$x = x$	*John sick*	*John (bounded sick) twice*	*John finish project*
Mary left before John arrived$_C$	*John frequently get sick*	*Bill (bounded asleep)*	*John go to opera*
	Bill sleep for an hour	*Bill (bounded sleep for an hour)*	*Bill get sick*
			Bill get sick twice

Figure 6.5: Aspectual Categories and Sample "Sentences"

Figure 6.5 shows sample "sentences" for each aspectual category. Recall that every tensed English sentence, e.g., "Mary left before John arrived," in combination with a context, is factual (i.e., unlocated). For untensed sentences, their stativeness depends on predicates (i.e., achievement/accomplishment versus state/process predicates), objects (e.g., a single countable object versus a mass object), and subjects (e.g., a non-collective individual versus a collection). By a co-occurrence restriction, factual formulas are unbounded, and telics are bounded. Statives are by default unbounded; e.g., JOHN SICK or JOHN HAVE A HEADACHE are considered unbounded. But they may become bounded and combine with frequency or cardinal adverbials.

Slightly more formally put, a formula is *bounded* if the episode it characterizes terminates in a distinctive result state. This is a property we ascribe to all telic episodes (such as accomplishments and achievements), as well as to some stative episodes (such as an episode of John's being asleep, at the end of which he is *not* asleep). Such episodes are called *closed* (or *inextensible*) episodes. Conversely, an episode is *open* (or *open-ended* or *extensible*) if it does not terminate in a distinctive result state. This is a property ascribed to *unbounded* states and processes, i.e., ones whose final point are described by the same description which characterizes the given episode as a whole. For instance, *was sick* in "John was sick when I saw him last week" is unbounded as the sentence does not imply that John was not sick right after the described episode. However, when we say "John was sick twice last year," we are talking about bounded "sick" episodes.

As a tentative formalization, we define a function `sit-type` ("situation type") by

$$(\text{sit-type } \eta) = (\text{Ke } \Phi) \quad \text{iff} \quad [\Phi ** \eta],$$

and introduce a function `result-type` from situation types to situation types. This will be constrained by axioms such as

$$(\text{result-type } (\text{Ke } [\tau \text{ wake-up}])) = (\text{Ke } [\tau \text{ awake}]),$$
$$(\text{result-type } (\text{Ke } [\tau \text{ asleep}])) = (\text{Ke } [\tau \text{ asleep}]), \quad \text{and}$$
$$(\text{result-type } (\text{Ke } [\tau \text{ (bounded asleep)}])) = (\text{Ke } (\neg [\tau \text{ asleep}])).$$

In keeping with its name, `result-type` satisfies the schema

$$[[\Phi ** \eta] \wedge [(\text{result-type } (\text{Ke } \Phi)) = (\text{Ke } \Psi)]] \rightarrow [\Psi * (\text{end-of } \eta)].$$

Then we can say η *terminates in a distinctive result state* iff

$$(\text{result-type } (\text{sit-type } \eta)) \neq (\text{sit-type } \eta).$$

The role of aspectual classes in the interaction between VPs and various kinds of temporal adverbials is crucial for interpretation of sentences and has been extensively discussed in the literature. For example, durative *for*-adverbials may combine with stative VPs only. When such stative adverbials are applied to telic VPs, iteration is implied. In contrast, frequency adverbials may combine with bounded VPs only; when they are applied to unbounded-stative VPs, those VPs need to be interpreted as bounded-statives. Cardinal adverbials also apply to bounded episodes only. Time adverbials specifying temporal locations, like *yesterday* or *last week*, may combine with either bounded or unbounded formulas (with unbounded ones, it suggests a *throughout* reading; with bounded ones, a *sometime during* reading). Thus, the adverbial rules shown earlier in this section need to be made more specific to accommodate the interaction between VPs and adverbials and possible shifts in aspectual classes, in the following direction.[10]

VP ← VP[stative] ADVL[dur]; (ADVL′ VP′)
e.g., Mary has *lived in London* <u>*for two years*</u>
VP ← VP[bounded] ADVL[cardinal]; (ADVL′ VP′)
e.g., Mary *visited Paris* <u>*twice*</u>
VP[unbounded, stat] ← VP[bounded] ADVL[freq]; (ADVL′ VP′)
e.g., John <u>*often*</u> *gets depressed*
VP[bounded] ← VP[unbounded, stative]; (bounded VP′)
e.g., John was *sick* twice ⇝ (<u>bounded</u> *sick*)
VP[unbounded, stative] ← VP[telic]; (iter VP′)
e.g., John *dated Mary* for two years ⇝ (<u>iter</u> *date Mary*)
VP[unbounded, stative] ← VP[telic]; (result-state VP′)
e.g., Mary temporarily *left the job* ⇝ (<u>result-state</u> *leave the job*)

[10]This kind of shifts in aspectual classes have already been discussed in literature; first in [Steedman, 1982], and subsequently in [Moens and Steedman, 1988; Smith, 1991].

Note that VPs headed by cardinal adverbials are considered bounded because we can say,

John {was sick twice} {in three years},

but

*John {was sick twice} {for three years}.

But VPs headed by frequency adverbials are considered unbounded; i.e.,

John {was frequently sick} {for three years};
*John {was frequently sick} {in three years}.

I now show some additional examples that illustrate these rules.

(6.47) a. Mary was sick for three weeks
b. (past ((adv-e (lasts (K1 ((num 3) (plur week))))) [Mary sick]))

(6.48) a. Mary was sick twice in December
b. (past ((adv-e (during (in-time *December*)))
((adv-n ((num 2) (plur episode))) [Mary (bounded sick)])))

(6.49) a. Mary received the (best dresser) award for three years
b. (past (The x: [x award] ((adv-e (in-span-of (K1 ((num 3) (plur year)))))
[Mary (iter (receive x))])))

(6.50) a. Mary became unconscious for five minutes
b. (past ((adv-e (lasts (K1 ((num 5) (plur minute)))))
[Mary (result-state (become unconscious))]))

(6.51) a. John wrote the book in three years[11]
b. (past (The x: [x book] ((adv-e (in-span-of (K1 ((num 3) (plur year)))))
[John (write x)]))

[11] Sentences like "John finished the book in three years" require somewhat distinct treatment from (6.51). In such sentences, it is its preparatory process (in the sense of Moens and Steedman [1988]) that occupied the three year-long interval, rather than the actual "finishing" event. That is, the sentence needs to be interpreted as "There was a three year-long interval such that over which John was writing the book and at the end of which the completion of writing occurred." Getting such an interpretation requires dealing with pragmatics and is a nontrivial task, but it seems unavoidable. As another example, "It took John two years to propose to Mary" needs to be understood as "At the end of a two year-long interval during which John had a relationship with Mary, he eventually proposed to her," rather than "John was engaged in a 'proposing' event over the two year-long interval."

6.3 From Preliminary Logical Form ULF to Episodic Logical Form ELF

Once an indexical logical form is obtained, the next phase is to combine it with a context structure for the utterance and "deindex" it. The main interest in this thesis is in tense-aspect deindexing. (Other aspects of deindexing, such as anaphoric processing, have not been worked out in detail.)

The temporal relationships between episodes are in general mediated by tense, aspectual auxiliaries, adverbials, surface order, and the context. The goal is a comprehensive account of how such temporal relationships are determined by syntax, semantics, and pragmatics. This will be discussed extensively in Chapters 7–9. Here, I will just indicate the flavor by showing the translation of one of the sentences we looked at in this chapter.

(6.12′) a. John will realize that Mary has left.

b. (decl (pres (futr [John realize (That (pres (perf [Mary leave]))))]))

c. $(\exists e_1\!: [e_1 \text{ same-time } \mathit{Now1}]$
$[[\texttt{Speaker} \text{ tell } \texttt{Hearer} \text{ (That}$
$(\exists e_2\!: [e_2 \text{ at-about } e_1]$
$[(\exists e_3\!: [e_3 \text{ after } e_2]$
$[[\text{John realize (That}$
$(\exists e_4\!: [e_4 \text{ at-about } e_3]$
$[(\exists e_5\!: [e_5 \text{ before } e_4] \; [[\text{Mary leave}] ** e_5]) ** e_4]))]$
$** e_3])$
$** e_2]))]$
$** e_1])$

6.4 Summary

I discussed computation of LFs from English phrase structures, with particular emphasis on VPs and PP adjuncts. Although the grammar fragment I showed here is tentative and far from being complete, lacking rules for questions, imperatives, exclamations, clausal adverbials, and many details of NPs and other phrases, it already handles a wide variety of constructs. Also, I should emphasize that future revisions and expansions of the grammar will not jeopardize the overall architecture proposed in this thesis, because of its modularity. In other words, both the *mechanism* for mapping surface form into meaning representations, and the syntactic, semantic and inferential properties of that representation, are invariant under changes in the grammar and semantic rules. The next three chapters provide the background and details of the deindexing stage that transforms indexical LFs into nonindexical ELFs.

Chapter 7

Locating Episodes in Time: Forerunners of Tense Trees

In Chapter 6, I discussed how to compute preliminary, indexical LFs for English sentences. Before actually getting into deindexing of LFs, however, I am going to digress a bit in this chapter to motivate the deindexing techniques to be developed in the next chapter. We are particularly interested in temporal deindexing here. I first discuss the phenomena we want to address, and then review previous work in this area to see what has been accomplished, what needs to be done, and how it might be done. The next two chapters describe a deindexing mechanism and a new type of context component called tense trees, as a basis for interpreting tense-aspect constructions and some time adverbials in a context-dependent way.

7.1 The Phenomena: Temporal Structure of Discourse

Narratives describe and relate episodes (states, events, eventualities, etc.). Those episodes successively introduced appear to "line up" in systematic ways as a function of the narrative/discourse structure, with tense and aspect playing a crucial role. Examples (7.1)–(7.4) in the following illustrate this familiar phenomenon.[1] In the diagrams below, u_a, u_b, $\ldots$, are utterance episodes for sentences a, b, $\ldots$, respectively; the direction of arrows indicates the progress of time; and the ordering of episode tokens at each node corresponds roughly to their temporal ordering such that the rightmost one is the most recent and ones piled up at the same location are about the same time with each other.

[1](7.1) is from [Hemingway, 1952]; (7.2) is from [Allen, 1987]; (7.3) is due to Len Schubert.

(7.1) a. They walked$_{\{e_1\}}$ up the road together to the old man's shack and
b. went$_{\{e_2\}}$ in through its open door.
c. The old man leaned$_{\{e_3\}}$ the mast against the wall and
d. the boy put$_{\{e_4\}}$ the box and the other gear beside it.
e. The mast was$_{\{e_5\}}$ nearly as long as the one room of the shack.
f. The shack was$_{\{e_6\}}$ made of the tough bud-shields of the royal palm and
g. in it there was$_{\{e_7\}}$ a bed, a table, one chair, ...

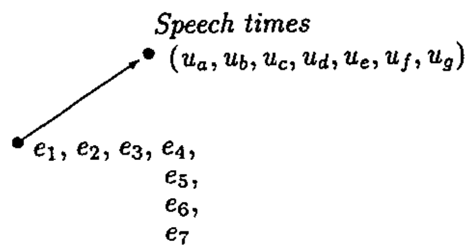

(7.2) a. Jack and Sue went$_{\{e_1\}}$ to a hardware store
b. as someone had$_{\{e_2\}}$ stolen$_{\{e_3\}}$ their lawnmower.
c. Sue had$_{\{e_4\}}$ seen$_{\{e_5\}}$ a man take it and
d. had$_{\{e_6\}}$ chased$_{\{e_7\}}$ him down the street, but
e. he had$_{\{e_8\}}$ driven$_{\{e_9\}}$ away in a truck.

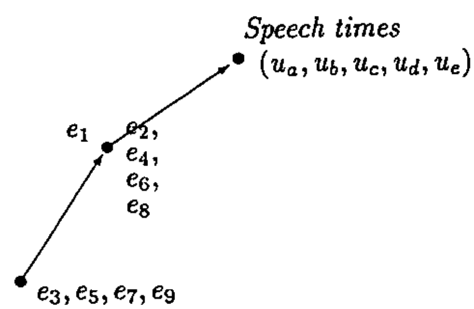

(7.3) a. Mary's note made up$_{\{e_1\}}$ John's mind.
b. He would$_{\{e_2\}}$ leave$_{\{e_3\}}$, and
c. would$_{\{e_4\}}$ let$_{\{e_5\}}$ her live her own life.
d. He began$_{\{e_6\}}$ to pack.

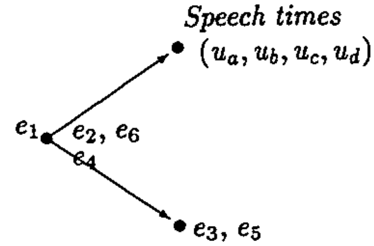

(7.4) a. John will find$_{\{e_1\}}$ this note when he gets home.

b. He will think$_{\{e_2\}}$ Mary has$_{\{e_3\}}$ left$_{\{e_4\}}$.

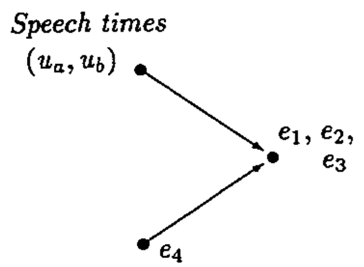

In (7.1), the past inflections and surface ordering place each of episodes e_1–e_7 "before" its respective utterance episode u_a, u_b, $\ldots$, u_g, and "after" (or "at about the same time" as) its immediately preceding episode (thus, each past-episode is "oriented" by its immediate predecessor).[2] By default, in a sequence of past episodes, time progresses, especially when the reported episodes are not stative, e.g., in (7.1), *walk up*$(e_1)\rightarrow$ *go in through door*$(e_2)\rightarrow$ *lean mast*$(e_3)\rightarrow$ *put things*(e_4). For stative episodes, usually time "stands still"; e.g., consider e_5, e_6, e_7 in (7.1).

In (7.2), the occurrence of *had* in (b) indicates the reference episode e_2 is the same, or at the same time, as the antecedent *going*-episode e_1, and places the *stealing*-episode e_3 "prior" to the reference episode e_2 (and similarly, places e_4, e_6, e_8 at the same time as e_1, and e_5, e_7, e_9 prior to e_4, e_6, e_8, respectively). Furthermore, (7.2) shows that the "orienting" relation between narrative episodes applies even to episodes reported in past perfect, i.e., among episodes e_3, e_5, e_7, e_9.

(7.3) and (7.4) show future in the past and present in the future, respectively. In (7.3), e_3 and e_5, episodes of John's leaving and letting Mary live her own life, are "after" their reference episodes e_2 and e_4 respectively, which are at about the same time as e_1, the past episode of Mary's note making up John's mind. (7.4) shows that in John's thinking Mary's leaving took place some time before the time of his thinking, rather than before the speech time.

These examples illustrate two things. First, in contrast to tenseless sentences of traditional logics, natural language sentences implicitly introduce episodes and various relationships among these episodes. The simplest and firmest is the temporal relationship between the episodes described by the sentence and the time the description is made, i.e., the relationship between the event or episode time and the speech time. Thus, the same sentence, say, "John got married last year," may have a different truth value depending

[2] We use the term "orientation" in a way used by Leech [1987] and Webber [1987a], not Smith [1978] for whom "orients" seems to mean simply "stands in some relation to." For example, in her tense interpretation system, the reference time is always oriented to the speech time.

on when it is uttered.

Second, they indicate the anaphoric nature of tense and aspect. That is, episodes successively described by narratives (or other types of discourse) are related to one another in some fashion. For instance, as (7.1) illustrates, an episode described in the simple past is normally linked to a "point of orientation" such as an immediately preceding event. That is, e_2 is after e_1, e_3 is after e_2, ..., e_6 is at the time of e_5, e_7 is at the time of e_6, etc. This phenomenon has been noted by Reichenbach [1947], Partee [1973; 1984], McCawley [1981], Hinrichs [1986], Leech [1987], Webber [1987a; 1988], and many others. Past, present and future perfect are usually held up as particularly clear examples of tenses involving anaphoric relatives to presupposed reference times. As seen in (7.2), the perfect aspect always involves a reference point and the newly described episode is either before or lasts until that reference point. In particular, each of e_3, e_5, e_7, and e_9 is before the reference point e_1. (Note that episodes are introduced not only by main verbs, but also by perfect and futural modal auxiliaries.) Similarly, in (7.3), (b) and (c) involve the future in the past, and it is important to locate the reference points of these future events. That is, a correct analysis should be able to identify that e_2 and e_4, the reference points for events e_3 and e_5 respectively, are in fact the same as e_1. Such temporal relations implicit in a narrative need to be made explicit to accurately represent the meaning of texts. I will begin my discussion of how this can be done by reviewing some previous work on tense and aspect, with particular emphasis on how the above two issues—the role of speech time and the "anaphoric" connections between events—are treated.

7.2 Tense-Aspect Interpretation: A Review of Previous Work

The goal in analyzing tense and aspect is to make explicit the truth conditions of English sentences involving tense and/or aspect. Despite numerous investigations into English tense-aspect semantics, the problem of formally interpreting tense and aspect has remained in large part unsolved. A formal solution requires (i) a well-defined *mapping* from a subset of English covering the most common tense-aspect constructions to a formal meaning representation (corresponding to ① in Figure 1.1) and (ii) a well-defined *denotational semantics* for the meaning representation, which accords with speakers' intuitions about the original text (corresponding to ② in Figure 1.1). With this in mind, we will look at some of the best-known approaches.

Most of the work in this field starts from theories developed by Reichenbach [1947] or Prior [1967]. Reichenbachian approaches transform tense into a set of three indices and assert that the indices stand in a certain relation. The result consists of "detensed" formulas in ordinary logic, including the explicit relationships among the time indices. Priorean approaches use tense operators that shift the time point at which formulas are evaluated. Unfortunately, neither approach as it stands is general enough to handle common English constructs (see [Kuhn, 1989] for a survey). Moreover, researchers pursuing these approaches paid little attention to the translation from English to logical form, i.e., mapping of ① above. Eventually these ideas were combined by Dowty [1979; 1982], who started out with a Prior-like method, and subsequently incorporated some of Reichenbach's concepts into his system. Also, staying within the framework of Montague Grammar, he showed how logical forms can be obtained. Tense-aspect has also been studied from a more computational perspective by AI researchers, e.g., Passonneau [1987; 1988], Webber [1987a; 1988] and Song and Cohen [Song, 1991; Song and Cohen, 1991]. Here, the emphasis has been on effective algorithms for extracting Reichenbach-like temporal relations from text, and temporally connecting successive sentences. As in the logical approaches, there has been relatively little concern with the mapping ①, and with integrating tense-aspect theories with detailed theories of logical form (e.g., see the articles in *Computational Linguistics 14*(2), 1988).

I now discuss these works and summarize what has been accomplished and what problems remain.

7.2.1 Reichenbach's Legacy

Reichenbach's [1947] major contribution was his recognizing the importance of what he called *reference time* in analyzing tense-aspect. Grammarians before him had taken for granted that tense would be successfully analyzed in terms of just the *speech time* and the *event time*. However, referring to sentences like "Peter had mailed the letter," he pointed out the need for a third index, a *reference time*, so that the given sentence can be analyzed as "There is some *reference time* which is before the speech time, and the event of Peter's mailing the letter took place before this *reference time*."

Recognizing the essential role of reference time in perfects, Reichenbach decided that this reference time is always present, even in simple past, present and future sentences. (The reference time R is what Reichenbach calls the temporal perspective from which an event is viewed.) This led him to his conception of tenses (both simple and complex) as always involving 3 kinds of times — the event time E, the speech time S and the reference time R. Of the various ways in which E, R, S may be temporally related, nine such ways are distinguished in the English tense system. These are *simple* past, present and future tense, past, present and future *perfect* tense, and *posterior* past, present and future tense. In addition, there is a progressive variant of each of the nine forms. Table 7.1 shows the possible configurations of these indices one could get in analyzing actual English sentences. (Commas correspond to concurrency, and $t_1 - t_2$ means t_1 is before t_2.) The relationship between S and R selects one of the past, present and future tenses: if the reference point R is at the *same time, before* and *after* the speech time S, then we obtain a present, past, or future tense, respectively.

Structure	*Tense*
S, R, E	simple present
$E, R - S$	simple past
$S - R, E$	simple future
$E - S, R$	present perfect
$E - R - S$	past perfect
$S - E - R$ $S, E - R$ $E - S - R$	future perfect
$S, R - E$	(posterior present)
$R - E - S$ $R - S, E$ $R - S - E$	(posterior past)
$S - R - E$	(posterior future)

Table 7.1: Reichenbach's Analysis of English Tense

The relative location of the event time E, with respect to the reference time R, is that E is either at the *same time as, before,* or *after* R. The simplest case is the one in which E is at the same time as R (noted as *simple* tenses in the above table), e.g., "Peter is hungry," "Peter ate sandwich," and "Peter will eat sandwich." In *perfect* variations, E is before R; e.g., "Peter had mailed the letter (when Mary called)" implies the event of Peter's mailing the letter took place before some reference point R (i.e., the time of Mary's calling).[3] For the case in which E is after R, he introduced the term "*posterior*" as English does not have a standard name for such relations. Examples are "Peter is going to leave" (*posterior present*), "Peter would leave" or "Peter was going to leave" (*posterior past*), and "Peter will be going to leave" (*posterior future*).

The notion of the reference time being at the center of the analysis, there naturally

[3]Actually, in English, the *perfect* aspect does not always imply that E *is before* R; E could extend until R. He notes that the English "present" perfect is often used in the sense of the corresponding extended tense (e.g., "Peter *has lived* in New York since 1977"), with the additional qualification that the duration of the event reaches *up to* the time of speech. However, this phenomenon of perfect is not limited to present perfect. Both past and future perfect may have an "*until* R"a reading as in "Peter had lived in New York for twelve years when he got a job offer from San Francisco" or "Peter will have lived in New York for thirty years by next March." Especially if the episode is a stative one, the reading in which the event time extends to the reference point often prevails, a point Reichenbach apparently overlooked. So, for the analysis of *perfect*, Reichenbach might have more accurately given both E *before* R and E *until* R analyses, or just E *before-or-until* R, which might later be particularized to one of the two readings. The dual analysis of the perfect has been discussed frequently by various authors, and we will have more to say about it later.

arises the question of how to locate this reference point. Reichenbach mentions a couple of rules that might be useful in determining the reference time of a clause. First, when several clauses are combined into a single sentence, the *permanence of the reference point* rule is recommended which says that the reference point is the same for all clauses within a single sentence. For example, in sentence "Peter had mailed the letter when John came," the reference time of the clause "John came" and that of the clause "Peter had mailed" are aligned at the same point as illustrated below.

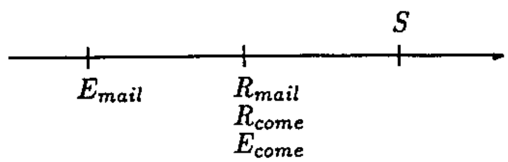

The rule also predicts sentences like "*Peter had mailed the letter when John has come" is unacceptable as the two reference times implied by this sentence are not aligned. On the other hand, the rule cannot explain why sentence "John telephoned before he came" is acceptable even though the two reference times do not coincide. (Obviously, the reference time of John's coming is earlier than that of his telephoning.) To handle such sentences, Reichenbach next suggests the rule of *the positional use of the reference point*, which does not require the reference times be located at the same point. However, he does not specify which rule should be applied when.

For time adverbials (called "time determiners" by Reichenbach) like *yesterday*, Reichenbach suggests that they uniformly refer to the reference point. He maintains, for instance, "When we say, 'I had met him yesterday,' what was yesterday is the reference point, and the meeting may have occurred the day before yesterday" [p. 294]. However, this is counterintuitive. In the dominant reading, it is the actual meeting event, rather than the reference time, which was during *yesterday*. This is even clearer in sentences like "I had first met him in 1965." This is a serious difficulty, since it was after all the perfect tenses which motivated his notion of reference time, and provided its strongest intuitive support.

Furthermore, Reichenbach had nothing to say about how to relate reference times across sentence boundaries. As we have seen, for a sequence of sentences within the same discourse segment, the reference time of a sentence is almost invariably connected to that of the previous sentence in some fashion. For example, in (7.1) we saw that the old man's leaning the mast against the wall was *right after* his (and the boy's) going into the shack, so this relation presumably holds between the corresponding reference times as well. Finding the right pairs of episodes involved in such orienting relations in successive sentences is of crucial importance for discourse/text understanding.

However, for all but sentences involving perfect *have*, it is not even clear we can coherently talk of "reference times" (other than speech time) at all. Successive sentences in a narrative, for instance, each seem to introduce a *new* time (for a new event) rather than merely referring (anaphorically) to an already existing, contextually salient time. Although these new times may be constrained by prior times in general, and easily shifted with such adverbials as "in the meantime," "before that," or "years later," they simply aren't *fixed* by the prior context. We do, of course, need to subscribe to reference times in the sense that sentences describe events or situations to which (or to whose times) we can subsequently *refer* in various ways (implicitly and explicitly). But the notion that simple past, present and future sentences each require their event times to coincide with a contextually determined reference time runs against intuitions about simple narratives. In fact, this difficulty with Reichenbach's reference times has been observed by several researchers. For instance, Harper and Charniak [1986] and Dalrymple [1988] contend that there is no reference event at all in simple tenses.[4]

Besides the problem of determining reference times and the problem of intersentential relations, Reichenbach's conception also involves fundamental problems concerning embedded clauses. For example, consider again the sentence (7.4 b) shown earlier in Section 7.1. Its temporal analysis according to Reichenbach's system will be

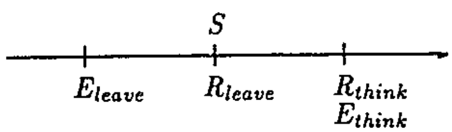

That is, according to Reichenbach's analysis, John will think that Mary's leaving took place some time before the time the speaker "uttered" sentence (7.4 b).[5] This is incorrect; it is unlikely that John would even know about the speaker's uttering sentence (7.4 b). In actuality, (7.4 b) only implies that John will think Mary's leaving took place some time before the time of his thinking as illustrated in Section 7.1. Thus, Reichenbach's system fails to take into account the local context created by syntactic embedding.

This problem arises from Reichenbach's view of sentences as flat, and tense as a top-level feature that uniformly indicates a relation with respect to the speech time S, no matter where in the sentence it appears. But tense indicates rather a relation relative to some time point given in the context. As additional examples, consider

(7.5) Peter said John *passed* the exam.

[4]Even for perfect, we ultimately want to say that, for example, a past perfect freely introduces a new past time — but because perfects are stative, that new past time tends to align itself with a previously introduced "reference time" (which is quite different from saying that it *is* a reference time). More on this in Chapter 8.

[5]Reichenbach himself discusses compound sentences involving relative clauses or embedded clauses [1947, pp. 293-4]. According to the examples he shows, each clause in a compound sentence has 3 times as usual, but, as a rule, S is common to all the clauses in the sentence. Since reference times are defined with respect to this common S, either *before*, *after* or *at the same time as* S (although there is a degree of freedom in "how much" before or after S), the above analysis appears to be the only possibility.

(7.6) Mary is going to move out tomorrow. In less than a week, she will realize that her move *was* a mistake.

(7.5) implies John's passing the exam took place (according to Peter) before Peter's saying so, rather than merely before the speech time. In (7.6), the past tense is used for a *future* event. This cannot be accounted for with a Reichenbachian analysis as it ignores the local environment in which the past tense occurs.

Another problem is his non-compositional approach of lumping together tense and aspect. Such non-compositional approaches, in conjunction with his analysis based on simple sentence types, run the risk of being ungeneralizable, and also are unsatisfactory from a modern compositional perspective. Above all, tense and perfect *have* need not co-occur. For instance, verbs and auxiliaries occur untensed in "*Failing* the exam after *having prepared* for it for a year almost devastated him" or "She is likely to *have forgotten*." It seems clear that English past, present, future and perfect are separate morphemes making separate contributions to syntactic structure and meaning. Thus, for instance, the temporal relations implicit in "John will have left" should be obtained not by extracting a *future perfect* and asserting relations among E, R and S, but rather by successively taking account of the meanings of the nested present, future and perfect operators in the logical form of the sentence. By the same token, *will* and *would* are best viewed as having separate tense (present or past) and modal (future) components, in terms of overall syntactic uniformity. This unifies the analyses of the modals in sentences like "He knows he *will* see her again" and "He knew he *would* see her again," and makes them entirely parallel to paraphrases in terms of *going to*, viz., "He knows he is going to see her again" and "He knew he was going to see her again." These latter "posterior tense" forms are patently hierarchical (e.g., *is going to see her* has 4 levels of VP structure, {*is* {*going* {*to* {*see her*}}}}, counting *to* as an auxiliary verb) and hence semantically composite on any compositional account. Moreover, *going to* can both subordinate, and be subordinated by, perfect *have*, as in "He is going to have left by then."

Attempts have been made to refine Reichenbach's theory. For instance, Hornstein [1977] tried to extend Reichenbach's theory to allow correct identification of reference events/times. However, staying within Reichenbach's noncompositional framework, his extension has the same problems as Reichenbach's. As well, Smith [1978] attempted to expand Reichenbach's system to properly handle adverbials and embedded *that*-clauses. She enumerates possible surface structure configurations for sentences with one time adverbial, sentences with two time adverbials, sentences with *that*-complements, etc., and supplies an interpretive rule for each configuration. Since the number of rules is very large (e.g., she has more than 50 rules just for sentences with *that*-complements), it is difficult to assess whether the proposed method would produce intuitively correct analyses or to tell whether her enumeration of possible configurations is exhaustive. In addition, as she adheres to Reichenbach's strategy of combining tense and aspect, she cannot handle untensed perfect VP constructs.

As will be seen later (in Section 7.2.3), researchers working in NLU, especially those primarily concerned with discourse structure, have almost invariably relied on some Reichenbach-like conception of tense (see, e.g., [Passonneau, 1988; Webber, 1988; Song and Cohen, 1991]). They adopt Reichenbach's notion that all tenses contain a reference time (in addition to speech time), and also tend to combine tense and perfect aspect into a single unit. They typically view tense and aspect as phenomena to be handled by a separate module, rather than as an integral part of a compositional mapping from text to meaning representations. It is easy to understand the appeal of this approach when one's concern is with higher-level structure. By viewing sentences as essentially flat, from which tense can be extracted as a top-level feature yielding a set of relations among E, R and S, i.e., one of nine possible tense values shown in Table 7.1, one can get on with the higher-level processing with minimum effort. However, once we descend to the lower levels of sentential structure, e.g., embedded clauses, the Reichenbachian view is ultimately unworkable.

In conclusion, while there is much that is right and insightful about Reichenbach's conception, in particular the notion of the reference time in interpreting perfects, his lumping together of tense and aspect and the assignment of E, R, S triples to all clauses are out of step with modern syntax and semantics. What is required is a *compositional* account in which operators corresponding to past, present, future and perfect contribute separately and uniformly to the meanings of their operands, i.e., formulas at the level of logical form, and a recursive technique that correctly passes the *contextual* information to the lower level sentences and analyzes syntactic constructs in a way that holds for *all* syntactic environments in which the constructs may occur.

7.2.2 The Logical Tradition

Priorean Tense Logic Approaches

A very different approach from Reichenbach's is that of various tense logics initially developed by Prior [1967]. Prior was interested in a tense logic *per se* as a modal logic, not in how to interpret natural language sentences. He was interested neither in a faithful reproduction of features of the semantic structure of any natural language, nor in its surface syntax. This is in contrast with Reichenbach who was much interested in correctly analyzing natural language sentences and their logical form although he did not go so far as to attempt to derive logical forms of English sentences.

Technically, the main difference between them is that in Prior's system evaluation of tensed sentences depends on the time they are evaluated. He did not use anything like event or time variables, but rather treated tenses as sentence operators.[6] Prior formulated

[6] Although Reichenbach is not explicit in how to incorporate tense into the logical form, we can easily

a number of simple tense logic systems; most of them have nonclassical, sentential modal operators P and F that indicate past tense and future tense respectively. Formulas can be thought of as being true or false *relative to a time* and the operators can be interpreted by the following conditions:

$$[\![P\phi]\!]^i = 1 \text{ iff for some } i', i' < i, [\![\phi]\!]^{i'} = 1\,;$$

$$[\![F\phi]\!]^i = 1 \text{ iff for some } i', i' > i, [\![\phi]\!]^{i'} = 1\,.$$

The following illustrates their use in approximating tense in English.

(7.7) a. Mary smiles
b. [MARY SMILE]

(7.8) a. Mary smiled
b. P[MARY SMILE]

(7.9) a. Mary will smile
b. F[MARY SMILE]

(7.10) a. Mary would have smiled
b. PFP[MARY SMILE]

However, many problems have been noted in trying to model English tense using Priorean tense operators. First, Kamp [1971] showed that Prior's system cannot express certain English constructs such as *now*, *until*, *since*, etc., when they occur in embedded clauses or in quantified expressions. For example, consider the following and their possible logical forms in Prior's method.

(7.11) a. Mary promised that she would be here *now*
b. P[MARY PROMISE PF[MARY HERE]]

(7.12) a. John married a girl who will exploit him
b. $P(\exists x$ [[x GIRL] $\wedge$ F[x EXPLOIT JOHN] $\wedge$ [JOHN MARRY x]])

(7.11b) fails to connect Mary's being here to the present. In fact, there does not exist an English sentence that corresponds to (7.11b), since the embedded past introduces a time even earlier than that of "promising." On the other hand, (7.12b) simply says "John married a girl who would exploit him." As a solution to this problem, Kamp proposed to

imagine that he would use tense as a predicate over existentially introduced, bound event variables. For example, I assumed in Chapter 2 that "John met Jeanne in Hollywood on Tuesday 8pm" would be translated as

$(\exists v)$ [meet(John, Jeanne)]*(v, Hollywood, Tuesday-8pm) $\wedge$ past(v).

generalize Prior's framework by allowing formulas to be evaluated at multiple "indices" instead of single times. I.e., he proposed a N operator (for "*Now*") which effectively resets the time of evaluation to the speech time. This amounts to keeping track of *two* times, i.e., the speech time and the other to evaluate expressions inside tense operators. Using Kamp's N operator, the above can be expressed as:

(7.11) c. P[MARY PROMISE N[MARY HERE]]

(7.12) c. $P(\exists x$ [[x GIRL] $\wedge$ NF[x EXPLOIT JOHN] $\wedge$ [JOHN MARRY x]])

Prior's system has further problems, though. As Dowty [1982] and many others have noted, it does not deal with sentences involving time adverbials or negation properly. For instance, consider the following sentences with their possible translations. 'Y' is a *yesterday* operator.

(7.13) a. John baked potatoes yesterday
b. YP[JOHN BAKE POTATOES], or
c. PY[JOHN BAKE POTATOES]

(7.14) a. John did not turn off the stove
b. $\neg P$[JOHN TURN OFF STOVE], or
c. $P\neg$ [JOHN TURN OFF STOVE]

Assuming the interpretation of Y is

$$[\![Y\phi]\!]^{i} = 1 \text{ iff for some } i', \text{ such that } i' \text{ falls within the day preceding } i, [\![\phi]\!]^{i'} = 1,$$

one can easily verify that neither of the translations above for (7.13) is correct. (b) asserts John's baking took place sometime *before yesterday*, while (c) asserts there is some time in the past such that John baked potatoes on one day prior to that. That's not what (7.13) means. And there is no other alternative translation of (7.13 a) in Prior's system. Similarly, neither of the translations above for (7.14) is correct; (a) implies that John has never turned off the stove, whereas (b) only implies that JOHN NOT TURN OFF STOVE is true some time in the past, which is trivially true unless all John has been doing since the beginning of time is turning off the stove!

Thus, tense logics based on the classic Priorean logic are inadequate for representing the logical form of tensed sentences in narratives. First, they are indexical, i.e., utterances interpreted at different times will have truth conditions dependent on those times—but there is no provision *in the logic* for recording those times. Second, they do not provide time or event variables; the latter are especially important to represent causal relations as we saw in Chapter 2. Third, embedded and other complex constructions are not properly handled.

Finally, they do not handle interactions between tense operators, temporal adverbials, and negation properly.

Dowty's First, Syntactic Approach to Tense-Aspect Interpretation

Noticing that previous work on tense and aspect is not adequate for the analysis of natural language sentences, Dowty came up with two successive proposals [1979; 1982], aimed at modelling tense, aspect and time adverbials in English. I will first discuss Dowty's earlier proposal [1979, Chapter 7].

In contrast to Reichenbach and Prior, Dowty was committed to formalized, systematic transduction from English sentences to their logical forms, following the tradition of Montague. Although Montague showed that logic can be successfully applied to the study of natural language, on the the principles of *compositionality* (the meaning of an expression is determined by the meaning of its parts) and *truth conditions* (the meaning of a declarative sentence determines the conditions under which that sentence is true), his tense-aspect analysis was rather rudimentary. Dowty's first system, which he calls a "syntactic" solution for a reason to be explained shortly, was developed as a part of his effort to extend Montague Grammar to handle tensed constructions correctly.

Recognizing that Prior's system fails for sentences with time adverbials, Dowty paid particular attention to the interaction between tense and time adverbials in his proposal. He came to believe that tenses are parasitic on time adverbials and that the solution lies in introducing tense and time adverbials together, syncategorematically, by a single syntactic rule. Thus, he let time adverbials introduce existentially quantified variables over time intervals into translations. For example, *yesterday* is translated as follows.

$$\text{YESTERDAY}' = \lambda T\, \exists t\, [t \subseteq \mathit{Yesterday} \wedge T(t)]$$

Tense operators may or may not introduce existential time variables. Specifically, they introduce existential variables over intervals if there is no time adverbial in the sentence; otherwise, they introduce λ-variables (as time adverbials already have brought in existential variables). For example, sentence ϕ in past tense is translated in one of the two ways as follows (as re-presented in [Dowty, 1982]):[7]

Translation of past tense sentences (Dowty'79).
For ϕ in past tense,
(i) for α a time adverbial, expression $\underline{\phi\,\alpha}$ is translated into
$\alpha'(\lambda t\, [t < t^* \wedge \text{AT}(t, \phi')])$;

[7]Dowty re-presented his earlier work in a slightly different form, albeit semantically equivalent. In his original proposal [1979], he used expressions PAST(t) and FUT(t), instead of $t < t^*$ and $t > t^*$.

(ii) if ϕ is *without* a time adverbial, it is translated into
$\exists t[t < t^* \wedge \mathrm{AT}(t, \phi')]$.

Here, t is a variable over time intervals, t^* is an indexical temporal constant, and AT is an operator such that for γ an expression denoting a time, $\mathrm{AT}(\gamma, \phi)$ means ϕ is true at γ. The following shows the interpretation of t^* and the AT operator and some of their logical properties.

For t, t' variables over time intervals and i a time interval,
(i) $[\![t^*]\!]^i = i$
(ii) $[\![\mathrm{AT}(\gamma, \phi)]\!]^i = 1$ iff $[\![\phi]\!]^{i'} = 1$, where $i' = [\![\gamma]\!]^i$
(iii) $\phi \equiv \mathrm{AT}(t^*, \phi)$
(iv) $\mathrm{AT}(t, \mathrm{AT}(t', \phi)) \equiv \mathrm{AT}(t', \phi)$

(i) says that at any time interval i, t^* denotes the interval i itself. The AT operator in (ii) is much like '$*$' in EL, except that γ denotes a time rather than a situation. I.e., an ELF $[\phi * \gamma]$ can be approximated as $\mathrm{AT}((\textit{time-of}\ \gamma), \phi)$ in Dowty's system. (iii) indicates that to say "ϕ is true AT t^*" is equivalent to simply saying "ϕ is true." (iv) shows that when two or more AT operators are nested, all but the innermost is vacuous (unless the time expression of the inner AT is t^*).[8]

The following examples illustrate the above rules.

(7.15) a. John left
b. $\exists t\,[t < t^* \wedge\ \mathrm{AT}(t,\ \textit{leave}(\textit{John}))]$

(7.16) a. John left on Thursday
b. $\exists t\,[t \subseteq \textit{Thursday}\ \wedge\ \ t < t^*\ \wedge\ \ \mathrm{AT}(t,\ \textit{leave}(\textit{John}))]$

Note that the final translations (b)-formulas still contain the indexical constant t^*, whose interpretation will be the time of evaluation, in this case, the speech time.

However, there are two serious problems as Dowty acknowledges. First, having separate rules depending on whether sentences are with or without time adverbials is rather unnatural as it amounts to treating tense plus a time adverbial as a single syntactic constituent. Second, the treatment does not work if a sentence has more than one time adverbial as in "John saw Mary *at noon yesterday*."

[8] Relative to a fixed value for t, the truth value of $\mathrm{AT}(t, \phi)$ does not depend upon the time at which this whole formula is evaluated, but only upon the truth value of the inner formula ϕ at the time denoted by t. Such formulas are called *eternal sentences* (in the terminology of EL, *atemporal* or *unlocated* sentences) because if they are true, they are true at all times. Note that Dowty is concerned with times only, not with places at which ϕ is true.

For *perfect*, Dowty again has two rules depending on whether there is a time adverbial, as shown below.

> *Translation of (present) perfect VPs (Dowty'79).*
>
> For β in (perfect) past participle form,
>
> (i) for α a time adverbial, the sequence of syntactic categories $\underline{\text{have } \beta\ \alpha}$ is translated into
>
> $\lambda x\, [\alpha' \,(\lambda t\, [\mathrm{XN}(t, t^*) \wedge \mathrm{AT}(t, \beta'(x))])]$;
>
> (ii) the sequence of syntactic categories $\underline{\text{have } \beta}$, without a following time adverbial, is translated into
>
> $\lambda x\, \exists t\, [\mathrm{XN}(t, t^*) \ \wedge\ \exists t_1 [t_1 \subseteq t \ \wedge\ \mathrm{AT}(t_1,\ \beta'(x))]]$,

where, if γ, γ' are expressions denoting time intervals, $\mathrm{XN}(\gamma, \gamma')$ means that γ is an "extended now" of γ' (i.e., γ denotes an interval beginning in the past and extending up to and including the interval denoted by γ').

According to these translations, the present perfect serves to locate an event within a *period of time that began in the past and extends up to the present moment* (called *extended now*; see [Bennett and Partee, 1978; McCoard, 1978]), and this period may further be specified by a time adverbial. For the analysis of past perfect, Dowty suggests one first apply either of the above two perfect rules, and then apply the past rule. Then, one would get an *extended now* (i.e., "until the reference time") analysis when there is a time adverbial; and an *embedded past* (i.e., "sometime before the reference time") analysis, otherwise. Here are some sample sentences in perfect and their translations.

(7.17) a. John has worked

b. $\exists t\, [\mathrm{XN}(t, t^*) \wedge \exists t_1 [t_1 \subseteq t \ \wedge\ \mathrm{AT}(t_1,\ \mathit{work}(\mathit{John}))]]$

(7.18) a. John has worked today

b. $\exists t [t \subseteq \mathit{Today} \wedge \mathrm{XN}(t, t^*) \wedge \mathrm{AT}(t, \mathit{work}(\mathit{John}))]$

(7.19) a. *John has worked yesterday

b. $\exists t [t \subseteq \mathit{Yesterday} \wedge \mathrm{XN}(t, t^*) \wedge \mathrm{AT}(t, \mathit{work}(\mathit{John}))]$

(7.20) a. John had left on Thursday

b. $\exists t' [t' < t^* \ \wedge\ \exists t [t \subseteq \mathit{Thursday} \ \wedge\ \mathrm{XN}(t, t') \wedge \mathrm{AT}(t, \mathit{leave}(\mathit{John}))]]$

Since (7.17a) does not contain a time adverbial, it gets the reading "There exists a time interval that extends to the speech time, and it has a *subinterval* in which JOHN WORK is true," as (7.17b) shows. Although for JOHN WORK, this may sound fine, such an analysis may not be satisfactory for sentence like "John has been crying." In contrast, (7.18a) contains a time adverbial and hence gets the reading "There exists a time interval that extends to 'now', i.e., t^*, and also is part of 'today,' and JOHN WORK is true over that interval." Obviously, such a reading will not be suitable for sentences like "John has called today." (7.19 b) reads "There exists a time interval that extends to 'now', which is also part of 'today,' and JOHN WORK is true over that interval." Note that this leads to a contradiction and indeed the original (7.19 a) is not acceptable English. Finally, (7.20 b) means John's leaving event extended till the reference time, t' in this case, which happens to be before the speech time t^* and at the same time was during Thursday. (Recall that t^* always evaluates to the time coordinate of the index at which it is evaluated.) This rather dubious *extended now* reading is forced because of the time adverbial contained in (7.20 a).

Together with the problems noted earlier, this splitting of the perfect analysis based on whether the clause contains a time adverbial, yielding intuitively unsatisfactory translations in many cases, makes Dowty's first proposal rather unattractive.

Dowty's Second Proposal with Double Indexing

As Dowty himself was dissatisfied with treating tense plus a time adverbial as a syntactic constituent and also with the lack of generality in the whole approach, he revised his position considerably and proposed several alternatives in [Dowty, 1982]. This "quasi-Reichenbachian" system is also reminiscent of Kamp's two-dimensional tense logic, but more flexible as will be seen. In this system Dowty develops an independent treatment of tenses and time adverbials and uses double indices resembling Reichenbach's reference time R and speech time S. Truth is then defined relative to these two indices, that is,

> *'ϕ is true at $\langle i, j\rangle$' is interpreted as 'ϕ is true when uttered at j and used to talk about the time i'.* [p. 37]

Like the earlier one, the new system makes use of the indexical constant t^* that always evaluates to the reference time, and the AT operator meaning "true at." The following clauses show the semantic details.

For i, j time intervals:

(i) $[\![\text{PRES } \phi]\!]^{i,j} = 1$ iff $[\![\phi]\!]^{i,j} = 1$ and $i = j$;

(ii) $[\![\text{PAST } \phi]\!]^{i,j} = 1$ iff $[\![\phi]\!]^{i,j} = 1$ and $i < j$;

(iii) $[\![\text{FUT } \phi]\!]^{i,j} = 1$ iff $[\![\phi]\!]^{i,j} = 1$ and $i > j$;

(iv) $[\![\text{WOULD } \phi]\!]^{i,j} = 1$ iff $[\![\phi]\!]^{i',j} = 1$,
where i' is some interval later than i ;

(v) $[\![\text{HAVE } \phi]\!]^{i,j} = 1$ iff $[\![\phi]\!]^{i',j} = 1$,
where i' is some interval of which i is a final subinterval ;

(vi) $[\![t^*]\!]^{i,j} = i$;

(vii) for t an expression denoting a time interval,
$[\![\mathrm{AT}(t,\phi)]\!]^{i,j} = 1$ iff $[\![\phi]\!]^{i',j} = 1$, where $i' = [\![t]\!]^{i,j}$.

The tense analysis in (i)–(iii) resembles Reichenbach's analysis of simple tense. Index i in (i)–(iii) corresponds to Reichenbach's reference time R, and j corresponds to speech time S. Dowty assumes that the reference time, i.e., i, is supplied by some kind of *narrative discourse rule*, especially for a sequence of past tense sentences. Unfortunately, by omitting Reichenbach's third index, i.e., event time E, Dowty could not compositionally handle expressions involving *would*. According to clauses (i)–(iii), [PAST [FUT ϕ]] yields $i<j$ and $i>j$, a contradiction. Thus, he had to introduce a separate rule to handle WOULD as shown in (iv). Note that the WOULD operator shifts the point of evaluation to some new reference time later than the embedded reference time. In (iv), i' is essentially equivalent to Reichenbach's event time E. As shown in analysis (v), HAVE also shifts the point of evaluation of its embedded sentences from i to i'. Again, i and i' correspond to Reichenbach's R and E, respectively. Note that Dowty has only one perfect analysis in the new system, namely, the *extended now* reading (without the *embedded past* one). Clauses (vi) and (vii) are essentially the same as in the earlier system.

Adverbials themselves are translated basically as in [Dowty, 1979]. Dowty treats both time adverbials and durative adverbials as directly modifying the reference time, similarly to Reichenbach's approach. The following are syntactic rules for combining adverbials with sentences.

Translation of sentences with time adverbials/durative adverbials (Dowty'82).

(i) for α a *time adverbial* and ϕ a formula, expression $\underline{\phi\,\alpha}$ is translated into $\alpha'(\lambda t[t = t^* \wedge \phi'])$;

(ii) for α a *durative adverbial* and ϕ a formula, expression $\underline{\phi\,\alpha}$ is translated into $(\alpha'\ \phi')$.

Here are examples of sentences with adverbials and their translations.

(7.21) a. John left on Thursday last week
b. $\exists t_1, t_2\ [t_1 \subseteq \text{Thursday} \wedge t_1 = t^* \wedge t_2 \subseteq \text{Last-week} \wedge t_2 = t^* \wedge$ PAST[leave(John)]]

(7.22) a. John worked for an hour
b. PAST[an-hour-long(t^*) $\wedge\ \forall t\ [t \subseteq t^* \rightarrow \mathrm{AT}(t, work(John))]$]

Suppose we evaluate (7.21b) at ⟨i, j⟩. Then, it is true iff $i < j$, i is contained in *Thursday* as well as in *Last-week*, and John actually left at i. Similarly, (7.22b) is true at ⟨i, j⟩ iff $i < j$, i is an hour long, and John worked at all subintervals of i. These analyses agree with our intuition. Note that Dowty now successfully treats multiple adverbials. Also note here that the information in the sentence and in the time adverbials are essentially interpreted *conjunctively*, which seems to be a nice feature.

However, the final logical form is still indexical, as in situation semantics. It is not only t^* and PAST which are indexical, but *every* expression, in the sense that to fix the semantic value of any expression, we need the values of the two indices, i and j. To be able to make inferences from sentences interpreted in different contexts, we will thus need to *deindex*. It is easy to see how to get rid of dependence on index j—we can use the clock time of the utterance and the AT operator. But reference time i is a problem. It appears, for instance, that in a text passage in the simple past, we already need to have a definite reference time for the *next* event to be reported. It is as if we already knew, given something like "John left on Thursday," *where in time* the next reported event must lie. But in view of such continuations as "He forgot to lock the door," "The house was empty without him," or "He returned two weeks later," this is thoroughly counterintuitive. In other words, the notion of "reference time" seems just as out of place in the analysis of simple senses here as in Reichenbach's version.

Let us now consider embedded clauses. For translation of embedded clauses (relative clauses or *that*-clauses), Dowty presents a few alternative S-complement rules. He always systematically introduces a new reference time for each subordinate clause, with the aid of the AT operator. So, in case of embedded *that*-clauses, he introduces an AT operator for each propositional object in the translation of attitude verbs. Here are the three alternative translation rules Dowty suggests.

Possible translations of S-complements. E.g., *believe* may be translated into

(i) $\lambda p\, \lambda x\, \exists t\, [\text{believe-that}(x,\ {}^\wedge\text{AT}(t,\ {}^\vee p))]$;

(ii) $\lambda p\, \lambda x\, [\text{believe}'({}^\wedge\text{AT}(t_k,\ {}^\vee p))(x)]$,
where t_k is a free variable, for any integer k;

(iii) $\lambda p\, \lambda x\, \exists t\, [t \leq t^* \wedge \text{believe-that}(x,\ {}^\wedge\text{AT}(t,\ {}^\vee p))]$.

According to rule (i), the meaning of, say, "John knew that Mary was sick" is equivalent to "There are certain times i and i' in the past (i.e., before speech time) such that Mary was sick at i', and he knew this at i." The trouble is that there is specified no relation between the value of t_1 (i.e., i') and i. That is, i', the time of Mary's being sick, could be even later than i, John's knowing it, which is not meant by the original English sentence. What's needed is to pass the reference time of the main clause to the embedded clause.

Apparently, Dowty recognizes this problem, and mentions some possible fixes shown in (ii) and (iii). In (ii), he treats subordinate clause reference times an indexical, namely, t_k. However, the difficulty is how to obtain t_k. Rule (iii) makes use of t^* instead of t_k. Let us now consider rule (iii) with an example.

(7.23) a. John will think Mary will visit Paris
b. FUT $\exists t[t \leq t^* \wedge$ think-that(John, $^\wedge$AT(t, [FUT visit(Mary, Paris)]))]

Suppose (7.23 b) is evaluated at $\langle i, j\rangle$, $i > j$. Then it is true iff there exists a value i', $i' \leq i$, for t such that

$$[\![\text{FUT visit(Mary, Paris)}]\!]^{i',j}$$

is true, i.e., $i' > j$, where $[\![t]\!]^{i,j} = i'$, and John thinks at time i that Mary visits Paris at time i'. The configuration of these times is as shown below.

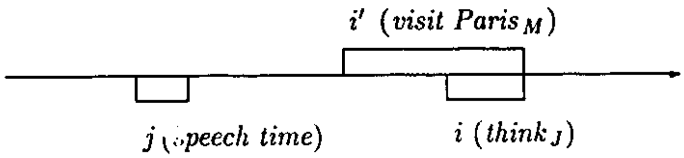

Contrary to (7.23 a), this analysis says that John will think that Mary's visit to Paris is just ending at the time of his thinking! Thus, Dowty's new system with double indices still has some problems for embedded sentences.

As another point, while Dowty aims at a systematic mapping and provided formalized, compositional accounts, he took "context" for granted. Obviously, Dowty realizes that context is essential to interpretation. For example, he points out that Partee's [1973] often cited sentence,

(7.24) I did not turn off the stove
PAST $\neg$[I turn off the stove]

can be interpreted as "At a certain past time which is talked about, I did not turn off the stove." To provide that "certain past time," he suggests that the reference time i be used as a contextual parameter, much like other indices such as speaker, hearer, utterance time, etc. But he does not provide any formal machinery for determining the values of this reference time, and as noted earlier, it is doubtful that the notion of "reference" time should even play a direct role in the truth conditions of tense and aspect operators.

In conclusion, Dowty's second proposal still involves problems with embedded sentences as seen above and with perfect analysis, which provides only an "extending until the reference time" reading. Also, as pointed out earlier, the final logical form is still indexical, to be evaluated at two indices i and j, with no mechanism available for locating one of these indices, namely, i, the reference time. In addition, the lack of event variables in the logical form weakens Dowty's proposal as a candidate for practical use in general NLU.[9] However, it is in many ways appealing—especially, in its compositional approach, truth conditional semantics, and conjunctive treatment of the information in the tense operators and in the time adverbials.

Hinrichs' Reichenbach-like Tense Logic

I will briefly remark on one more tense interpretation system, one developed by Hinrichs [1987; 1988]. Hinrichs tries to amalgamate into his system various approaches discussed so far, i.e., Reichenbach's, Prior's, and Kamp's. He tries to repair Prior's tense logic—especially in the treatment of the interaction between tense, negation, and adverbials—by introducing multiple indices like Kamp and Dowty. Whereas Kamp and Dowty each used two indices, however, he nominally makes use of all three of Reichenbach's indices. Thus, his system could be described as a Reichenbachian system reformulated as a tense logic, with a possible world approach. On closer inspection, however, one finds that his so-called reference time plays a role quite different from the Reichenbachian one. It is, in fact, more like a "time frame" to which the event time is confined. For instance, the past operator is defined as follows.

> For ts, tr, and te standing for speech time, reference time, and event time, respectively,
> $[\text{PAST p}]_{ts,tr,te} = 1$ iff $[\text{p}]_{ts,tr,t'} = 1$ for some time t' such that $t' < ts$ and $t' \subseteq tr$.

Note that the event time t' for the past-embedded formula is required to fall *within* the external reference time, tr, and that the externally supplied event time, te, is ignored. With this formulation, Hinrichs is able to handle sentences like (7.13a) that Prior couldn't. In actual implementation, ts and tr are used in the object language as indexical constants, and events times are existentially quantified variables over times. For instance, *arrived* is translated into $\lambda x[\exists t'[t' < ts \ \& \ t' \subseteq tr \ \& \ \text{arrive}'(x)(t')]]$. Note also that the above rule does not have explicitly $tr < ts$ Reichenbach originally had in his system. Hinrichs instead assumes that tr can be provided by the context, i.e., this time frame-like reference time will be updated after processing each sentence in a discourse (see [Hinrichs, 1986]).

One interesting point though is his treatment of NPs as temporally indexical. Following Enç [1981], he takes evaluation of verbal and nominal predicates to be independent of one another, and assigns tense scope only over the main verb of the sentence. For instance, he

[9] Dowty's semantics, just like Montague's, was taken on a possible worlds approach. Intuitively, as mentioned in Chapter 2, a possible world is a *complete*, i.e., *total*, specification of how things are, or might be, in every semantically relevant detail. Sentences denote—merely—truth values in worlds, usually at a certain time. Hence, problems arise as soon as attitudes are involved.

translates "Every admiral graduated from Annapolis" into "Every admiral *in a restricted set* graduated from Annapolis." For this, he introduces a predicate R that ranges over properties that are salient in a given context. The translation of "every" then contains the predicate R as well as an existential time variable t, as follows:

$$\text{EVERY}' = \lambda P \lambda Q \forall x[\exists t[P(x)(t) \ \& \ R(x)(t)] \rightarrow Q(x)] \, .$$

R could be instantiated, e.g., as the intension of the set of individuals x who are in the Pacific Fleet at a time that is the same as ts, i.e., by

$$\lambda t \lambda y[\text{be-in}'(\text{Pac-Fleet}')(y)(t) \ \& \ t = ts] \, .$$

Obviously, finding out the right intension to instantiate R is a nontrivial task, and Hinrichs does not offer any suggestions. However, this concept turned out to be important in our research on quantifier deindexing, and I will say more about this in Chapter 9 where I discuss deindexing of quantifiers.

Finally. Hinrichs limits his discussion to simple sentences and has nothing to say about how to interpret embedded clauses. Furthermore, like Reichenbach, he lumps together tense and perfect aspect into PRESPERF, PASTPERF, etc. [Hinrichs, 1987, p. 29], and hence cannot handle untensed perfect constructs.

7.2.3 AI Approaches

So far, I have discussed various approaches to tense-aspect interpretation proposed by semanticists and philosophers. I now turn to approaches with a more computational orientation, proposed mainly by AI researchers. I will discuss works by Passonneau, Webber, and Song and Cohen, in that order. Related work by Lascarides *et al.* will be discussed in Chapter 8.

Situation Types and Temporal Structures a la Passonneau

Passonneau [1987; 1988] developed a system for intrasentential temporal structure analysis whose distinctive feature is the use of a simple taxonomies of situation types (i.e., aspectual classes or "Aktionsarten") to refine a Reichenbach-like analysis of simple tenses and perfect and progressive aspects. Her goal was to analyze the temporal structure of simple declarative sentences. She limits her analysis to actual situations, i.e., situations that have already occurred or are occurring.

She classifies situations into three major categories: states, processes and transition events. This classification is in turn to be understood in terms of two sorts of distinctions among "time interval arguments" associated with event descriptions: an *active/stative* distinction, and a *bounded/unbounded/unspecified* distinction.[10] The active/stative distinction corresponds to the presence or absence of *changes* over that time interval;[11] the bounded/unbounded/unspecified distinction corresponds to the possibility of the extending the time interval over which the given event description holds. So, a *state* is a situation that holds over a stative, unbounded interval. A *process* is a situation that holds over an active interval which is either unbounded or unspecified. If the verb is in progressive form, the interval is unbounded; otherwise, unspecified. For instance, the interval associated with the situation described by "The alarm was sounding at 8am" is unbounded, as it may be extended both forward and backward in time. In contrast, the interval for "The alarm sounded at 8am" is unspecified as one cannot infer a temporal extension. (The alarm may not have sounded before 8am.) A *transition event* is a complex situation that consists of a *process* which culminates in a new state or process. Thus, the temporal structure of a transition event consists of an active interval bounded by another interval that may be either active or stative.

The task for temporal analysis is then to find out, for each sentence, the relation between the interval argument of the described situation and the corresponding Reichenbachian event time ET. (Passonneau uses Reichenbach's ET, RT, ST triples in her analysis.) For this, she defines relations for each pair of situation types (i.e., time interval argument types) and Reichenbachian event time ET as shown below. (The distinction between 'includes' and 'has' below seems to be that of 'properly contains' and 'contains'.)

Situation Type		Time Argument		Event Time
State		stative	unbounded	includes ET
Process	(progressive)	active	unbounded	includes ET
	(nonprogressive)	active	unspecified	has ET
Transition event		active	bounded	unifies with ET

Next, using Reichenbach's ET, RT, ST configuration, further specification of the temporal structure is possible. In the case of *present-unbounded*, ET is located sometime within the interval coincident with ST. Processes or transition events in present tense are eliminated as she does not cover habitual or repetitive events. In the case of *past-unbounded*, ET is prior to ST *within* a persisting interval, and the same situation extends unchanged forward towards the present and back into the past.[12] In the case of *past-unspecified*, as

[10] We attempted a formalization of an aspectual system similar to Passonneau's in Section 6.2.4.

[11] It seems that "time intervals" as here understood are uniquely associated with events; i.e., even concurrent events have distinct "time intervals" associated with them. Thus, to say that a time interval is stative is not to say that the entire world stands still during it.

[12] Passonneau's intention here is to be able to explain sentences like "The pump was failing, and is still failing." Past tense in reference to unbounded interval does not apply to the whole duration, but to the ET within the interval.

in "The pump operated," ET may or may not be an endpoint of an interval. In the case of *transition event–past*, as in "The pump failed," ET locates an endpoint as shown in the above table.

In the PUNDIT system, which was the context for Passonneau's work (see also [Dahl *et al.*, 1987]), verbs are lexically decomposed, and the decomposition incorporates information about situation types, as in [Dowty, 1979]. For instance, the decomposition of (intransitive) verb "fail" is of form

become(inoperative(patient( ))).

The temporal component of the system analyzes each sentence by first checking its progressive aspect and verb category (situation types), and then simply looking in the above table to find out the corresponding temporal relation. Additional Reichenbachian relations between RT, ET and ST are obtained from tense and perfect aspect. I now show a couple of sample sentences and the temporal structures that would be obtained in this method.

(7.25) The pump failed

Lexical aspect = *transition event*
Progressive = *no*
Situation representations:
event(f1, become(inoperative(patient(pump1))), moment(f1))
state(f2, inoperative(patient(pump1)), period(f2))
where *start(moment(f1), period(f2))*
Event Time: *ET is unified with f1*
Perfect = *no*
Tense = *past*; thus, ET *is* $RT < ST$

(7.26) The pump is failing

Lexical aspect = *process*
Progressive = *yes*
Situation representation:
process(f1, become(inoperative(patient(pump1))), period(f1))
Event Time: ET = *moment(f1) such that includes(period(f1), moment(f1))*
Perfect = *no*
Tense = *pres*; thus, ET *is* $RT = ST$

Note in (7.26) that Passonneau assumes that if something is failing, it necessarily fails. In the above examples, RT does not have a role, but in a context where there are two or more clauses/sentences connected together, it would play one. (Passonneau gives some hints as to how simple temporal adverbial clauses may be handled.)

Notice that this approach is methodologically quite different from the ones discussed earlier. Its advantage is that it produces a simple analysis for simple sentences without much difficulty. However, the range of expressions that can be handled by this method seems to be very limited. Note that Passonneau eliminates sentences involving future, embedded clauses, negation, modals, intension, or repetitions. Thus, one question that remains is whether, and how, her method can be extended to handle a wider range of natural language constructs. Another question is whether, and how, her ontological distinctions (events, time interval arguments, etc.) and the concepts underlying her taxonomy of situations can be put on a formal, model-theoretical footing.

Tense as Discourse Anaphor and Webber's Heuristics

While Passonneau focused on temporal structure within individual sentences, Webber [1987a; 1988] looked at the problem in a discourse context. She was particularly interested in the anaphoric nature of tense and aspect in narratives, and in the context dependency to which intersentential temporal connections give rise.

She assumes that a listener's developing discourse model represents, among other things, the events and situations being discussed, along with their relationships with one another. She calls this the **E/S (event/situation) structure.** The question then arises where in the evolving E/S structure to integrate the event or situation described in the next clause. Webber argues that at any point in the discourse, there is one entity in the E/S structure that is most attended to, and hence most likely to stand in an anaphoric relation with the tense of the next clause. To capture this idea, she introduces a discourse-level focus mechanism called *temporal focus* (TF). TF is a dynamically moving focus that resembles the *discourse focus* (DF) of Grosz and Sidner [1986]. The novelties are that TF grounds the context-dependency of tense and that focus management heuristics can be used to track the movement of TF.

A couple of tools she uses need be explained before we discuss her method in detail. First, she employs a Reichenbach-like configuration of times ET, RT, ST for representing tense. Among these, RT is taken as the basis for anaphora; i.e., the event described by a new clause is linked to the existing E/S structure via its RT. TF is usually the ET of a previous utterance and is most likely to be used as the referent of the RT of the next utterance. Second, she adopts the tripartite event ontology [Moens and Steedman, 1988]. That is, events have a structure consisting of three phases: preparatory, culmination, and consequence.[13] The movement of TF depends on which phase of the event structure is

[13]One needs to understand that the term "preparatory" (or "consequence") phase is used here in a fairly loose way. It does not have to be an intrinsic part of the event structure, and could be a lengthy time period that precedes (or follows) the event. Consider, for instance, "John left for Europe last Friday. He will be back in three months." As will be seen, according to Webber's analysis, the event of John's coming back is at sometime during the consequence phase of his leaving event, which requires the consequence

attended to. When the consequence phase of a prior event is attended to, TF moves forwards to that consequence phase; when the preparatory phase is attended to, TF similarly moves backwards; and when the culmination phase is attended to, TF does not move. RT of the new sentence is linked to the event that is the current TF.

So, in a successive pair of sentences that belong to the same discourse segment, there may be as many as three ways of linking the new event (more precisely, its reference time RT) to the current E/S structure. If the new clause uses the perfect, TF always stays at the same place. In other cases, one of the three options may be selected based on inferences made from world knowledge. In a simple narrative, TF is likely to gradually move forward.

Let us now consider a simple example from [Webber, 1988].

(7.27) a. John went into the florist shop.
b. He had promised Mary some flowers.
c. He picked out three red roses.

When (7.27 b) is processed, TF will be at event E_a of John's going to the florist shop. As mentioned above, when (7.27 b) is encountered, TF does not move in the E/S structure, as the sentence uses the perfect. Thus, R_b is linked to E_a, and E_b is entered in the E/S structure such that $E_b < R_b$, as required by the perfect. Next, (7.27 c) is in past tense, and at this point, there are three possibilities in terms of the movement of TF. With inferences based on world knowledge, a listener is likely to move TF forward to Conseq(E_a), the consequence phase of E_a, to which R_c is linked. Since the new sentence is simple past, we get $E_c = R_c <$ ST. The temporal structure then looks as follows.

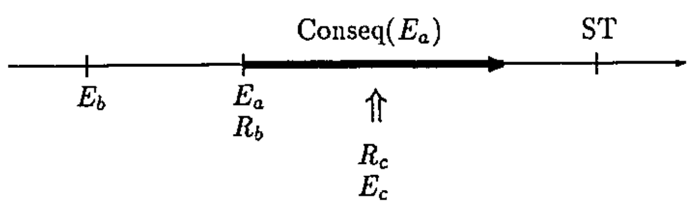

We will look at another example; this time, one in which TF moves backwards.

(7.28) a. John gave Mary three red roses.
b. He bought them at the florist shop on the corner.

In this case, the most plausible choice at (7.28 b) for focus maintenance will be TF moving backwards, thus linking R_b to Prep(E_a), as shown below.

phase to be interpreted as large enough to overlap the event of his coming back. The same holds for the preparatory phase; cf., "John will be back next month. He went on a world tour two months ago."

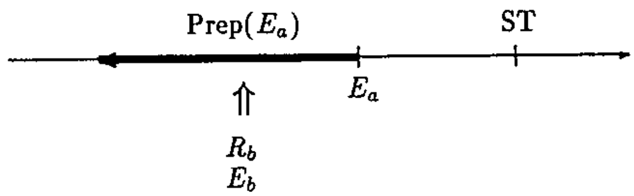

Although in general time progresses in simple narratives, the above shows this isn't always the case. Predicting TF movements is a nontrivial task, and recently it has received much attention (cf., [Song, 1991; Song and Cohen, 1991]). I will have more to say about this later.

The two examples we saw involve local movements of TF within a discourse segment. Webber also discusses complicated discourse situations in which TF may move across discourse segment boundaries. To keep track of and predict the movement of TF, she proposes an *embedded discourse heuristic* and *focus resumption heuristics*, in addition to a *focus maintenance heuristic* that predicts regular local movements. (We have seen the behavior of latter already.) She shows the following variation of (7.27).

(7.29) a. John went into the florist shop.
b. He had promised Mary some flowers.
c. *She said she wouldn't forgive him if he forgot.*
d. He picked out three red roses.

Webber notes that (7.29 c) is most sensibly interpreted with respect to the "promising" event E_b. She assumes that when the listener recognizes an embedded discourse segment, s/he stores ("caches") the current TF for possible resumption later. That is, TF moves from its current position E_a to E_b, *caching* E_a for possible resumption later. Following this gross movement, local TF movement maintenance technique may be applied as usual. Next, at (7.29 d), the embedded segment comes to an end, and the flow returns to the main segment that has been in suspension. Thus, the previously cached E_a is resumed, after which local movement maintenance techniques may be applied further. The following configuration shows the result of the analysis.

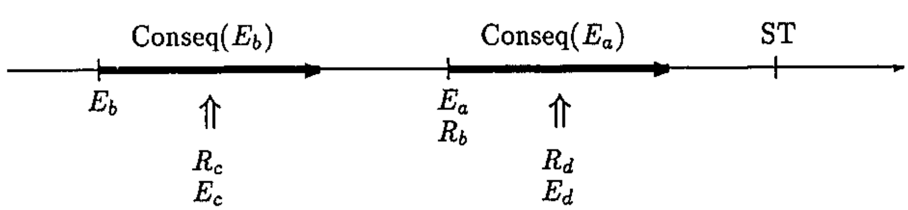

This method may well become practical once a reliable way of detecting segment boundaries becomes available. Important as it is, however, detecting discourse segment boundaries remains a largely unsolved problem, and that is another issue to which I will return.

Song and Cohen's Deterministic Approach

Webber's proposal may be considered nondeterministic in that it leaves open (i) what phase of the currently focused event is to be selected as a linking point for the tense of the next sentence and (ii) when to start an embedded segment and when to resume the previous segment. Webber's assumption concerning (i) was that the most plausible interpretation may be obtained by inference. Song and Cohen [Song, 1991; Song and Cohen, 1991] worked toward making the interpretation process deterministic, especially with respect to (i). Their deterministic algorithm is based on two main ideas: elimination of incoherent tense sequences, and heuristics for choosing among alternative ways of moving TF, especially preference of progression over elaboration for a sequence of the same tense.[14]

In their technical apparatus, they adhere to Passonneau's [1987; 1988] classification of situations and Reichenbach's SRE tense system. For temporal focusing, they use a SRE triple called a *temporal focus structure* (TFS), which is similar to Webber's TF but allows both ET and RT to be tracked. Whereas Webber considered only RT as anaphoric, they take both ET and RT as anaphoric. The following sort of example provides some support for their position.

"John had spotted Mary, and had followed her down the street."

Note that here not only the reference times are linked, but in addition the "following" event is linked to (or, oriented by) the "spotting" event. (Also recall the examples seen in Section 7.1.)

The concept of coherent tense sequence applies to successive sentences in simple narratives. For instance, Song and Cohen assume that a present tense sentence may be followed by a past tense sentence, but not by a past perfect sentence. The latter produces an incoherent tense sequence in English, and their algorithm halts for such input.

Exceptions to their rules do seem rare in a statistical sense, though it is not hard to make up counterexamples. Below are many of the tense sequences they assume incoherent, with my counterexamples.

- *Present — Past perfect*
  "Mary is altering a dress. Her grandmother had made it for her wedding."
  "Mary is angry about the accident. The other driver had been drinking."

[14] They use the term "tense sequence" to refer to a succession of tenses as they occur in successive sentences.

- *Present — Future perfect*
  "John is on his way to Mary's. But she will have left for the airport by the time he arrives at her apartment."

- *Present — Posterior past*
  "Mary is watching TV. She was going to go out tonight. But she changed her mind at the last moment."

- *Present perfect — Past perfect*
  "John and Mary have divorced each other. They had been fighting like cats and dogs for years."
  "They have bought the house. They had scrimped for years to save up the down payment."

- *Present perfect — Future perfect*
  "John has worked for this company since he was a young man. In fact, he will have worked for exactly 30 years next December."

None of these mini-stories sound unusual, but will be rejected by their algorithm as being incoherent. Indeed it is doubtful that one can *reliably* detect temporal incoherence on purely syntactic grounds.

For disambiguating temporal relations, the essence of their algorithm is as follows.

(1) If a new sentence has the same tense-aspect as the one currently in focus, then (i) if the situation described by the new sentence is unbounded, ET_{new} (event time of the new sentence) is the same as ET_{focus} (event time of the one currently in focus), (ii) otherwise, $ET_{focus} < ET_{new}$.

(2) If (i) the sentence in current focus is present perfect and the new sentence is past, or (ii) the currently focused sentence is posterior present and the new sentence is future, then $ET_{new} \leq ET_{focus}$.

(3) The rest are mostly simple cases. They consist of tense sequences present—future, present—past, present—present perfect, past—past perfect, etc. So the relation between ET_{focus} and ET_{new} can be trivially asserted. E.g., from present—future, it follows that $ET_{focus} < ET_{new}$.

Note that after each sentence is processed, the TFS is updated with the *SRE* structure of the new sentence. In most cases, the immediately preceding sentence is in focus when a new one is processed. But it is also possible that some other earlier sentence becomes the current focus. (See [Song and Cohen, 1991] for details of this.)

How accurate is this algorithm? Thanks to its allowance for aspectual classes (boundedness in (1)), it seems to be reasonably accurate in a statistical sense for simple narratives.

However, note that in (1-ii), the case of *non*-unbounded situations, they take time progression as a default. Though often correct, this choice leaves no room for reverse progression of time (recall Webber's TF movement to a preparatory phase of an event). Moreover, for some tense sequences such as *present perfect—present perfect*, this is inappropriate as a default analysis. For example,

"Mary has cleaned the bathroom. John has cleaned the living room."

in no way implies that John cleaned the living room *after* Mary cleaned the bathroom. (This phenomenon of present perfect as *indefinite* past will be discussed in Chapter 8.)

Next, (2-i) seems like a good rule of thumb. As they say, a past tense sentence following a present perfect sentence often elaborates the event described by the present perfect sentence. However, it is not always so, especially if the preceding present perfect sentence is stative, and the subsequent past sentence is telic. For instance, consider

"I have lived in this house for 25 years.
My next door neighbor moved here only last year."

The assumption in (2-ii) is even more dubious, as the following counterexample illustrates.

"Mary is going to leave tonight. She will return next Tuesday."

Finally, one intersting point is that Song uses the (weakened) algorithm in a planning environment. The weakened algorithm is basically the same as the one introduced above except that for the "progression" procedure, he makes no commitment to any ordering between the ETs of two situations [Song, 1991, p. 136]. The actual relation between the two situations can be different depending on which candidate plan it belongs to. When the tense analysis gives an unambiguous temporal analysis, this may be used as a constraint to eliminate inconsistent candidate plans. If there are ambiguities in temporal analysis, he uses the prestored temporal constraints in the candidate plans to help resolve the ambiguities.

Although Song and Cohen make many interesting observations, some of their assumptions and generalizations are in need of refinement or qualification. Making default inferences is fine, but if the algorithm does not provide a way to assess plausibility or a way to recover from an incorrect decision, it may be of very limited utility. As Webber [1987a, p. 150] says, all one can say is that one way of construing the temporal order in a story might be more *plausible* than the others. To get a more accurate algorithm for temporal analysis, one needs to make use of the discourse structure and world knowledge (e.g., about causal relations), etc.[15] Finally, we need to reiterate that Song and Cohen restrict them-

[15]Cf., the discussion of the work of Lascarides *et al.* in Chapter 8.

selves to very simple narratives. All sentences describe actual situations, and embedded clauses and modal, intentional, negated and frequentative contexts are excluded. Also the problem of mapping from English to the semantic representation is left open.

I have discussed some AI approaches to the analysis of tense and aspect. They are similar in that they rely on Reichenbach's tense analysis. As well, they are noncompositional insofar as they do not derive the meanings of "complex tenses" from the meanings of the constituent aspectual auxiliaries, main verbs, and verb inflections. But this is symptomatic of a more fundamental limitation: they perform tense-aspect analysis more or less in isolation, without attempting to integrate it with a detailed theory of logical form, let alone providing a mapping from English syntactic structure to such a logical form. And finally, they avoid many phenomena found in virtually all real texts, such as embedded clauses, intensional contexts, untensed perfects, adverbials, quantifiers, etc.

7.3 Compositionality and Context: Background of Tense Tree Development

We have considered some of the best-known proposals for interpreting tense and aspect. Let us now look back at the two points mentioned at the beginning of Section 7.2. The proposals we considered were more formally explicit about ② than about ①. The early work of Reichenbach and Prior, in particular, had little to say about ①, the mapping from utterance and context to semantic representation. As well, we noted that most AI-motivated work on tense and aspect, adopting a Reichenbachian approach, does not go very far in formalizing ①. Some of the more recent work in linguistic semantics, such as Dowty's and Hinrichs', is aimed at a systematic mapping for ①. Nevertheless, even these formal, compositional accounts tend to take "context" for granted — the focus is on truth conditions, and so the model is simply assumed to supply the needed reference times.

What is lacking in all the above, besides full compositionality, is then a formal notion of context, integrated with a detailed theory of logical form. In particular, the analysis of tense and aspect requires formalization of the notion of temporal context and its relation to logical form. Some recent work that tried to identify reference time systematically was Schubert and Pelletier's [1989]. As tense trees were influenced by this work, I will briefly consider how they tried to locate reference events automatically.

Schubert and Pelletier's Reference Vector

Adhering to a compositional approach, Schubert and Pelletier separate tense and perfect aspect. They further assume that tense always takes wide scope relative to time adverbs. Evaluation is done at a world-time pair $\langle i, w\rangle$, so in that sense, there is only one time index. However, the evaluation of formulas also appeals to the values, in the current interpretation, of certain indexical constants, including a constant "now" and "time reference vector" r. The semantic value of $[r]$ of this reference vector is a sequence of time intervals. These indexically accessed times in effect provide something like the double or triple indices used by Dowty, Hinrichs, and others. The semantics of tense/aspect operators involves a *shift* in the time of evaluation to some other, indefinite (existentially quantified) time, constrained in certain ways relative to the given time of evaluation. For instance, consider

(7.30) a. Mary left
b. PAST(leave(Mary))

In the evaluation of the above (7.30 b) at time i, PAST checks whether i is the current value of *now* and also whether the embedded formula, (leave(Mary)), holds at "some time" j before i. Thus, there is a "backward" shift from i to j.

Their innovation though lies in the way they make use of the time reference vector (and other indexicals) as a formal context-dependent parameter. A formula is evaluated in a context $[\,]$, and evaluation of the formula in turn transforms the context as a side effect. In the time reference vector, one element is distinguished as "being in focus." That is, $[r]$ is of the form of $\langle i_1, \underline{i_2}, \ldots\rangle$, where the underlined element indicates that it is in focus. (The bold symbol **r** is used in the object language to pick out the focal time. Thus, in the preceding example, $[\mathbf{r}] = i_2$.) There is an operation on r which "shifts focus" one element to the right ($\overrightarrow{r}$, forward in time) or one element to the left ($\overleftarrow{r}$, backward in time), relative to **r**.

I now show some sample rules, e.g., rules for *predication, past, perfect,* and two default adverbials, THEREUPON and AT-THAT-TIME. In the following, π is a predicate, α a term, and Φ a formula. $[\,]_{\overleftarrow{r}}$ denotes the valuation function based on the same interpretation as $[\,]$, except that $[r]$ is modified so that the focus is shifted to the left (and similarly so for $[\,]_{\overrightarrow{r}}$).

A. Evaluation rules
1. $[\pi(\alpha)] = [\pi]^{[\alpha]}$
2. $[\text{PAST}(\Phi)]^{i,w} = 1$ iff $i = [now]$ and $[\Phi]^{j,w}_{\overleftarrow{r}} = 1$ for some j before i;
3. $[\text{PERF}(\Phi)]^{i,w} = 1$ iff $i = [\mathbf{r}]$ and $[\Phi]^{j,w}_{\overleftarrow{r}} = 1$ for some j before i;
4. $[\text{THEREUPON}(\Phi)]^{i,w} = 1$ iff $[\Phi]^{i,w} = 1$ and i is immediately or shortly after $[\mathbf{r}]$;
5. $[\text{AT-THAT-TIME}(\Phi)]^{i,w} = 1$ iff $[\Phi]^{i,w} = 1$ and $i =$ end of $[\mathbf{r}]$.

B. Context transformation rules

Notation: $\Phi^{i,w}[\,]$ designates the (new) context which is generated from having already "processed" (i.e., evaluated) Φ in the (old) context $[\,]$, at time i in world w.

1. $\pi(\alpha)^{i,w}[\,] \;=\; [\,]_{r:i}$, if π is monadic, and $[\pi(\alpha)]^{i,w} = 1$
   (i.e., r is assigned the unit vector containing just time i);
   $=\; [\,]$, if π is n-adic, where $n \geq 1$.
2. $(\mathrm{PAST}(\Phi))^{i,w}[\,] \;=\; (\Phi^{i,w}[\,]_{\overleftarrow{r}})_{\overrightarrow{r}}$ for some j before i such that $[\Phi]^{j,w}_{\overleftarrow{r}} = 1$,
   if there is such a j and $i = [now]$;
3. $(\mathrm{PERF}(\Phi))^{i,w}[\,] \;=\; (\Phi^{i,w}[\,]_{\overleftarrow{r}})_{\overrightarrow{r}}$ for some j before i such that $[\Phi]^{j,w}_{\overleftarrow{r}} = 1$,
   if there is such a j and $i = [\mathrm{r}]$;
4. $(A(\Phi))^{i,w}[\,] = [\Phi]^{i,w}[\,]$ for A a time adverb, if $[A\Phi]^{i,w} = 1$.

Note that one needs only consider cases where the new element in focus is the one immediately to the left of the current one in focus (which corresponds to 'shifting backward in time one notch') or where the new element in focus is the one immediately to the right of the current one in focus (which corresponds to 'shifting forward in time one notch"). For illustration, consider the passage (7.31) and their logical forms, one slightly simplified from their example in [1989, p. 260]. I will omit discussing evaluation of the formulas. Instead, I will trace the reference vector, which is the main item of interest here, showing snapshots of it at various points during the process of this mini-passage. In the reference vector, '_' indicates a focal time whose value is as yet undefined.

(7.31) a. John entered the room.

$${}_{\uparrow a}\mathrm{PAST}(\,{}_{\uparrow b}(\text{The } x\text{: room}(x))\ \text{enter}(x, \text{john}))\,{}_{\uparrow c}$$

Reference vector at a: $\langle \underline{now_1} \rangle$
b: $\langle \_\ \ now_1 \rangle$
c: $\langle i_{enter}\ \ \underline{now_2} \rangle$

b. Mary had taken down his painting.

$$\mathrm{PAST}(\,{}_{\uparrow d}\mathrm{PERF}(\,{}_{\uparrow e}(\text{The } y\text{: johns-painting}(y))\ \text{take-down}(y, \text{mary}))\,{}_{\uparrow f}\,)\,{}_{\uparrow g}$$

Reference vector at d: $\langle \underline{i_{enter}}\ \ now_2 \rangle$
e: $\langle \_\ \ i_{enter}\ \ now_2 \rangle$
f: $\langle i_{takedown}\ \ \underline{i_{enter}}\ \ now_2 \rangle$
g: $\langle i_{takedown}\ \ i_{enter}\ \ \underline{now_3} \rangle$

c. He groaned.

$\text{PAST}(\,{}_{\uparrow h}(\text{THEREUPON}(\,{}_{\uparrow i}\,\text{groan(john)})))\,{}_{\uparrow j}$

Reference vector at h: $\langle i_{takedown}\ \ \underline{i_{enter}}\ \ \text{now}_3\rangle$
i: $\langle i_{takedown}\ \ \underline{i_{enter}}\ \ \text{now}_3\rangle$
j: $\langle i_{takedown}\ \ i_{groan}\ \ \underline{\text{now}_4}\rangle$

Note that after processing (7.31 a), i.e., at point c, the reference vector contains i_{enter}. At point d, this i_{enter} gets the focus, and serves automatically as the reference time for the following PERF operator.

However, there is a problem with the reference vector in that PAST and (present) PERF events/times may "overwrite" one another, and similarly PAST (FUTR ...)) can overwrite PRES. For instance, the following kind of passages are not satisfactorily handled.[16]

(7.32) a. The train started to move.

${}_{\uparrow a}\text{PAST}(\,{}_{\uparrow b}(\text{The } x\text{: train}(x))\ \text{start-move}(x))\,{}_{\uparrow c}$

Reference vector at a: $\langle\underline{\text{now}_1}\rangle$
b: $\langle\_\ \ \text{now}_1\rangle$
c: $\langle i_{move}\ \ \underline{\text{now}_2}\rangle$

b. John waved his hand at Mary.

$\text{PAST}(\,{}_{\uparrow d}\,\text{THEREUPON}(\,{}_{\uparrow e}\,\text{wave-hand-at(john, mary)}))\,{}_{\uparrow f}$

Reference vector at d: $\langle\underline{i_{move}}\ \ \text{now}_2\rangle$
e: $\langle\underline{i_{move}}\ \ \text{now}_2\rangle$
f: $\langle i_{wave}\ \ \underline{\text{now}_3}\rangle$

c. He would be back as a lawyer.

$\text{PAST}(\,{}_{\uparrow g}\,\text{AT-THAT-TIME}(\,{}_{\uparrow h}\,\text{FUTR}(\,{}_{\uparrow i}\,\text{back-as-lawyer(john)})))\,{}_{\uparrow j}$

Reference vector at g: $\langle\underline{i_{wave}}\ \ \text{now}_3\rangle$
h: $\langle\underline{i_{wave}}\ \ \text{now}_3\rangle$
i: $\langle i_{wave}\ \ \underline{\text{now}_3}\rangle$
j: $\langle i_{wave}\ \ \underline{i_{back}}\rangle$

[16]Since they do not provide a rule for FUTR, I assumed a rather obvious one in analyzing (7.32 c), i.e.,
$[\text{FUTR}(\Phi)]^{i,w} = 1$ iff $i = [now]$ and $[\Phi]^{j,w}_{\overrightarrow{r}} = 1$ for some j after i;
$(\text{FUTR}(\Phi))^{i,w}[\,] \;=\; (\Phi^{i,w}[\,]_{\overrightarrow{r}})_{\overrightarrow{r}}$ for some j after i such that $[\Phi]^{j,w}_{\overrightarrow{r}} = 1$,
if there is such a j and $i = [now]$.

Note that at point j, now_3 is overwritten by i_{back}, which will in turn be overwritten by now_4 should the story has another sentence following. Thus, we would get something like i_{back} is immediately after now_3, and possibly lose i_{back} altogether from the reference vector. This problem arises because the reference vector is flat, being a linear array of time tokens. However, the reference vector, together with the two default adverbials THEREUPON and AT-THAT-TIME, serves as a crude way of getting orienting relations. This reference vector in fact provided the basis for the tense tree development.

7.4 Summary

I discussed some of the well-known previous work on tense and aspect, in particular, the work by Reichenbach, Prior and Dowty, motivating a compositional and truth-functional approach to tense-aspect interpretation and the need for a context structure. I also discussed several approaches taken by AI researchers, namely, Passonneau's, Webber's, and Song and Cohen's. Then I discussed a work by Schubert and Pelletier that makes use of a context called a reference vector, which motivated the *tense tree* component of our context structure. The next chapter describes the tense tree in detail and present deindexing rules that make use of the tense tree.

Chapter 8

Deindexing with Tense Trees ; Or, A Theory of English Tense-Aspect Semantics

In this chapter, I first introduce the intuitive idea behind the method for deindexing LFs. The method transforms indexical LFs to nonindexical ELFs in a principled way, using tense trees as context structures. Tense trees provide the points of orientation needed to interpret tense and aspect, and are transformed in the course of deindexing LFs. I then provide a formal statement of tense-aspect deindexing rules. Some of the important features of the deindexing mechanism developed in this chapter are (1) that it produces completely deindexed logical forms, thus providing a basis for inference, (2) that it automatically introduces episodic, i.e., situational, variables into the logical form, transforming tense and aspect operators into relations between those episodic variables, (3) that it makes use of well-defined dynamic context structures to locate orienting episodes, and (4) that it is formulated systematically with a small number of recursive equations. In this chapter, I concentrate on rules for deindexing LFs and associated tense tree transformations for declarative speech acts, present and past tense, future modality, and perfect and progressive aspect, and illustrate how they actually work with a couple of examples. Deindexing rules for other operators, i.e., logical operators like conjunctions, negation and quantification, adverbials, and special EL functors like Ka and Ke, are dealt with in the next chapter.

8.1 Tense Trees: A New Type of Context Component

We saw in Chapter 7 that compositionality and context are key issues for correct interpretation of tense and aspect. With this in mind, two goals in this research on tense and aspect processing have been first, to explore the possibility of treating `past`, `pres`, `perf` and `futr` operators as separate sentential operators, and second, to formally characterize their effect on both meaning and context.

One question that arises in formalizing the effect of tense and aspect is how new episodes (or times) are evoked, and what their status is in the resulting logical forms. Are they represented by existentially quantified variables? Free variables or parameters? Constants? The position taken here is that each tense operator, as well as future and perfect operators, introduces a new existentially quantified episode variable. However, given the semantics of existential quantification in EL, these episode variables can subsequently play essentially the role of anaphorically accessible parameters. In addition, various logical operators also introduce episodes, including negation, conjunction, and quantification. This will be discussed in detail in Chapter 9, but the relevant point here is that this will allow us to introduce subepisodes, further analyzing composite episodes, like the ones characterized by (2.16) and (2.17) seen in Chapter 2.

In the deindexing algorithm to be presented in the next section, the preliminary, indexical LF is deindexed by processing it in the current context (recall C, denoting a context in Figure 6.1), and the context is simultaneously transformed. Context-dependent constituents of the LF, such as operators `pres`, `past` and `perf` and adverbs like *today, regularly, twice* or *earlier*, are replaced by explicit relations among "quantified" episodes. The episodes (or times) introduced by the various operators are intuitively not equally "accessible" in the interpretation of a given constituent within a sentence. For instance, in interpreting the verb phrase *have left* in "When Mary arrives, John will *have left*," the episode or time of Mary's arrival seems temporarily accessible, providing a determinate temporal reference for that verb phrase. That episode again seems in "clear view" when we interpret the phrase *be surprised* in a subsequent sentence, "She will *be surprised*." Various examples in Chapter 7 have illustrated these interactions. It therefore appears crucial that the context structure provide access to "the right episodes at the right time" (a dynamic temporal focus) in building a context-independent representation of a text from an indexical LF. Essentially, what is needed is a more adequate version of the "time reference vector," and transformation rules, proposed in [Schubert and Pelletier, 1989].

8.1.1 The Structure of Tense Trees

Before describing tense trees in detail, I should remind the reader of other possible components of a context structure C. First, it has a clock that generates a succession of *Now*-points. In implementation, the system automatically generates the *Now*-term with a unique subscript, Now_i, and assigns a unique system-time and the calendar-time as its properties. In addition, a context structure may in general contain whatever contextual information a more complete discourse theory may call for, e.g., a nested segment structure which records the evolving set of relationships among discourse segments, including the current time, hearer and speaker parameters for those segments, salient temporal and spatial frames, tokens for salient referents other than episodes (possibly in the form of *history lists* with recent referents, or *focus lists* of entities referenced within them), etc. (see [Allen, 1987, Chapter 14]). However, I will not attempt a full specification because I am not addressing all aspects of deindexing as mentioned earlier. The only part of the context structure that is of concern in this thesis is the stack of tense trees, specifically, the *current* (most recent) tense tree.

Tense trees differ from simple lists of Reichenbachian indices in that they organize episode tokens (for described episodes and the utterances themselves) in a way that *echoes the hierarchy of temporal and modal operators* of the sentences and clauses from which the tokens arose. In this respect, they are analogous to larger-scale representations of discourse structure which encode the hierarchic segment structure of discourse. (As will be seen in Chapter 9, the analogy goes further.) Tense trees for successive sentences are "overlaid" in such a way that related episode tokens typically end up as adjacent elements of lists at tree nodes. The traversal of trees and the addition of new tokens is simply and fully determined by the logical forms of the sentences being interpreted.

I now describe tense trees more precisely. The form of a tense tree is illustrated in Figure 8.1. A tense tree node may have up to three branches—a leftward *past* branch, a downward *perfect* branch, and a rightward *future* branch. *Present* does not involve branching; every root node amounts to the present tense. Each node contains a stack-like list of recently introduced episode tokens (or, simply, episodes). As an aid to intuition, the nodes in Figure 8.1 are annotated with simple sentences whose indexical LFs would lead to those nodes in the course of deindexing.

In addition to the three branches, the tree may have (horizontal) *embedding links* to the roots of embedded tense trees. There are two kinds of these embedding links, both illustrated in Figure 8.1. One kind, indicated by dashed lines, is created by subordinating constructions such as VPs with *that*-complement clauses. The other kind, indicated by dotted lines, is derived from the surface speech act (e.g., telling, asking or requesting) implicit in the mood of a sentence.[1] The two kinds of embedding links require slightly

[1] As seen in Chapter 3, the utterances of a speaker (or sentences of a text, etc.) are ultimately represented

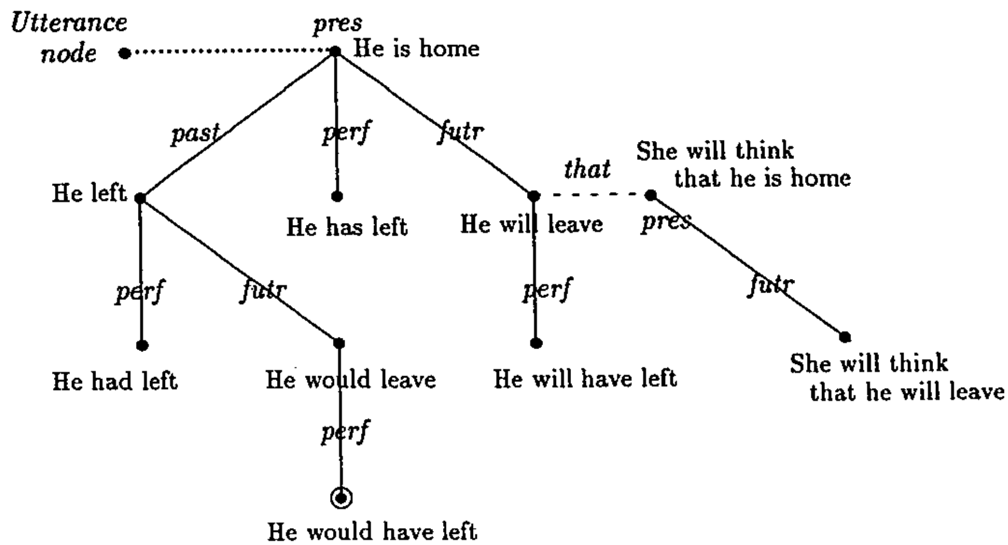

Figure 8.1: A Tense Tree

different tree traversal techniques since one kind is "inserted" by the deindexing of the top-level surface speech act operator (indicating the surface speech act, such as `decl` or `ques`), while the other is directly given in the input sentence. This will be discussed in detail in the next section.

A set of trees connected by embedding links is called a *tense tree structure* (though it is often loosely referred to as a tense tree). This is in effect a tree of tense trees, since a tense tree can be embedded by only one other tree. As an indexical LF is processed in a recursive top-down manner to deindex its tense and aspect operators, adverbials, etc., the tense tree is traversed in a way dictated by the operators encountered. The position of the current traversal is called the *focal node*, or *focus*, of the tense tree structure and is indicated as ◉. At any time, exactly one node of the tense tree structure for a discourse is in focus. Where branches to be traversed do not exist, they are created, and as new episode variables are introduced into the LF, copies of these variables are added to lists maintained at the nodes of the tense tree. Note that the "tense tree" in Figure 8.1 is in fact a tense tree structure, with the lowest node in focus.

---

in terms of modal predications expressing these surface speech acts like [Speaker tell Hearer (That Φ)].

8.1.2 Locating Episode Tokens on Tense Trees

The major function of tense trees is to allow simple, systematic interpretation (by deindexing) of tense, aspect, and time adverbials in texts consisting of arbitrarily complex sentences, and involving implicit temporal reference across clause and sentence boundaries. This includes certain relations implicit in the ordering of clauses and sentences, of the type observed in Section 7.1. Typically, the relation is one of temporal precedence or concurrency, depending on the *aspectual class* or *aktionsart* involved (cf., "John closed his suitcase; He walked to the door" *versus* "John opened the door; Mary was sleeping"). However, in "Mary got in her Ferrari. She bought it with her own money," the usual temporal precedence is reversed (based on world knowledge). Also, other discourse relations could be implied, such as *cause-of, explains, elaborates,* etc. Whatever the relation may be, finding the right pair of episodes involved in such relations is of crucial importance for discourse understanding. As described in Chapter 6, the predicate constant `orients` is used which subsumes all such relations, and `orients` predications can later be used to make probabilistic or default inferences about the temporal or causal relations between the two episodes (more on this later).

By default, an episode added to the right end of a list at a node is "oriented" by the episode which was previously rightmost. For episodes stored at different nodes, one can read off their temporal relations from the tree roughly as follows. At any given moment, for a pair of episodes e and e' that are rightmost at nodes n and n', respectively, where n' is a daughter of n, if the branch connecting the two nodes is a past branch, [e' before e].[2] If it is a perfect branch, [e' impinges-on e], which yields entailments [e' before e] for telic e' (as in "John has left") and [e' until e] for stative e' (as in "John has been working"). If it is a future branch, [e' after e]. If it is an embedding link, [e' at-about e]. These orienting relations and temporal relations are not extracted *post hoc*, but rather are automatically asserted in the course of deindexing using the rules shown later.

As a preliminary example, consider the following passage and a tense tree annotated with episodes derived from it by deindexing rules that will be discussed shortly:

(8.1) a. John picked up the phone.
b. He *had* told Mary that he *would* call her.
c. He nervously began to dial.
d. He hoped she *would* remember him.

[2]Or, sometimes, *same-time* (cf., "John noticed that Mary *looked* pale"). This is not decided in an *ad hoc* manner, but as a result of systematically interpreting the context-charged relation bef_T.

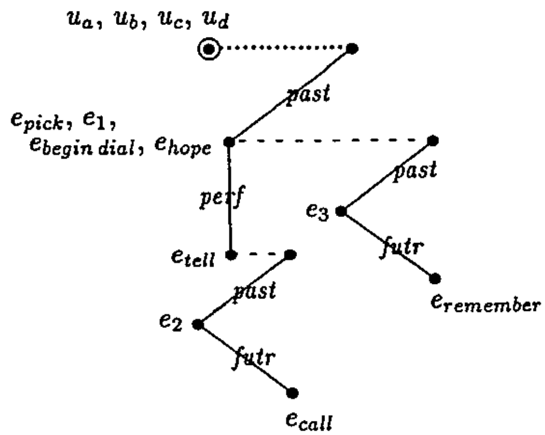

Intuitively, the temporal content of passage (8.1) is that the event of John's *telling*, e_{tell}, took place *before* some time e_1, which is at the same time as the event of John's *picking* up the phone, e_{pick}; and the event of John's *calling*, e_{call}, is located *after* some time e_2, which is the at the same time as the event of John's *telling*, e_{tell}. Similarly, Mary's *remembering* episode is *after* John's *hoping* episode. Thus, e.g., [e_{pick} orients e_1], [e_{tell} before e_1], [e_{call} after e_2], [$e_{begin\,dial}$ orients e_{hope}], and [$e_{remember}$ after e_3]. For the most part, this information can be obtained from the tense tree.

In the above tree, u_a–u_d are utterance episodes for sentences (8.1 a)–(8.1 d) respectively. A crucial observation about this is that the tokens for the "picking" event in (8.1 a) and the "telling Mary" event in (8.1 b) are *not* placed at the same node—whereas the "beginning to dial" event in (8.1 c) does end up at the same node as "picking." If we assume that the "reference episode" (e_1) of perfect episode "telling" can be inferred to be at approximately the same time as the picking event e_{pick} (and it turns out that it can), then the "beginning to dial" event $e_{begin\,dial}$ can be inferred to be shortly after the "picking". As well, note that the "beginning to dial" and "hoping" episodes are adjacent at the *past* node, so that [$e_{begin\,dial}$ orients e_{hope}]. So the collocation of episode tokens in the tense tree is such that one can "read off" the relationships implicit in the tense-aspect operators and surface ordering of sentences. In addition, the deindexing rules yield [e_2 same-time e_{tell}] and [e_3 same-time e_{hope}]. From these, one may infer [e_{tell} before e_{pick}], [e_{call} after e_{tell}], and [$e_{remember}$ after e_{hope}], assuming that the *orients* relation defaults to *same-time* here.

How does [e_{pick} orients e_1] default to [e_{pick} same-time e_1]? In the tense tree, e_1 is an episode evoked by the past tense operator which is part of the meaning of *had* in (8.1 b). It is a *stative* episode, since this past operator logically operates on a sentence of form (perf Φ), and such a sentence describes a *state* in which Φ has occurred—in this instance, a state in which John has told Mary that he will call her. It is this stativity of e_1 which (by default) leads to a *same-time* interpretation of *orients*.[3] Thus, on our account, the tendency of past perfect "reference time" to align itself with a previously introduced past event is just an instance of a general tendency of stative episodes to align themselves with their orienting episode. This is the same tendency noted previously for "John opened the door. Mary was sleeping." Further comments about particularizing the *orients* relation will be provided in the next section.

Earlier, it was mentioned that the relation [e_2 same-time e_{tell}] is obtained directly from the deindexing rules. This will become clear when Past and Futr rules are discussed; here I note only that e_2 is evoked by the past tense component of *would* in (8.1b), and denotes a (possible) *state* in which John will call Mary. Its stativity, and the fact that the subordinate clause in (8.1b) is "past-dominated," causes [e_2 bef$_T$ e_{tell}] to be deindexed to [e_2 same-time e_{tell}]. (A node is *past-dominated* if there is a *past* branch in its ancestry, where embedding links also count as ancestry links. *bef*$_T$ is a context-charged relation that is to be particularized to *before* or *at-about*; see Footnote 2 and Section 8.2.2.)

8.1.3 Traversing Tense Trees: Glimpses of the Deindexing Algorithm

I now briefly discuss how deindexing of LFs may be done. The processing of the (indexical) LF of a new utterance always begins with the root node of the current tense tree (structure) in focus. The processing of the top-level operator immediately pushes a token for the surface speech act onto the episode list of the root node. As we have seen in many examples, a typical indexical LF looks like this:

(8.2) a. John knew that Mary had not left.
b. (*decl* (*past* [John know (*That* (*past* (*perf* (¬ [Mary leave])))]))

Given our compositional approach, the operators decl, past, That, ¬, perf, and other constituents of this LF contribute separately to its meaning. As an LF is recursively transformed, the tense and aspect operators encountered, past, perf, and futr, in particular, cause the focus to shift "downward" along existing branches (or new ones if necessary). That is, processing a past operator shifts the current focus down to the left, creating a new branch if necessary. Similarly perf shifts straight down, and futr shifts down to the right. Each of those operators implicitly introduces exactly one *episode*. This is in contrast with Reichenbach's method, in which every tensed clause always introduces three times R, S, E whether it is a simple past clause like "John laughed" or a composite one like "Mary had not left." Certain operators embed new trees at the current node, (e.g., That), or shift focus to an existing embedded tree, (e.g., decl). At the same time, new episode tokens are added to the lists of tokens at the nodes as rightmost element of its

[3]More accurately, the default interpretation is [(end-of e_{pick}) same-time e_1], in view of examples involving a longer preceding event, such as "John painted a picture. He was pleased with the result."

episode list. As each node comes into focus, its episode list and the lists at certain nodes on the same tree path provide explicit reference episodes in terms of which past, pres, futr, perf, time adverbials, and implicit "orienting" relations are rewritten nonindexically. Eventually the focus returns to the root, and at this point, we have a nonindexical LF, as well as a modified tense tree with new branches and/or new episode tokens. It is worth emphasizing that each of the operators is treated uniformly in deindexing and context change. More specifically, they drive the generation and traversal of tense trees in deindexing.

8.2 Tense-Aspect Deindexing Rules

Suppose that we already have a tense tree structure T, with a particular node in focus, as a result of processing previous inputs and partially processing the current input. (If not, we generate a one-node tree using the "new-tree" function $\triangle$.) Deindexing of an LF relative to a tense tree T is defined by equivalences shown in Table 8.1. Each rule in Table 8.1 consists of two equivalences: one for modifying LFs and the other for the tree transformation (this is achieved as a side effect in the implementation). Each equivalence deindexes the top-level operator, pushing the dependence on context one level deeper into the LF. I first explain the various symbols used in Table 8.1, and then discuss each of the rules.

e_{T} denotes a new episode token uniquely defined as a function of T. (Actually, for a given C, but the rest of C is neglected here.) In short, it is the "next episode variable not yet used in T." For instance, it might be the letter e followed by the least numeric suffix i such that e_i does not occur anywhere in T, but the only point of importance is that it must be a well-defined function on tense tree structures T. Notice the first six rules in Table 8.1, i.e., Decl, Past, Pres, Futr, Fpres and Perf, each "create" a new episode token.

$\mathtt{Now}_{\mathrm{T}}$ denotes the utterance time uniquely defined for the current context, T, i.e., it refers to the speech time for the most recent utterance in T. As discussed in Chapter 2, a Now_i point is typically given properties like

(clock-time-of Now_1) = (interval (tuple 1992 7 22 9 35 46)
(tuple 1992 7 22 9 35 51)).

$\mathtt{Last}_{\mathrm{T}}$ is the last-stored episode variable at the *focal node* of T. So, for a succession of simple past-tensed sentences, each episode generated will orient the next one. The *orients* predications can later be used to make (probabilistic or default) *narrative inferences* about the temporal or causal relations between the two episodes. (More on this later.)

`Decl:` $(\text{decl } \Phi)_T \leftrightarrow (\exists e_T{:}[[e_T \text{ same-time } \texttt{Now}_T] \wedge [\texttt{Last}_T \text{ immediately-precedes } e_T]]$
$[[\texttt{Speaker} \text{ tell } \texttt{Hearer} \; (\text{That } \Phi_{\hookrightarrow \circ T})] ** e_T])$

Tree transformation: $(\text{decl } \Phi) \cdot T = \leftarrow (\Phi \cdot (\hookrightarrow \circ T))$

`Past:` $(\text{past } \Phi)_T \leftrightarrow (\exists e_T{:}[[e_T \text{ bef}_T \; \texttt{Emb}_T] \wedge [\texttt{Last}_{\swarrow T} \text{ orients } e_T]] \; [\Phi_{\circ \swarrow T} ** e_T])$

Tree transformation: $(\text{past } \Phi) \cdot T = \uparrow(\Phi \cdot (\circ \swarrow T))$

`Pres:` $(\text{pres } \Phi)_T \leftrightarrow (\exists e_T{:}[[e_T \text{ at-about } \texttt{Emb}_T] \wedge [\texttt{Last}_T \text{ orients } e_T]] \; [\Phi_{\circ T} ** e_T])$

Tree transformation: $(\text{pres } \Phi) \cdot T = (\Phi \cdot (\circ T))$

`Futr:` $(\text{futr } \Phi)_T \leftrightarrow (\exists e_T{:}[[e_T \text{ after } \texttt{Last}_T] \wedge [\texttt{Last}_{\searrow T} \text{ orients } e_T]] \; [\Phi_{\circ \searrow T} ** e_T])$

Tree transformation: $(\text{futr } \Phi) \cdot T = \uparrow(\Phi \cdot (\circ \searrow T))$

`Fpres:` $(\text{fpres } \Phi)_T \leftrightarrow (\exists e_T{:}[[e_T \text{ after } \texttt{Emb}_T] \wedge [\texttt{Last}_{\searrow T} \text{ orients } e_T]] \; [\Phi_{\circ \searrow T} ** e_T])$

Tree transformation: $(\text{fpres } \Phi) \cdot T = \uparrow(\Phi \cdot (\circ \searrow T))$

`Perf:` $(\text{perf } \Phi)_T \leftrightarrow (\exists e_T{:}[[e_T \text{ impinges-on } \texttt{Last}_T] \wedge [\texttt{Last}_{\downarrow T} \text{ orients}_T \; e_T]] \; [\Phi_{\circ \downarrow T} ** e_T])$

Tree transformation: $(\text{perf } \Phi) \cdot T = \uparrow(\Phi \cdot (\circ \downarrow T))$

`Prog:` $(\text{prog } \Phi)_T \leftrightarrow (\text{prog } \Phi_T)$

Tree transformation: $(\text{prog } \Phi) \cdot T = \Phi \cdot T$

`That:` $(\text{That } \Phi)_T \leftrightarrow (\text{That } \Phi_{\mapsto T})$

Tree transformation: $(\text{That } \Phi) \cdot T = \leftarrow (\Phi \cdot (\mapsto T))$

`Pred:` For π an atomic predicate and $\tau_1, \ldots, \tau_n$ terms,
where τ_i, $1 \leq i \leq n$, is atomic, except possibly for τ_{n-1}:

$[\tau_n \; \pi \; \tau_1 \; \tau_2 \; \ldots \; \tau_{n-1}]_T \leftrightarrow [\tau_n \; \pi \; \tau_1 \; \tau_2 \; \ldots \; \tau_{n-1_T}]$

Tree transformation: $[\tau_n \; \pi \; \tau_1 \; \ldots \; \tau_{n-1}] \cdot T = \tau_{n-1} \cdot T$

Table 8.1: Basic Temporal Deindexing Rules

$\mathtt{Emb_T}$ denotes the last-added episode at the node which directly *embeds* the tree containing the focal node of T. If there is no embedding node, $\mathtt{Emb_T}$ denotes what $\mathtt{Now_T}$ denotes. $\mathtt{Emb_T}$ is usually an episode corresponding to a surface speech act or attitude verb.

$\circ$T denotes the tense tree which is just like T except that the new token e_T has been added to the focal node.

$\swarrow$T, $\downarrow$T and $\searrow$T: '$\swarrow$T' is T with the focus displaced to the left (i.e., `past`) daughter, with creation of a new daughter if necessary. Similarly, '$\downarrow$T' and '$\searrow$T' signify T with the focus displaced to the strictly downward (i.e., `perf`) daughter, and to the right (i.e., `futr`) daughter, respectively, with creation of a new daughter as necessary.

$\mapsto$ and $\hookrightarrow$: These indicate focus shifts to the root of an embedded tree. '$\hookrightarrow$' indicates *retraversal* of the *last* embedding link added at the current focus. If no embedding link exists at the current focus, it creates a new tree, embeds it at the focal node, and shifts focus to the root of the embedded tree. '$\mapsto$' adds a *new* embedding link at the current focal node, and shifts focus to the root of the newly embedded. Note that '$\mapsto$' *always* adds a new embedding link, whereas '$\hookrightarrow$' does so only when there exists no embedding link at the current focal node. The dotted link in Figure 8.1 corresponds to a '$\hookrightarrow$' link, and the dashed link corresponds to a '$\mapsto$' link. The reason behind this distinction will be explained shortly when the `That` rule is discussed.

$\uparrow$ and $\leftarrow$: These symbols indicate focus shifts to a parent node and an embedding node respectively. '$\uparrow$' signals upward displacement of the focus to the mother, and $\leftarrow$ indicates focus restoration to the embedding node. These restore the focus to its original position, assuming that recursive processing of operand Φ has no net effect on the focus (which it doesn't, thanks to the way deindexing rules work). Note that the function composition in, for example, "$\uparrow(\Phi \cdot (\circ\swarrow T))$" is read "from the inside to the outside," as usual.

In the tree transformations, the dot operator, '$\cdot$', denotes the transformation yielding a modified tense tree from an indexical formula (its left operand) and an initial tense tree (its right operand). That is, it symbolizes the tree-structure transformation function,

$$\cdot : \text{LF-expressions} \times \text{Tense-tree structures} \rightarrow \text{Tense-tree structures}$$

In essence, $(\Phi \cdot T)$ yields the tree one would expect if the shift and store operations specified in the deindexing rules were carried out on a *global* data structure (and the focus reset to the node originally focused in T).

Let us now look at each of the deindexing rules in Table 8.1 in more detail.

8.2.1 The Decl-Rule for Surface Speech Acts

$$\texttt{Decl}: \quad (\text{decl } \Phi)_{\text{T}} \leftrightarrow (\exists e_{\text{T}}{:}[[e_{\text{T}} \text{ same-time } \texttt{Now}_{\text{T}}] \wedge [\texttt{Last}_{\text{T}} \text{ immediately-precedes } e_{\text{T}}]] \; [[\text{Speaker tell Hearer (That } \Phi_{\hookrightarrow \circ \text{T}})] ** e_{\text{T}}])$$

$$\textit{Tree transformation}: (\text{decl } \Phi) \cdot \text{T} = \leftarrow (\Phi \cdot (\hookrightarrow \circ \text{T}))$$

As discussed in Chapter 6, the top sentential level LF (input to the deindexer) is augmented with surface speech act operators such as decl for declarative sentences, ques for interrogative sentences, excl for exclamatory sentences, etc., so that the utterances of a speaker (or sentences of a text, etc.) are ultimately to be represented in terms of modal predications expressing these surface speech acts, such as [Speaker tell Hearer (That Φ)] or [Speaker ask Hearer (YN-q Φ)].[4] Although these speech acts are not explicitly part of what the *speaker* uttered, they are part of what the *hearer* gathers from an utterance. Speaker and Hearer in the deindexed formula are to be replaced by constants (the speaker and hearer parameters of the utterance) obtained from the context structure C (exclusive of T). Thus, deindexing decl brings in the (preliminary) surface speech act of type *tell* into the LF. (Rules for other kinds of surface speech act operators are omitted in this thesis.)

The Decl-rule "creates" a new episode token e_{T}, which it asserts is exactly at the time of utterance $\texttt{Now}_{\text{T}}$ (which amounts to Reichenbach's speech time), shifts focus to the root of the embedded tree (if there is no embedded tree, it first embeds a new one), and states that the formula [Speaker tell Hearer (That Φ)], after recursive deindexing of Φ (with e_{T} now "stored" (as indicated by '∘') at the embedding tree, and focus shifted to the embedded tree), *characterizes* the new utterance episode.

Note that 'That' in the RHS of the rule is a "facsimile" of the sentence nominalization operator That derived explicitly from the text (as in, e.g., "John thinks *that* Mary is pretty"). That is, the former is semantically indistinguishable from the latter, but treated slightly differently in tense tree maintenance since it was in effect "inserted" by the hearer rather than directly given in the input sentence.

In the tree transformation, after the deindexing of Φ is completed recursively, focus is restored to the embedding node as indicated by ←.

---

[4]In [Hwang and Schubert, 1992c], we used [Speaker ask Hearer (Whether Φ)] instead of [Speaker ask Hearer (YN-q Φ)], taking Whether as a nominalization operator that abbreviates Answer-to (YN-q ...). However, we found it somewhat confusing as the English word "whether" is ambiguous between a *question* and an *answer to a question*. For instance, consider

(1) a. John asked Mary *whether she had submitted her assignment.*
   b. *Whether there exist creatures on Mars* is a good question.

(2) a. John knows *whether Mary has submitted her assignment.*
   b. It's hard to tell *whether there exist creatures on Mars.*

In (1), the *whether*-clause means a *question*, not an *answer* to it. So, unambiguous YN-q is now used instead of ambiguous Whether.

8.2.2 The Past-Rule and More on the Orients Relation

Past: $(\text{past } \Phi)_T \leftrightarrow (\exists e_T{:}[[e_T \text{ bef}_T \text{ Emb}_T] \wedge [\text{Last}_{\swarrow T} \text{ orients } e_T]] \; [\Phi_{\circ\swarrow T} \, ** \, e_T])$
Tree transformation: $(\text{past } \Phi) \cdot T = \uparrow(\Phi \cdot (\circ\swarrow T))$

The Past-rule works like this: (i) it creates a new episode token e_T; (ii) predicates it to be before-or-at-about (i.e., 'bef_T') the embedding event, which would be either a surface speech act episode or an attitude/modal episode; (iii) shifts focus to left daughter ('$\swarrow$'); (iv) designates the last added episode at the current (i.e., new) focus, '$\text{Last}_{\swarrow T}$', as the point of orientation for the new episode e_T; (v) drops the new episode token e_T at the focus ('$\circ$'); and then (vi) states that the formula Φ on which past operates, after recursive deindexing, *characterizes* the new episode e_T ('$** \; e_T$'). In the tree transformation equation, '$\uparrow$' restores the focus to its original position, assuming that recursive processing of Φ has no net effect on the focus (which it doesn't, thanks to the way the remaining rules work).

Two things in the Past-rule require further comment: first, the ambiguous interpretation of past as indicated by 'bef_T', and second, the orients relation. The 'bef_T' relation is both indexical (dependent on T) and "context-charged," with different probable consequences, either *before* or *same-time* (or, *at-about*), depending on the aspectual class of its first argument and whether the focal node of T is past-dominated. (As explained already, a node is "past-dominated" if there is a *past* branch in its ancestry.) If the current focal node of T is *not* past-dominated, bef_T is deindexed to before (a noncontext-charged relation). If the focal node *is* past-dominated, however, bef_T is deindexed to at-or-before, where [e at-or-before e'] strongly "suggests" [e same-time e'] for *stative* e, and [e before e'] for *telic e*. (Progressive and perfect VPs and negations are normally considered stative.) The following examples illustrate this point.

(8.3) Mary *was* tired

(8.4) John will realize Mary *told* him a lie

(8.5) John thinks Mary *was* in love with him

(8.6) John *realized* Mary *was* watching him

(8.7) John *realized* Mary *took* his pen

In (8.3) and (8.6)–(8.7), the top-level "being tired," "realizing" and "thinking" episodes are *before* the embedding utterance episodes (since they are not past-dominated). Also, in (8.4)–(8.5), the *non*past-dominated "telling" or "being in love" episode is *before* its embedding "realizing" or "thinking" episode. But the past-dominated "watching," "being in love," and "taking" episodes in (8.6)–(8.7) could be either *at the same time as* or *before* their embedding episodes, depending on the aspectual class of the embedded episode and other factors.

In this case, Mary's watching John (which is stative) is likely to be at the *same time* as John's realizing, whereas Mary's taking John's pen (which is telic) is likely to be *before* John's realizing it. (In fact, the "past" tense of *took* in (8.7) may be considered as "lazy past perfect," i.e., a reduction of *had taken*. Such a reduction of past perfect to past for telic episodes is frequently observed in embedded clauses, especially in American English.) Note, however, that these are only plausible inferences. Sometimes the relationship is indeed ambiguous, or an *earlier*-episode reading is strongly preferred even if the embedded episode is stative, or a *concurrent*-event reading is strongly preferred even for telic embedded episodes. Consider, for instance, variants like "John noticed that Mary winked at him" that forces a concurrent-event reading and "John remembered that Mary was in love with Jack in high school" that forces an earlier-episode reading.

Because of the `orients` relation introduced by the `Past`-rule, for a succession of simple past-tensed sentences, each episode generated will orient the next one. (A new episode variable is generated for the orienting predication when no orienting episode is found in the tense tree.) This `orients` relation and certain others derived from context, e.g., `has-preferred-antecedents` discussed in Chapter 6, are context-charged, i.e., their meaning is "discharged" using *uncertain* (plausible) inference. The `orients` relation is essentially an indicator that there could be a more specific discourse relation between the argument episodes. For instance, the fact that [e orients e'] may *suggest*, among other possibilities, immediate temporal succession, a subepisode relation, a causal or explanatory relation, or any of the discourse relations that have been discussed in the literature (e.g., [Hobbs, 1979]). I will discuss this issue in Section 8.3.

8.2.3 The `Pres`-Rule

`Pres`: $(\text{pres}\ \Phi)_{\text{T}} \leftrightarrow (\exists e_{\text{T}}{:}[[e_{\text{T}}\ \text{at-about}\ \texttt{Emb}_{\text{T}}] \wedge [\texttt{Last}_{\text{T}}\ \text{orients}\ e_{\text{T}}]]\ [\Phi_{\circ\text{T}} ** e_{\text{T}}])$
Tree transformation: $(\text{pres}\ \Phi)\cdot\text{T} = (\Phi\cdot(\circ\text{T}))$

The `Pres`-rule works much like `Past`, except that there is no focus shift involved. That is, it "creates" a new episode token e_{T}, which it predicates to be *at about* the time of embedding event (e.g., the utterance of the sentence being processed or a modal embedding episode), designates the last-stored episode at the current focus as the point of orientation for the new episode, and states that the formula Φ on which *pres* operates, after recursive deindexing (with e_{T} now stored at the focus), *characterizes* the new episode.

Note that since the new episode e_{T} is predicated to be at about $\texttt{Emb}_{\text{T}}$, rather than at about $\texttt{Now}_{\text{T}}$, it does not make the mistake of equating "present" with the "speech time." Thus, in the example below, Mary's loving episode in (8.8) and the reference episode of the perfect in (8.9) are interpreted as at the same time as the embedding, John's thinking episode (rather than at the same time as the speech time).[5]

(8.8) John will think Mary *loves* him

(8.9) John will think Mary *has* left him

(Recall that in the Past-rule, as well, "past" did not necessarily mean "before speech time.") In general, thanks to the $\mathtt{Emb_T}$ in Pres and Past rules, we get correct interpretations for embedded sentences. For instance, consider the following sentences.[6]

(8.10) a. One day you will regret that you *treated* me like this
b. One day you will regret that you *were* treating me like this
c. One day you will regret that you *have* treated me like this
d. ?One day you will regret that you *treat* me like this
e. ?One day you will regret that you *are treating* me like this

(8.11) a. Smith will claim on the witness stand that he *was* in Mexico
b. ??Smith will claim on the witness stand that he *is* in Mexico

Suppose the sentences in (8.10) are talking about the "treating" episode that is taking place right now. In (a)–(b), this treating episode is described using past tense, which is fine in our mechanism as past only means before the embedding "regretting" episode. Similarly in (c), the present perfect only indicates that the embedded "treating" episode is some time before the "regretting" episode, rather than before the speech time. As for (d), our mechanism would analyze it such that the "treating" episode is taking place at the time as the embedding "regretting" episode, which is not quite correct. However, notice that (d) is felicitous only if *treats* is construed habitually, and the habitual "treating" episode extends far enough into the future to reach the point of John's "regretting" episode. Similar observations are applicable to the sentences in (8.11). "Being in Mexico" in (8.11 b) certainly refers to the time of Smith's "claiming," i.e., the meaning of (8.11 b) is distinct from that of (8.11 a). Thus, if we set aside habitual/generic readings, our deindexing rules for Past and Pres seem correct, agreeing with our intuitions about English.

Sentence (8.10 e) shows an interesting phenomenon though. Here, the subordinate clause seems ambiguous, i.e., it allows both a habitual reading and a nonhabitual one. In the latter reading, it clearly refers to the utterance time rather than the time of the embedding "claiming" episode. This seems to be a similar phenomenon to that of "past

[5]Incidentally, the Pres-rule will interpret "John thought Mary *loves* him" the same as "John thought Mary *loved* him." I think this is in fact what the speaker meant though one may get the implication that John would still think that Mary still loves him. Since *loving* is a stative episode, this kind of plausible inference is always allowed in EL.

[6]Sentences in (8.10) are slightly modified ones from Parsons' sentence discussed in [Dowty, 1982, p. 50], and those in (8.11) are Dowty's [1982, p. 51].

as lazy past perfect" as discussed in Section 8.2.2, i.e., a case of using present progressive instead of the more grammatically correct past or past progressive tense. However, the usage can't really be called *lazy* as there is no reduction involved in (8.10e) relative to (8.10a) or (8.10b) (if anything, the opposite). Intuitively, it seems more likely that a "perspective shift" is involved here—the hearer is abandoning the "futuric perspective" induced by the first three words, "You will regret ... ," and shifting the perspective back to the present in midsentence. The perspective shift here may be abetted by the progressive form of the subordinate clause, since the progressive is often used to emphasize that an event *is* in progress, bringing the hearer closer to the "scene." Also, the amount of amendment required could be a matter of degree depending on the nature of the verb in the embedding sentence, e.g., people *regret* only prior (i.e., past) events.[7] This is just a speculation, however. I will have more to say about perspective shifts in Chapter 9.

8.2.4 The Futr-Rule

Futr: $(\text{futr } \Phi)_T \leftrightarrow (\exists e_T\colon [[e_T \text{ after } \mathtt{Last}_T] \wedge [\mathtt{Last}_{\searrow T} \text{ orients } e_T]]\ [\Phi_{\circ\searrow T} ** e_T])$
Tree transformation: $(\text{futr } \Phi) \cdot \text{T} = \uparrow(\Phi \cdot (\circ\searrow \text{T}))$

Before discussing the rule, it needs to be emphasized that on our conception of tense there are only two tenses in English: past and present. Note that these alone are manifested in English as verb inflections. *Will* and *would*, in their future and future-in-the-past readings, are regarded as consisting logically of tense plus the futr modal operator.

To understand the Futr-rule, recall the following lexical rules for *will* and *would*, introduced in Chapter 6:

V[aux, pers, numb, pres, _VP[base]] ← *will*; λPλx<pres (futr [x P])>
V[aux, pers, numb, past, _VP[base]] ← *would*; λPλx<past (futr [x P])>

Note that the futr operator is encountered *only after* its implicit tense operator has been processed, so that a characterizing ('**') relation will already embed the (futr Φ) expression. Here, $\mathtt{Last}_T$ will be the episode introduced by the embedding past or pres tense operator.[8]

Thus, the effect of the Futr rule is quite analogous to that of Pres and Past, except that the temporal location of the new episode e_T is specified relative to the episode $\mathtt{Last}_T$ characterized by "having Φ true in its future," rather than relative to an "embedding

[7]But it's not clear; e.g, "I'll report that you are treating me like this" does not sound very bad (especially if the intended reporting is imminent). Also, *intention* and futr seem to behave in a very similar way; cf., ??Smith *intends* to claim on the witness stand that he is in Mexico.

[8]Technically, however, it is possible for $\mathtt{Last}_T$ to be a future or perfect episode; for example, in "When Mary eventually comes back around next month, John *will be going to* move to Chicago with his new girlfriend in a few days."

episode" $\texttt{Emb}_\mathrm{T}$. Note that, **futr** is interpreted *relatively* in our mechanism, which is one of the important differences from Reichenbach's analysis.

One may also note that according to the Futr-rule, (futr $\Phi)_\mathrm{T}$ specifies a *stative* type of episode. That is, if a Φ-episode lies in the future of some reference episode (the one designated as $\texttt{Last}_\mathrm{T}$ here), it also lies in the future of all subepisodes of that reference episode (inward persistence). On the other hand, (futr $\Phi)_\mathrm{T}$ is not in general atemporal (unlocated), since there may be a final Φ-episode in a world, and no episodes after it have a Φ-episode in their future. (Of course, if Φ is atemporal, so is (futr $\Phi)_\mathrm{T}$.)

Before moving on to the perfect operator, I briefly discuss the Fpres-rule.

Fpres: (fpres $\Phi)_\mathrm{T} \leftrightarrow (\exists e_\mathrm{T}{:}[[e_\mathrm{T}$ after $\texttt{Emb}_\mathrm{T}] \wedge [\texttt{Last}_{\searrow\mathrm{T}}$ orients $e_\mathrm{T}]]\ [\Phi_{\circ\searrow\mathrm{T}} ** e_\mathrm{T}])$
Tree transformation: (fpres $\Phi) \cdot \mathrm{T} = \uparrow(\Phi \cdot (\circ\searrow \mathrm{T}))$

This tentative rule is to handle "future-oriented present" — the present tense used to refer to the future, as in (8.12) below.

(8.12) a. John leaves tomorrow
b. (decl (fpres ((adv-e (during *Tomorrow*)) [John leave])))
c. $(\exists e_1{:}[e_1$ same-time *Now1*]
[[Speaker tell Hearer (That
$(\exists e_2{:}[e_2$ after $e_1]$ [[[e_2 during *Tomorrow*] $\wedge$ [John leave]] $** e_2]))]$
$** e_1])$

I showed a rough deindexed logical form above, but will omit discussing the logical form as adverbial deindexing has not been covered yet. The Fpres rule assumes, though, that the disambiguation of the present tense between **pres** and **fpres** can be done relatively easily, which I think is the case, but I will not pursue the matter here.

8.2.5 A Pragmatically Ambiguous Perf-Rule

Perf: (perf $\Phi)_\mathrm{T} \leftrightarrow (\exists e_\mathrm{T}{:}[[e_\mathrm{T}$ impinges-on $\texttt{Last}_\mathrm{T}] \wedge [\texttt{Last}_{\downarrow\mathrm{T}}$ orients $e_\mathrm{T}]]\ [\Phi_{\circ\downarrow\mathrm{T}} ** e_\mathrm{T}])$
Tree transformation: (perf $\Phi) \cdot \mathrm{T} = \uparrow(\Phi \cdot (\circ\downarrow \mathrm{T}))$

Perfect may occur either in present tense ("John has left today"), in past tense ("John had left yesterday"), or untensed ("John will have left by tomorrow morning," "John is believed to have left yesterday" or " 'Tis better to have loved and lost than never to have loved at all"). Thus, the Perf rule resembles the Futr rule in that the **perf** operator is normally encountered either after an implicit tense or future modality operator or a kind forming operator has been processed, which would cause the perf-clause to be already embedded within '**' at the point when it is deindexed. Recall that lexical rules for perfect auxiliary *have* are much like those for the future modal auxiliary *will*, i.e.,

$V[aux, 3per, sing, pres, \_VP[-en]] \leftarrow$ *has* ; $\lambda P \lambda x$<pres (perf [x P])>
$V[aux, pers, numb, past, \_VP[-en]] \leftarrow$ *had* ; $\lambda P \lambda x$<past (perf [x P])>
$V[aux, pers, numb, \text{-ing}, \_VP[-en]] \leftarrow$ *having* ; $\lambda P \lambda x$(perf [x P])

Consequently, the deindexing of perf works much the same way as that of futr, except that the focus shift is "straight downward" rather than rightward. So I omit describing the behavior of the Perf-rule. Note that like futr-formulas, perf-formulas are always stative.

In the Perf-rule, Last_T is analogous to the Reichenbachian reference time R for the perfect. But there is one important difference: Last_T is a new episode evoked by the tense or kind-forming operator that embeds the perf-formula (cf., (*pres* (perf Φ)) or (*Ke* (perf Φ))), rather than one supplied by the immediately prior context. It "aligns" itself with such a prior episode, because of its stativity (in the default "discharging" of the orients relation), but is not *identical* with it.

The impinges-on relation confines its argument e_T (the situation or event described by the sentential operand of perf) to the temporal region preceding the argument Last_T. It is a context-charged relation like orients, and its more specific import depends on the aspectual types of its arguments. Specifically, if e_T is an (unbounded) open (i.e., extensible) episode, impinges-on entails that the state or process involved *persists* to the reference episode, i.e., [e_T until Last_T]. If e_T is a (bounded) closed (i.e., inextensible) episode, impinges-on entails that it occurred sometime *before* the reference time, i.e., [e_T before Last_T], and by default its result state persists to the reference time. That is, unbounded-stative "goings-on" extend all the way to the reference time while telic or bounded-stative ones need not. The (default) persistence of a result state to the reference time can be written as

For Ψ a formula,
$(\exists e{:}\,[[e \text{ episode}] \wedge [(\text{result-type } (\text{sit-type } e)) \neq (\text{sit-type } e)]$
$[(\text{result-type } (\text{sit-type } e)) = (\text{Ke } \Psi)]]$
$(\exists e'{:}\,[[e' \text{ episode}] \wedge [e \text{ impinges-on } e']]$
$\rightarrow_{.8, e, e'} (\exists t{:}\,[[t \text{ time}] \wedge [(\text{end-of } e) \text{ subep-of } t] \wedge [t \text{ until } e']]\ [\Psi * t])$.

As an example, consider the following (neglecting orienting relations).

(8.13) a. John has been sleeping
b. (pres (perf (prog [John sleep])))
c. ... [(perf (prog [John sleep])) ** e_5] ...
d. ... [($\exists e_6$: [e_6 *impinges-on* e_5] [(prog [John sleep]) ** e_6]) ** e_5] ...
e. ... [($\exists e_6$: [e_6 *until* e_5] [(prog [John sleep]) ** e_6]) ** e_5] ...

(8.14) a. John has woken up

b. (pres (perf [John wake-up]))

c. ... [(perf [John wake-up]) ** e_{11}] ...

d. ... [($\exists e_{12}$: [e_{12} *impinges-on* e_{11}] [[John wake-up] ** e_{12}]) ** e_{11}] ...

e. ... [($\exists e_{12}$: [e_{12} *before* e_{11}] [[John wake-up] ** e_{12}]) ** e_{11}] ...

The **pres** in the (b)-formulas above introduces e_5 and e_{11} as shown in the (c)-subformulas. Then **perf** in the (c)-subformulas introduces e_6 and e_{12}, asserting [e_6 impinges-on e_5] and [e_{12} impinges-on e_{11}], as seen in the (d)-subformulas. Notice that e_5 and e_{11} are the reference episodes of the perfect. Next, by ampliative inference, we have in (8.13e) [e_6 until e_5] (since e_6, John's sleeping, is an open, extensible—unbounded-stative—episode). That is, John has been sleeping *until now*. This is the desired inference. On the other hand, in (8.14e), we have [e_{12} before e_{11}], as John's waking up episode, e_{12}, is a closed one. Since further (result-type (Ke [John wake-up])) = (Ke [John awake]), this "suggests" there is a time t which has (end-of e_{12}) as subepisode and lasts until e_{11} such that [[John awake] * t], and in particular

[[John awake] * (end-of e_{12})] and [[John awake] * (begin-of e_{11})].

Note that this is *only* an implication, since we can perfectly well say, without contradiction, "John has woken up but fallen asleep again."

For another example, consider sentence pairs like

(8.15) John has become well

(8.16) John has been ill

In (8.15), the episode described by the *have*-complement, John's becoming well, is a closed episode, and hence *precedes* the present. In (8.16), however, the corresponding episode of John's being ill extends to the (speaker's) present in the preferred reading. In the non-preferred reading of (8.16), John's episode of being ill strictly *precedes* the present. Such a reading may be obtained by allowing a "lexical extension rule" to be applied to lexical verbs marked as unbounded-statives, transforming them into (less readily available) bounded-stative verbs. As discussed in Chapter 6, the corresponding semantic transformation applies an operator **bounded** to the logical translation, yielding a bounded-stative predicate as in (past (perf [John (*bounded ill*)])).

Note that the alignment of reference episodes of perfect with contextually supplied episodes is automatically achieved in the tense tree mechanism. Consider, for instance, the following:

(8.17) John inferred that Mary *had left*.

(8.18) a. John *entered* the room.
b. Mary *had* taken down his paintings,
c. and *had* hung up Schwarzenegger posters.
d. He *groaned*.

In (8.17), the (past-dominated) "past" of *had* will generate an episode in the past relative to the embedding, "inferring" episode. Since perfect reference episodes are stative (because the state of being *after* some given type of event can persist indefinitely, and holds of the subintervals of any intervals of which it holds), the past episode of *had* (serving as reference episode for the perfect) is interpreted *at-about* the time of the embedding "inferring" episode. This amounts to picking the "inferring" episode as the reference episode of the embedded perfect.

In (8.18b), the (wide-scope) "past" of *had* generates an episode in the past relative to the time of speech and oriented by the "entering" episode in (8.18a). It is this past "entering" episode with which the reference episode for the perfect aligns itself. As already mentioned during discussion of the `Past` rule, the key to the seemingly anaphoric character of the perfect is this: if [e_1 orients e_2], and e_2 is *stative*, then there is a strong suggestion that e_2 is either concurrent with e_1 (namely, when e_1 is stative as well) or contiguous with its end point (when e_1 is telic). Since the "entering" episode is telic and the perfect reference episode is stative, the orients relation between them is interpreted as *right after*, i.e., the perfect reference episode is contiguous with John's entering. This is tantamount to making the (end) time of John's entering the "reference time" for the perfect. So we get the desired "anaphoric" effect, *without* treating past as anaphoric, and without singling out the past in past perfect for special treatment. Perhaps this is the most important point about the `Perf`-rule — that the reference episode is introduced simply by the normal effect of operators (usually `past`, `pres` or `futr`) "exterior" to it.

The `Perf` rule also solves a number of problems in the interaction of perfect aspect with tense, adverbials, and the aspectual class of the complement (without resorting to separate methods for the various forms of tensed and untensed perfect, perfect progressive, etc., as is often done in computational linguistics).

Finally, I should mention that certain well-known problems involving temporal adverbials in perfect sentences, such as the inadmissibility of "*John has left yesterday" is not accounted by the above `Perf` rule. This and some other issues related to perfect are now discussed.

Idiomatic Present Perfect?

English present perfect is unique: it describes an episode that extends up to the speech time or that has taken place in some indefinite past, but it does not allow the actual event time to be specified. For instance, (8.19) and (8.20) below are fine, but (8.21) isn't.

(8.19) Mary has been here this morning

(8.20) Mary was here at 11:30 this morning

(8.21) *Mary has been here at 11:30 this morning

The reason (8.21) sounds odd is that it has the event time specified, whereas in (8.19) only the "time frame" containing the event time is specified. As mentioned in Chapter 7, this time frame is also called "extended now," or, XN for short [McCoard, 1978, Chapter 4]. As a rule, definite time adverbials in present perfect VPs modify this XN, not the actual events. (Indefinite time adverbials can freely modify the event time. Consider, for example,

> Mary has jogged *at dawn* twice this month.
> Mary has visited Paris *in spring*.

The indefinite time adverbials, i.e., generic adverbials, *at dawn* and *in spring* above modify event times, rather than the XN.) Under that view, *this morning* in (8.19) refers to the time frame XN that contains Mary's "(bounded) being here" episode and lasts until the speech time. In (8.21), *at 11:30 this morning* is *prior* to the speech time and also *coincides* with the XN that extends to the speech time, leading to a contradiction. Note, however, that this is an issue only for bounded sentences because, for unbounded ones, the episode time fills up the XN, and so there is no distinction between the event/episode time and the XN.

A phenomenon that is closely related to this XN theory, as well as to the theory of perfect as *indefinite past*, is that episodes described in present perfect do not involve strict orienting relations. This is in comparison with simple tense which evokes episodes easily accessed anaphorically, especially in a sequence of past tensed sentences. The following examples illustrate this point.

(8.22) a. The guests *began* to arrive. Mary *set up* drinks on the table.
b. Mary *set up* drinks on the table. The guests *began* to arrive.

(8.23) a. The guests *have* begun to arrive. Mary *has* set up drinks on the table.
b. Mary *has* set up drinks on the table. The guests *have* begun to arrive.

In (8.22), where sentences are in past tense, reversing the sentence order has the effect of also reversing the implied sequence of events; thus, (a) and (b) are different stories. But in (8.23), where sentences are in present perfect, sequence (a) and sequence (b) do not differ significantly.[9] This phenomenon of present perfect has been observed by many including Leech [1987, p. 41] and ter Meulen [1991].

[9] In contrast, sentences in past perfect may have orienting relations as seen in (7.2), Chapter 7.

Splitting the Perf*-Rule?*

The Perf-rule shown earlier is not capable of accounting for some of the phenomena discussed. More specifically, it has the following problems: first, it asserts an orients relation regardless of whether there exists one; second, it admits sentences like "*Mary has left yesterday"; and third, it does not properly analyze telic sentences involving adverbials that modify XN, e.g., "Mary has visited London once *this year*." In fact, earlier in this research, the perfect was assumed to be ambiguous— that is, the present perfect auxiliary, *has* or *have*, was logically translated as perf1, but when occurring untensed or in combination with past, the perfect is ambiguous, i.e., logically either perf1 or perf2. The main difference between them is that perf1 is sensitive to the bounded/unbounded distinction, while perf2 is not. The possibility of having two kinds of perfect operators remains open although its relative complexity limits its appeal. The two versions could be written in the following sort of way (this is very tentative).[10]

Perf(1): $(\text{perf}_1\ \Phi)_T \leftrightarrow (\exists e_T{:}[e_T \text{ until } \mathtt{Last}_T][\Phi_{O\downarrow T} ** e_T])$, for Φ unbounded
$(\text{perf}_1\ \Phi)_T \leftrightarrow (\exists e_T{:}[e_T \text{ until } \mathtt{Last}_T][\Phi_{O\downarrow T} * e_T])$, for Φ bounded
Tree transformation: $(\text{perf}_1\ \Phi)\cdot T = \uparrow(\Phi \cdot (O\downarrow T))$

Perf(2): $(\text{perf}_2\ \Phi)_T \leftrightarrow (\exists e_T{:}[[e_T \text{ before } \mathtt{Last}_T] \wedge [\mathtt{Last}_{\downarrow T} \text{ orients } e_T]]\ [\Phi_{O\downarrow T} ** e_T])$
Tree transformation: same as Perf(1)

In the Perf(1)-rule, the episode e_T temporally coincides with the XN time frame.[11] When Φ is unbounded, this e_T is the actual episode of type Φ; for bounded Φ, however, the actual event of interest is a temporal part of e_T (notice '$*$'). In the latter case, the actual event is not evoked, let alone being stored in the tense tree. Note that in the Perf(1)-rule, e_T is not in the orients relation. The second variant Perf(2) is simpler, amounting to a "relative past" or a "past-in-past" reading, as in "John realized that Mary *had* left." Note the orients relation in Perf(2).

The distinction between perf1 and perf2 helps to account for the following contrasts:

(8.24) a. *John has left yesterday
b. John had left the day before
c. John will have left the day before

[10]New lexical rules for the perfect auxiliary *have* will then be like the following:
V[aux, 3per, sing, pres, _VP[-en]] ← *has*; λPλx<pres (perf₁ [x P])>
V[aux, pers, numb, past, _VP[-en]] ← *had*; λPλx<past (perf₁ [x P])>
V[aux, pers, numb, past, _VP[-en]] ← *had*; λPλx<past (perf₂ [x P])>

[11]Note that I use the term XN in a rather broad sense, i.e., it indicates a time frame that extends up to the reference episode (as opposed to just the speech time).

On the proposed account, only perf1 is available in (8.24 a); so, we do not get "John's leaving" episode stored in the tense tree, but only an episode *containing* John's leaving and lasting till now — but that cannot possibly be contained in *yesterday*. In (8.24 b) and (8.24 c), no difficulties are encountered because of the availability of the perf2 reading.

Unfortunately, the new rules introduce new problems. A few things that are lost, besides simplicity, are first, the default persistence of result states; second, analysis of sentences involving event-modifying adverbials such as "Mary has visited Paris *in spring*"; third, analysis of some complex sentences, e.g., ones involving conjunction in the sentential operand of perf, such as "The president has *arrived* and *seen* the scene already." In this sentence, individual subepisodes and their relation to each other (e.g., "arriving" episode orients "seeing" episode) cannot be expressed by Perf(1).

These shortcomings could be remedied with a further complication. For instance, for the first problem, we may introduce $e_{T'}$ for the actual event and assert an impinges-on relation on the RHS of the bounded version of Perf(1). The second and third problems may be solved by letting bounded Φ introduce its own episode $e_{T'}$ and store them in an embedded tense tree. However, the interaction between perfect and time adverbials will remain problematic, and sentences like "Mary has jogged *at dawn* twice *this week*" seem to call for yet another set of features. Clearly further thought is required, but I will leave this to future research.

8.2.6 The Prog-Rule

Prog: $(\text{prog } \Phi)_T \leftrightarrow (\text{prog } \Phi_T)$
Tree transformation: $(\text{prog } \Phi) \cdot T = \Phi \cdot T$

The deindexing rule for the prog operator is trivial; it just passes the tense tree T inward, bypassing the prog operator itself. There is no tree transformation involved.

One point of interest is that the formula (prog Φ) is always considered unbounded-stative, while the sentential operand of prog, Φ, may be either telic or stative. That is, Φ could be an activity as in "Mary was playing the piano" or "Mary was watching TV," or a habitual/generic as in "When I first met her, Mary was smoking/running a mile every afternoon." Or, Φ could be an achievement or accomplishment; e.g., as in "Mary was baking a cake," "Mary was climbing the mountain," etc. Thus, a progressive form of a telic formula is ambiguous in that the formula may be interpreted repetitively or as an atomic episode which is "in progress." For instance, "Mary was blinking" could mean she was blinking repeatedly or she was in the middle of one "blink." (For the latter interpretation, consider "Mary's eyes look closed on the photo because she was blinking when the flash went off.")

Some tentative meaning postulates concerning statives and progressives are

MP 8.1. For Φ stative or factual,

$$(\forall e: [\Phi ** e]\ (\forall t: [[t \text{ time}] \wedge [t \text{ during } e]]\ (\exists e': [e' \text{ same-time } t]\ [\Phi ** e'])))$$

MP 8.2. For Φ telic,

$$(\forall e: [\Phi ** e]\ (\forall t: [[t \text{ time}] \wedge [t \text{ proper-during } e]]\ (\exists e': [e' \text{ same-time } t]\ [(\text{prog } \Phi) ** e'])))$$

For instance, MP 8.2 allows us to infer from [[John blink] ** E5] and [E6 proper-during E5] and from [[John eat Dinner] ** E7] and [E8 proper-during E7],

$$(\exists e_1: [e_1 \text{ same-time E6}]\ [(\text{prog [John blink]}) ** e_1]) \quad \text{and}$$
$$(\exists e_2: [e_2 \text{ same-time E8}]\ [(\text{prog [John eat Dinner]}) ** e_2]),$$

respectively. Incidentally, formulas like [(prog Φ) ** e] involves a caveat; one should take into consideration the granularity associated with Φ. That is, [(prog [John eat Dinner]) ** e] does imply that John is engaged in an "eating" activity at every single instant during e; he must need a short break between each bite!

8.2.7 The That-Rule for Embedded Sentences

That: $(\text{That } \Phi)_{\text{T}} \leftrightarrow (\text{That } \Phi_{\mapsto \text{T}})$

Tree transformation: $(\text{That } \Phi) \cdot \text{T} = \leftarrow (\Phi \cdot (\mapsto \text{T}))$

The deindexing rule for *That*-nominalization adds a *new* embedding link to the root of a *new* embedded tree, shifting the focus to the root node of the new tree, as indicated by the $\mapsto$. This is in contrast to $\hookrightarrow$ introduced by surface speech act rules, like `Decl` introduced earlier, which create a new embedding link *only if* no embedding link exists at the current focus. As with `Decl`, focus is restored to the embedding node after its argument Φ is deindexed.

Intuitively, the distinction between two kinds of embedding, one for a sentence nominalization derived explicitly from the text and the other introduced through an implicit speech act, is motivated by the following sort of contrast:

(8.25) a. Mary finished the homework.
b. She ordered a pizza.

(8.26) a. Mary said that she finished the homework.
b. She also said that she ordered a pizza.

In (8.26), the embedded sentences are objects of attitudes, and it is much less clear than in (8.25) whether they refer to "successive" episodes. Note that the LFs for (8.25a,b) will have speech act predicates `decl`, deindexing of which will let the embedded "finishing" and "ordering" episodes stored at the same node in the same embedded tree as the embedding link will be *re-traversable*. But in (8.26 a,b), although the top-level "saying" episodes are stored in the same node, the embedded episodes are stored in different embedded trees. Thus it will be *harder* to establish a connection between them (i.e., tense structure alone will not establish a connection though inference based on other information still might).

8.2.8 The Pred Predication Rule

Pred: For π an atomic predicate and $\tau_1, \ldots, \tau_n$ terms, where τ_i, $1 \leq i \leq n$, is atomic, except possibly for τ_{n-1}:
$[\tau_n\ \pi\ \tau_1\ \tau_2\ \ldots\ \tau_{n-1}]_{\mathrm{T}} \leftrightarrow [\tau_n\ \pi\ \tau_1\ \tau_2\ \ldots\ \tau_{n-1_{\mathrm{T}}}]$
Tree transformation: $[\tau_n\ \pi\ \tau_1\ \ldots\ \tau_{n-1}] \cdot \mathrm{T} = \tau_{n-1} \cdot \mathrm{T}$

The last rule discussed in this chapter is the Pred-rule for deindexing atomic predications. The rule shown above is a slightly simplified version which assumes the predicate π and all of its arguments — except possibly the second last — are atomic. The second last argument τ_{n-1} could be an embedded *That*-nominal or a kind nominal (especially of Ka- or Ke-types).

However, in view of sentences with complex arguments at other positions, e.g., "That Mary is intelligent is well known" or "That John took someone else's money was even worse than that he failed the exam," or sentences with non-atomic predicates, e.g., $\lambda x[[x$ scarce$] \wedge [x$ expensive$]]$, a more realistic rule would be

$$[\tau_n\ \pi\ \tau_1\ \tau_2\ \ldots\ \tau_{n-1}]_{\mathrm{T}} \leftrightarrow [\tau_{n_{\mathrm{T}}}\ \pi_{\tau_n \cdot \mathrm{T}}\ \tau_{1_{\mathrm{T1}}}\ \tau_{2_{\mathrm{T2}}}\ \ldots\ \tau_{n-1_{\mathrm{T(n-1)}}}],$$

where $\mathrm{T1} = \pi \cdot (\tau_n \cdot \mathrm{T})$, $\mathrm{T2} = \tau_1 \cdot \mathrm{T1}$, $\mathrm{T3} = \tau_2 \cdot \mathrm{T2}$, $\ldots$, $\mathrm{T(n-1)} = \tau_{n-2} \cdot \mathrm{T(n-2)}$, with a tree transformation rule,

$$[\tau_n\ \pi\ \tau_1\ \tau_2\ \ldots\ \tau_{(n-1)}] \cdot \mathrm{T} = \tau_{n-1} \cdot (\cdots (\tau_2 \cdot (\tau_1 \cdot (\pi \cdot (\tau_n \cdot \mathrm{T}))))).$$

8.3 Particularizing the ORIENTS Relation

We have seen that various deindexing rules introduce the orients relation into the formula. I also mentioned that the orients relation is essentially an indicator that there could be a more specific discourse relation between the argument episodes; i.e., $[e$ orients $e']$ suggests, among other possibilities, immediate temporal succession, a subepisode relation, a causal or explanatory relation, or any of the discourse relations. Existing proposals for getting these discourse relations right appear to be of two kinds. The first uses the aspectual classes of the predicates involved to decide on discourse relations, especially temporal ones, e.g., [Partee, 1984], [Dowty, 1986] and [Hinrichs, 1986]. The second approach emphasizes inference based on world knowledge and discourse/rhetorical structure, e.g., [Hobbs, 1985a] and [Lascarides and Asher, 1991; Lascarides and Oberlander, 1992]. The work by Lascarides *et al.* is particularly interesting in that it makes use of a default logic and is capable of retracting previously inferred discourse relations. I now discuss this issue in some detail.

Our approach aims to fully combine the use of aspectual class information and world knowledge. In the earlier example, "Mary got in her Ferrari. She bought it with her own money," I noted that the default interpretation of `orients` is reversed by world knowledge: one owns things *after* buying them, rather than before. But sometimes world knowledge is mute on the connection. For instance, in "John raised his arm. A great gust of wind shook the trees," there seems to be no world knowledge supporting temporal adjacency or a causal connection. Yet we tend to infer both, perhaps attributing magical powers to John (precisely because of the lack of support for a causal connection by world knowledge). So in this case default conclusions based on *orients* seem decisive. In particular, we would assume that if e and e' are telic episodes, where e is the performance of a volitional action and e' is not, then [e orients e'] suggests [e right-before e'] and (less firmly) [e cause-of e']. Thus, what is needed to handle this kind of case is a plausible inference based on aspectual classes of the sentences involved. The approach to plausible inference in EL in general, and to such default inferences in particular, is probabilistic. The hope is that we will be able to "weigh the evidence" for or against alternative discourse relations (as particularizations of `orients`). Though details on this process have not been settled, the following kind of discourse axioms may be used.

Causal Connection.

$$(\exists e_1\!: [e_1 \text{ event}]\ (\exists x\!: [[x \mid e_1] \text{ action}]\ (\exists e_2\!: [[e_2 \text{ event}] \wedge [e_1 \text{ orients } e_2]] \\ \neg(\exists y\ [[y \mid e_2] \text{ volitional-action}])))) \\ \rightarrow_{.6, e_1, e_2} [e_1 \text{ immed-cause-of } e_2]$$

where [e immed-cause-of e'] means that [e cause-of e'] and [e right-before e'], and the .6 is the lower bound on the probability of the conclusion (see Chapter 4). This axiom allows us to infer from "John greeted Mary. She was startled" that John's greeting caused Mary's being startled, with minimum degree of belief .6. The probabilities could be modified as new evidence comes in, supporting or disconfirming the discourse relations inferred.

Another point that needs to be mentioned is that we would eventually like to have other kinds of entities on the tense tree as well, not just the episodic variables derived during the tense and aspect deindexing process. In particular, event nominals seem to be prime candidates for placement on tense trees, such as the *accident* nominal in (8.27) below. Note that it is the *accident* that orients the reference episode of the next past perfect sentence (i.e., the drinking episode was prior to the *accident*).

(8.27) Mary is angry about the *accident*. The other driver *had* been drinking.

It is of interest to briefly describe the work of Lascarides *et al.* [Lascarides and Asher, 1991; Lascarides and Oberlander, 1992; Lascarides *et al.*, 1992] at this point, since it can quite naturally be viewed as "picking up where this thesis leaves off" in the analysis of intersentential temporal relations in narratives. Here, we have been satisfied to determine which episodes are connected by ("context-charged") `orients` relations and which ones are not. The main problem we addressed was getting these relations right for arbitrary clausal embedding. In essence, our technique is to "pile up" episodes connected by `orients` relations at particular tense tree nodes. We left to further research the problem of specifying how `orients` relations become particularized to specific temporal, causal, or other relations, as a result of plausible inference based on world knowledge. (But see the discussion of interpretation by "implicit question answering" in Chapter 11. This constitutes a sketch of an alternative approach to that of Lascarides *et al.*)

By contrast, Lascarides *et al.* put aside the problem of arbitrarily nested clauses, and instead set out to determine *specific* intersentential temporal and discourse relations (rather than something like `orients` relations) among sequences of simple narrative clauses. To the extent that they succeed, their techniques could probably be used to establish temporal and discourse relations *among the episodes "piled up" at a tense tree node* (i.e., those linked into chains through `orients` relations).

More specifically, they are interested in explaining why a sequence of sentences with syntactically same structure may imply different temporal ordering, as illustrated by

(8.28) a. Max fell.
b. John pushed him.

(8.29) a. John pushed.
b. Max fell.

In both (8.28) and (8.29), a telic past sentence is followed by a telic past sentence. However, people understand (8.28) as $E_a > E_b$, but $E_a < E_b$ in (8.29), where "<" and ">" mean "before" and "after" respectively.

The following discussion is based largely on [Lascarides and Asher, 1991].[12] Their central assumption is that particular temporal/causal relations are closely correlated with particular discourse relations. Thus, if the discourse relations are known, temporal/causal relations are easily deducible (and vice versa). They mention the following sort of correspondences between discourse relations and temporal/causal relations. (The clause α appears in the text before β, and E_α denotes the main event described by α.)

[12]The recency of the 1992 papers precludes their full evaluation in the context of this thesis (whose emphasis in any case is on representation).

`Discourse Relations`	`Temporal/Causal Relations`
Narration(α, β), α, β events	E_β after E_α
Elaboration(α, β), α, β events	E_β part-of prep. phase of E_α
Explanation(α, β), α, β events	E_β cause-of E_α
Background(α, β), α event, β state	E_α and E_β overlap
Result(α, β), α event, β event or state	E_α cause-of E_β

This reduces the problem of finding implicit temporal/causal relations to finding the right discourse relations. For this, they use both world knowledge and linguistic knowledge. Here are sample *laws* encoding such knowledge.

- If clauses α and β are `discourse-related`, and α describes an event e of x falling, and β describes an event e' of y pushing x, then normally e' is the cause of e. [defeasible causal law]
- If clauses α and β are `discourse-related`, and α describes an event e of x switching off the light, and β describes an event e' of room y becoming dark, then normally e is the cause of e'. [defeasible causal law]

From the causal connections established by these laws, Lascarides *et al.* in turn infer the *Explanation* discourse relation for (8.28). (See the discourse rule for Explanation below.) Here saying that α, β are `discourse-related` is similar to saying [E_α `orients` E_β] in EL terminology. (But there *are* differences: `discourse-related` clauses may be only indirectly linked by a sequence of `orients` relations; and clauses linked directly by `orients` need not be discourse-related.) As these laws show, most causal laws represent a mixture of linguistic knowledge (re "discourse-related") and world knowledge (the "physics" of pushing, light switches, etc.). A possible criticism here is that these "laws" seem much too specific to be plausible as explicitly stored knowledge. Rather, one feels that the two rules should really be instances of a more general schema, something like this:

> If clauses α and β are `discourse-related`, and the *type* of event described by α, after abstraction of specific arguments, is known to be capable of causing the *type* of event described by β, then normally (the specific event) E_α is the cause of (the specific event) E_β.

For instance, as applied to (8.29), this warrants conclusion [$E_{JohnPushedMax}$ cause-of $E_{MaxFell}$], since an event of type E_{xPushy} is known to be capable of causing an event of type E_{yFall}. (Perhaps we would also want to involve the known *types* of x and y here — i.e., they seem to be *persons*.) The advantage of such a schematic approach would be that we would no longer be "mixing" discourse knowledge with specific world knowledge. For instance, the causal connection between pushing and falling, or between switching off a room light and the room getting dark, could now be stated independently of any discourse notions.

Lascarides *et al.* also show the following "maxim of interpretation."

Narration. If α and β are discourse-related, Narration(α, β) holds unless there is information to the contrary. [defeasible law]

Axiom for Narration. If Narration(α, β) holds, and α and β describe the events e and e' respectively, then e occurs before e'. [nondefeasible law]

The defeasible law *Narration* tries to capture Gricean-style pragmatic maxims. They also provide rules for various discourse and temporal/causal relations such as the following. (I use Disc-rel(α, β) to abbreviate "α and β are discourse-related.")

1. If Disc-rel(α, β) and E_β is a preparatory event for E_α, then Elaboration(E_α, E_β) holds.
2. If Disc-rel(α, β) and E_β is a state, then E_α and E_β overlap.
3. If Disc-rel(α, β), α event, β state, and E_α and E_β overlap, then Background(α, β) holds.
4. If Disc-rel(α, β) and E_β is cause of E_α, then Explanation(α, β) holds.
5. If Disc-rel(α, β) and E_α is cause of E_β, then Result(α, β) holds.
6. If Disc-rel(α, β) and E_β is *not* prep for E_α, then Elaboration(α, β) does not hold.
7. If Elaboration(α, β) holds, but Elaboration(α, γ) does *not* hold, then Narration(β, γ) does not hold.

Note that rule 5 together with the previously cited causal laws allows inference of the *Result* discourse relation (*provided* that α and β are known to be discourse-related). Similarly, laws concerning conditions for "preparatory events" feed into rules 1 and 6. Apparently, more than one discourse relations (possibly conflicting with each other) can hold between two sentences. Sometimes a particular choice among the alternatives is dictated by known causal, part-of or temporal relations; if not, then Narration(α, β) holds. A nonmonotonic logic called MASH is used for this reasoning.

For the inference process to get under way, it still remains to figure out which pairs of sentences can be discourse-related, i.e., what the possible attachment sites for a sentence are. After all, the cited laws and maxims all *presuppose* knowledge about which sentences are discourse-related. For instance, consider

(8.30) a. Guy experienced a lovely evening last night.
b. He had a fantastic meal.
c. He ate salmon.
d. He devoured lots of cheese.
e. He won a dancing competition.

Now with what sentence is **e** discourse-related? I.e., what is the attachment site for **e**? To provide a sufficiently strong basis for inference, Lascarides *et al.* propose an additional set of conditions, called *constraints* on discourse structure. E.g., the following is a (simplified) definition for possible attachment sites (i.e., "open" sites).

> Let R be Explanation or Elaboration; Then the current sentence can be discourse-related only to (i) the previous sentence α, (ii) a sentence β such that $R(\beta, \alpha)$, or (iii) a sentence γ such that $R(\gamma, \beta)$ and $R(\beta, \alpha)$.

According to this constraint, only **d**, **b**, **a** are open for attachment of **e**. But intuitively, **d** cannot be related to **e**, as dancing is not normally part of a meal. The following law, whose form is much like that of the earlier (causal law) examples, expresses this intuition:

- If clauses α and β are discourse-related, and α describes an event e of x having a meal, and β describes an event e' of x win a dance competition, normally e' is *not* preparatory phase of e.

Through discourse rules 6 and 7, their nonmonotonic logic is then able to rule out **d** as an attachment site. This leaves **a** and **b** as the only attachment sites, i.e., **e** must be discourse-related to **a** or **b**.

In summary, this appears to be a promising approach to the determination of particular temporal and discourse relations in narratives, especially if the overly specific laws can be replaced by more general ones, backed up by world knowledge and inference mechanisms which can "weigh the evidence."[13] However, since Lascarides *et al.* do not provide a detailed theory of logical form computation, their proposals cannot yet be fully evaluated from the perspective of this thesis, or exploited in our implementation.

8.4 Tense Trees and Deindexing Rules at Work: Examples

I now illustrate (more fully than in Section 8.1.3) how the deindexing mechanism works. I also show how tense trees are modified as discourse is processed, in particular, how episode tokens are stored at appropriate nodes of the tense tree, and how deindexed LFs, with `orients` and temporal ordering relations incorporated into them, are obtained. As each node comes into focus, its episode list and the lists at certain nodes on the same tree path provide explicit reference episodes in terms of which `decl`, `past` and `perf`, and implicit "orienting" relations are rewritten nonindexically. Eventually the focus will return to the root, and at this point, we have a nonindexical LF, as well as a modified tense tree. I will show two traces.

[13] In personal communication, Jon Oberlander has mentioned that probabilistic approaches may eventually be needed.

Example 1

Deindexing involves application of a fixed set of equivalences, namely, the *deindexing rules* just seen. Let us assume (6.8), repeated below as (8.31), was uttered right after the sentence "Pluto was walking slowly," and the tense tree component of the context after processing it was like T shown below. The task here is to deindex (8.32), the indexical logical form for (8.31), using a context structure containing T as its current tense tree.

(8.31) John realized that Pluto was tired.

(8.32) (decl (past [John realize (That (past [Pluto tired]))]))$_{\mathrm{T}}$

T = 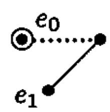

In T, e_0 corresponds to the added speech act — i.e., the speaker's utterance of the preceding sentence, and e_1 denotes the described episode, "Pluto was walking slowly." An indexical formula Φ in combination with a context structure C is normally written as Φ_C, but since all but the current tree T is ignored here, it is written instead as Φ_T.

First, we need to deindex the topmost operator `decl`. An application of `Decl`-rule to (8.32), relative to T, gives us (8.33) shown below. Notice that the `Decl` rule introduces into the formula the new utterance event, e2, and transforms the tense tree T into T1, moving the tree annotation inward toward the *past*-clause.

(8.33) (∃e2: [[e2 same-time Now2] ∧ [e0 immed-precedes e2]]
[[**Speaker** tell **Hearer** (That
(past [John realize (That (past [Pluto tired]))])$_{\mathrm{T1}}$)]
** e2])

T1 = 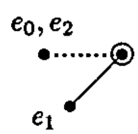

Here, the underlined part is the subformula that still needs to be deindexed (in the modified context T1). **Speaker** and **Hearer** are to be replaced by the speaker and the hearer parameter of the context C. Next, the **Past**-rule is applied to (8.33), resulting in the modified LF (8.34) and the tense tree T2.

(8.34) (∃e2: [[e2 same-time Now2] ∧ [e0 immed-precedes e2]]
[[**Speaker** tell **Hearer** (That
(∃e3: [[e3 before e2] ∧ [e1 orients e3]]

[[John realize (That (past [Pluto tired]))]$_{T2}$
** e3]))]
** e2])

T2 = 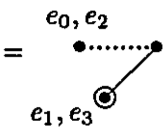

Note that Past introduces e3 for John's "realizing" event, and asserts that it is *before* the utterance episode e2 (since the current focal node is *not* past-dominated) and oriented by e1. As mentioned earlier, the orients relation is assumed to have certain "default" consequences, dependent on the aspectual classes of the episodes they relate. Since e1 (Pluto's walking slowly) is stative and e3 (John's realizing) is telic, the inference from [e1 orients e3] is that John's realization was *during* Pluto's walking slowly, i.e., [e3 during e1].[14] (As well, a causal relation, [e1 cause-of e3], can be tentatively inferred.)

Next, the Pred-rule and the That-rule are applied consecutively, with the resulting LF (8.35) and the tense tree T3. Deindexing of the atomic, nonindexical argument *John* and the atomic predicate *realize* by the Pred-rule is trivial, but (That (past [Pluto tired])) needs further, recursive application of deindexing rules. First the That-rule generates a new tree and embeds it at the current focal node with a *non* re-traversable link, and then shifts the focus to the root node of the newly embedded tree. In T3, a *re-traversable* embedding link is indicated with '·····', and a *non* re-traversable one is indicated with '- - -'.

(8.35) (∃e2: [[e2 same-time Now2] ∧ [e0 immed-precedes e2]]
[[Speaker tell Hearer (That
(∃e3: [[e3 before e2] ∧ [e1 orients e3]]
[[John realize (That (past [Pluto tired])$_{T3}$)]
** e3]))]
** e2])

T3 = 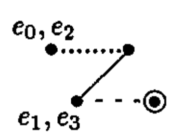

[14]Bonnie Webber has pointed out that "realize" is ambiguous between a telic reading and a stative reading (e.g., "I realize you have a Volvo"). A stative reading is especially easily available in present tense sentences. I set aside this predicate disambiguation problem here, and assume that we are concerned with deindexing the telic reading of (8.31). (A verb like "inferred," in place of "realized," would have avoided the ambiguity.)

It then remains to process the innermost tensed clause (past [Pluto tired])$_{T3}$. One more application of the Past-rule to (8.35), but keeping in mind that we are now at a past-dominated node, converts (8.35) to:

(8.36) (∃e2: [[e2 same-time Now2] ∧ [e0 immed-precedes e2]]
[[Speaker tell Hearer (That
(∃e3: [[e3 before e2] ∧ [e1 orients e3]]
[[John realize (That
(∃e4: [e4 at-or-before e3] [<u>[Pluto tired]</u>$_{T4}$ ** e4]))]
** e3]))]
** e2])

T4 =
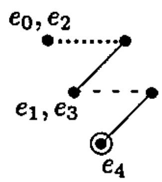

Assume the orienting predication is omitted when there is no orienting episode. Note that though e_4 (Pluto's being tired) is predicated to end during e_3 (John's realizing it), it may well be an initial segment of a much longer episode of the same type; i.e., Pluto may continue to be tired. (On the other hand, if the embedded episode were *telic*, as in "Mary noticed that John winked at her," it would indeed be prevented from extending beyond the embedding episode.)

It remains to apply the Pred-rule once more, to the underlined subformula of (8.36), and to interpret *Pluto* relative to the tense tree T4 as just *Pluto* (a nonindexical constant), with overall result shown below. This completes the deindexing process. The final, nonindexical ELF is shown in (8.37), with the final tree structure as shown in T5. Note that the focus has been shifted to the root node by recursively shifting it back to the mother node.

(8.37) (∃e2: [[e2 same-time Now2] ∧ [e0 immed-precedes e2]]
[[Speaker tell Hearer (That
(∃e3: [[e3 before e2] ∧ [e1 orients e3]]
[[John realize (That
(∃e4: [e4 at-or-before e3] [[Pluto tired] ** e4]))]
** e3]))]
** e2])

T5 = 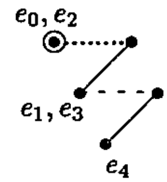

Thus we have a fully context-independent representation of (8.32), which can be used freely for ampliative inference. First, the context-charged `at-or-before` relation has certain default consequences dependent on the aspectual classes of the episodes they relate. Since e_4 is stative (given its characterization [Pluto tired]), the ampliative inference from the context-charged predication [e4 at-or-before e3] is that e_4 is *concurrent* with e_3 (i.e., the same time as John's realization), in the absence of contrary information. Also the context-charged relation [e_1 orients e_3] will lead to the tentative inference that John's realization was *during* Pluto's walking slowly, in view of the fact that e_1 is stative and e_3 telic. Again, a causal relation can also be tentatively inferred. The results of deindexing thus seem to be in complete accord with intuition.

What is important here is that episodes evoked by successive sentences, or by embedded clauses within the same sentence, are correctly connected to each other. In particular, note that the orienting relation between the episode e_3 of John's realizing that Pluto was tired and the episode e_1 of Pluto's walking slowly is automatically incorporated into the above deindexed formula. (Also, we could plausibly *particularize* this orienting relation to [e_3 during e_1], based on the aspectual class of these episodes.) Thus we have established inter-clausal connections automatically, which in other approaches require heuristic discourse processing. This was a primary motivation for tense trees.

Example 2

We will see another example, this time without tracing the deindexing process in detail. Consider sentences (8.38 a)–(8.40 a).

(8.38) a. John went to the hospital.

b. (decl $\uparrow_a$ (past $\uparrow_b$ [John goto Hospital])) $\uparrow_c$

c. (∃ e1: [e1 same-time *Now1*]
  [[`Speaker` tell `Hearer` (That
    (∃ e2: [e2 before e1] [[John goto Hospital] ** e2]))]
  ** e1])

(8.39) a. The doctor told John he had broken his ankle.

b. (decl $\uparrow_d$ (past $\uparrow_e$ [Doctor tell John (That $\uparrow_f$
(past $\uparrow_g$ (perf $\uparrow_h$ [John break Ankle]))))])) $\uparrow_i$

c. (∃ e3: [[e3 same-time *Now2*] ∧ [e1 immed-precedes e3]]
[[Speaker tell Hearer (That
(∃ e4: [[e4 before e3] ∧ [e2 *orients* e4]]
[[Doctor tell John (That
(∃ e5: [e5 at-or-before e4]
[(∃ e6: [e6 before e5] [[John break Ankle] ** e6])
** e5]))]
** e4]))]
** e3])

(8.40) a. He gave him a crutch.

b. (decl $\uparrow_j$ (past $\uparrow_k$ [Doctor give John Crutch])) $\uparrow_l$

c. (∃ e7: [[e7 same-time *Now3*] ∧ [e3 immed-precedes e7]]
[[Speaker tell Hearer (That
(∃ e8: [[e8 before e7] ∧ [e4 *orients* e8]]
[[Doctor give John Crutch] ** e8]))]
** e9])

The LFs before deindexing are shown in (8.38 b)–(8.40 b), where the labelled arrows mark points I will refer to; the final, context-independent ELFs are in (8.38 c) and (8.40 c). The transformation from (b)-formulas to (c)-formulas and the corresponding tense tree transformations are done with the deindexing rules shown earlier. Anaphoric processing is presupposed here.

The snapshots of the tense tree while processing (8.38 b) and (8.40 b), at points $\uparrow_a$–$\uparrow_l$, are shown in Figure 8.2 (with a null initial context).

A couple of comments are in order. First, note that when the second past operator in (8.39 b) is processed, the current focal node is past-dominated as shown in the snapshot at point g. So, the Past-rule asserts a context-charged relation [e_5 at-or-before e_4] as indicated in (8.39 c). This relation is then further particularized into [e_5 at-about e_4] by ampliative inference since e_5, being perfect, is stative. Second, though the resultant tree happens to be unary, additional branches would be added by further text, e.g., a future branch by "It will take several weeks to heal."

What should be noted here is that *Reichenbach-like relations are introduced compositionally*. That is, [e6 before e5], i.e., the event e_6 of John's breaking of the ankle is *before* the state e_4 John is in at the time of the doctor's talking to him. In addition, our recursive

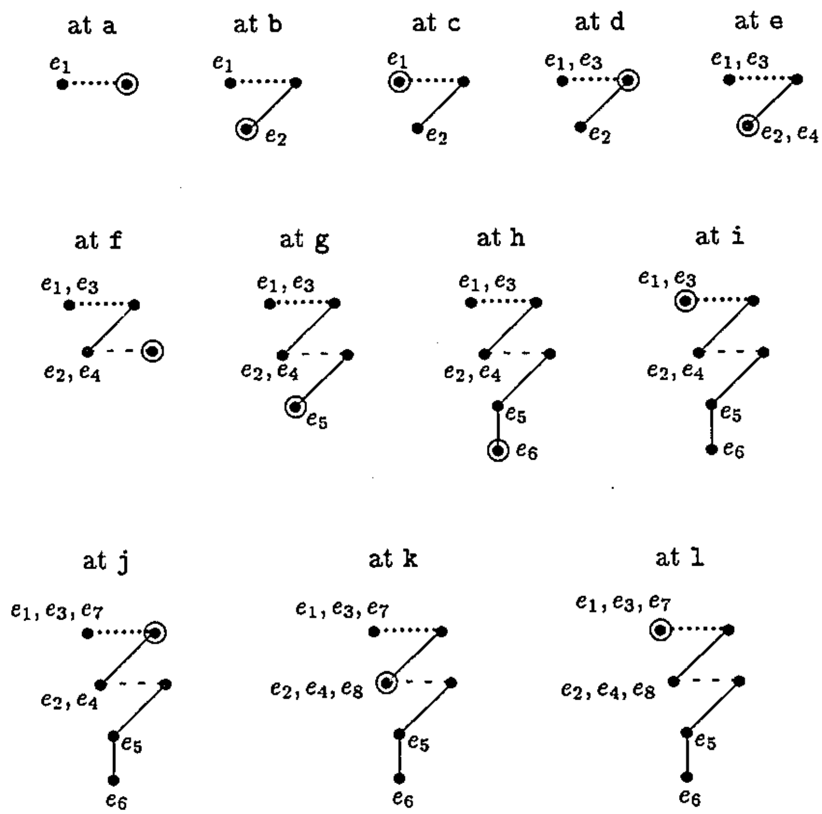

Figure 8.2: Snapshots of Tense Trees

deindexing rules take correct account of embedding. For instance, the embedded present perfect in a sentence such as "John will think that Mary has left" will be correctly interpreted as relativized to John's (future) thinking time, rather than the speech time, not as in a Reichenbachian analysis. Also, as verified already with the previous example, the episodes evoked by successive sentences or by embedded clauses within the same sentence are correctly connected to each other. In this example, the orienting relation between John's going to the hospital, e2, and the doctor's giving a diagnosis, e4, and the orienting relation between the doctor's diagnosis, e4, and giving John a crutch, e8, are automatically incorporated into the deindexed formulas. See [e2 orients e4] in (8.39 c) and [e4 orients e8] in (8.40 c). These orienting relations may then be plausibly particularized to [e4 after e2] and [e8 after e4], based on their aspectual class and coherence seeking inference.

8.5 Summary and Remarks

I have described a new, principled approach to tense and aspect interpretation within a compositional framework for language understanding. (Deindexing of operators other than tense and aspect is described in the next chapter.) The central concept is that of a tense tree structure as part of a more general context structure. This provides a straightforward and easily visualized way of converting originally indexical LFs to representations of the meaning of an utterance which are context-independent at least as far as the episodic relations implicit in the tense-aspect structure are concerned. This includes the most common "orienting" relations between episodes, and some of the more subtle relations conveyed by perfect aspect.

The mechanism is compositional in the sense that operators corresponding to past (`past`), present (`pres`), future (`futr`) and perfect (`perf`) contribute separately and uniformly to the meanings of their operands, i.e., formulas at the level of LF. Thus, for instance, the temporal relations implicit in "John will have left" are obtained not by extracting a "future perfect" and asserting relations among E, R and S, but rather by successively taking account of the meanings of the nested `pres`, `futr` and `perf` operators in the LF of the sentence.

As it happens, each of the operators, `decl`, `past`, `futr`, `perf`, etc. (with the exception of `prog`), implicitly introduces exactly "one" episode. Thus, a simple present sentence like "John is tired" would introduce only *one* episode concurrent with the speech time, not *two* as in Reichenbach's analysis. (In fact, we believe Reichenbach's reference time in simple tense is redundant.) Even more importantly for present purposes, each of `pres`, `past`, `futr` and `perf` is treated uniformly in deindexing and context change. More specifically, they drive the generation and traversal of *tense trees* in deindexing. As we have seen, the way tense trees are used allows the *orienting* episodes to be located automatically from the tree.

The scheme is easy to implement, and has been successfully used in the TRAINS interactive planning advisor at Rochester [Allen and Schubert, 1991]. A Common Lisp implementation of the deindexing process [Hwang, 1992] allows rules to be straightforwardly represented and easily edited; sample sentences of the type shown in this chapter run in roughly a tenth of a second on a Sun SPARCstation 1.

Chapter 9

Temporal Deindexing of Various Logical Operators and Adverbials

In Chapter 8, I introduced tense trees and showed basic deindexing rules adequate for some simple sentences. In this chapter, I discuss some more "advanced" rules that are needed to deindex complex sentences involving logical operators such as quantifiers, negation and coordination, infinitives, and various kinds of adverbials. After that, I briefly discuss possible ways of extending the tense tree and the deindexing mechanism to handle discourse involving perspective shifts. My aim in this chapter is mostly to provide food for thought and pointers for future work; thus, the discussion will be both tentative and informal. All the same, the proposals made give reason to think that the deindexing scheme developed so far can be extended to deal with the implicit episodic relations in arbitrarily complex formulas.

9.1 Deindexing of Logical Compounds

The basic idea in the deindexing rules for logical compounds is to make them evoke episodes corresponding to the syntactically embedded formulas, and relate them appropriately to the superordinate episode. One point that always requires care is the need to assert the intuitively required orienting relations, while avoiding assertion of orienting relations which have no intuitive support. The tree-embedding mechanism can be used to good effect here, since episodes within an embedded tree are co-accessible but generally "hidden" from the embedding environment. I now discuss possible ways of deindexing various logical operators, i.e., connectives (conjunction and disjunction), negation, and quantifiers, in that order.

9.1.1 Conjunction

Conjunction is indexical in the sense the *ordering* of the conjuncts contributes to the meaning of the conjunctive formula. For instance, we understand "The child looked around and took the candy" and "The child took the candy and looked around" differently. Below are listed some sentences involving conjunction, with their preliminary (unscoped) ULFs.

(9.1) a. John lost his job and his wife.
b. (decl [John <past lose> <∧ <The λx[[x job] ∧ [x of-genitive John]]>
<The λx[[x wife] ∧ [x of-genitive John]]>>])

(9.2) a. John loved Mary and she loved him.
b. (decl <∧ [John <past love> Mary] [Mary <past love> John]>)

(9.3) a. John packed and left.
b. (decl [John <∧ <past pack> <past leave>>])

(9.4) a. Mary sang and danced for an hour.
b. (decl ((adv-e (lasts (K1 ((num 1) hour))))
[Mary <∧ <past sing> <past dance>>]))

(9.5) a. John has bought but will soon sell the house.[1,2]
b. (decl [John <but λx<pres (perf [x buy <The house>$_i$])>
λx<pres (futr [x sell <The house>$_i$])>>])

(9.6) a. Mary left John, and he misses her.
b. (decl <∧ [Mary <past leave> John] [John <pres miss> Mary]>)

The connectives and tense operators in the above ULFs need to be scoped. Normally tense operators are raised to have the widest scope possible. Sentence (9.1) has only one tense operator, with the coordination over NPs; so, aside from the definite quantifier The, there is only one possible scoping. The rest of the examples, however, involving conjunction of tensed VPs, have more than one tense operator. In particular, each of the sentences (9.2)–(9.5) has the same tense operators in its conjuncts, while (9.6) has different ones. Here the question arises whether the tense operators need to be raised above the connectives. One seemingly effective method, which I adopt here, is to raise them as high as possible even over the connectives, as long as there is no clash of tenses. For instance, in (9.2)–(9.5), the

[1] Note that "but" is kept in the translation because of pragmatics.

[2] The index in <The house>$_i$ indicates that the two occurrences of this term must either be scoped as a single quantifier, or as two quantifiers with disjoint scopes. Also, if the meaning of "house" can be particularized in more than one way, the same particularization must be used in the disjoint-scope case (see [Schubert and Pelletier, 1982]).

tense operators are raised above the connective as well as above their immediate clauses, and "factored" out into one.

One advantage of this factoring is that it implicitly introduces a complex super-episode characterized by the conjoined sentence, with a subepisode for each conjunct (either concurrent with the super-episode or a temporal part of it).[3] The super-episode is automatically obtained at the deindexing stage via the common tense operator. Note that in (9.4), in particular, it is this super-episode that lasted for an hour. For formulas with different tenses such as (9.6), the connective acts as scope trap. (If tense operators were raised over the coordinator, they would clash.) Thus, two different conjunction deindexing rules are proposed depending on whether the conjuncts are tenseless or tensed.

I first show a possible deindexing rule for conjunction of tenseless formulas.

∧(tenseless): For Φ and Ψ tenseless formulas,

$$[\Phi \wedge \Psi]_T \leftrightarrow [(\exists e_T{:}[e_T \text{ subep-of } \texttt{Last}_T]\ [\Phi_{\circ\mapsto T} ** e_T]) \wedge (\exists e_{T'}{:}[[e_{T'} \text{ subep-of } \texttt{Last}_T] \wedge [e_T \text{ orients } e_{T'}]]\ [\Psi_{\circ T'} ** e_{T'}])],$$
$$\text{where } T' = \Phi \cdot (\circ\mapsto T)$$

Tree transformation: $[\Phi \wedge \Psi] \cdot T = \leftarrow(\Psi \cdot (\circ(\Phi \cdot (\circ\mapsto T))))$

Here, the tense operator has already been deindexed, leading to a conjunctive formula characterizing an episode, of the sort $[[\Phi \wedge \Psi] ** \eta]$. Note that the tenseless-**∧** rule generates subepisodes characterized by each conjunct. The episode characterized by Φ orients the episode characterized by Ψ, and both are subepisodes of a larger one characterized by $\Phi \wedge \Psi$. This allows, for example, that in (9.2) John's loving Mary becomes the orienting episode for Mary's loving John (once the pronouns are disambiguated); and in (9.3), John's packing orients his leaving. As discussed in Chapter 8, for two open-ended episodes, the `orients` relation suggests that the episodes are concurrent (as in John and Mary's loving each other), while for closed ones, one after the other (as in the case of John's packing and leaving). But this is only an implicature; e.g., (9.1) does not necessarily imply that John's losing job took place earlier than his losing wife. The main point here is that conjunction introduces subepisodes, enabling any implicit temporal relations between the subepisodes to be made explicit.

Note that the tenseless-**∧** rule embeds a new tree, and stores both e_T and $e_{T'}$ in the newly embedded tree as episodes "subordinate" to the outer episode ($\texttt{Last}_T$), rather than on a par with it. Since subepisodes are stored in the embedded tree, they are less easily available in processing subsequent sentences. For instance, if (9.1) is followed by "But he won a lottery," this episode is oriented by the conjoined episode of his losing his job and wife, rather than by the second subepisode of his losing his wife.

[3]Aside from this practical convenience, factoring of tense in conjoined VPs actually occurs in some languages, e.g., Korean and Japanese.

Finally, note that Ψ is processed relative to a context structure which has already been modified by Φ, as is necessary for making explicit any implicit temporal relations. In short, the context created, or augmented, by Φ is used in the interpretation of Ψ. I now illustrate the tenseless-**∧** rule by showing ELFs of a couple of the sentences shown above. The (c)-formulas below are scoped LFs, and the (d)-formulas are deindexed ELFs. I neglect most `orients` relations (except the ones obtained from the conjunction), and omit showing tense trees.

(9.1′) a. John lost his job and his wife.

c. (decl (past [(The x: [[x job] ∧ [x of-genitive John]] [John lose x]) ∧
(The y: [[y wife] ∧ [y of-genitive John]] [John lose y])]))

d. (∃e_1: [e_1 same-time *Now1*]
[[Speaker tell Hearer (That
(∃e_2: [e_2 before e_1]
[[(∃e_3: [e_3 subep-of e_2]
[(The x: [[x job] ∧ [x of-genitive John]] [John lose x])
** e_3]) ∧
(∃e_4: [[e_4 subep-of e_2] ∧ [e_3 orients e_4]]
[(The y: [[y wife] ∧ [y of-genitive John]] [John lose y])
** e_4])]
** e2]))]
** e_1])

(9.2′) a. John loved Mary and she loved him.

c. (decl (past [[John love Mary] ∧ [Mary love John]]))

d. (∃e_1: [e_1 same-time *Now1*]
[[Speaker tell Hearer (That
(∃e_2: [e_2 before e_1]
[[(∃e_3: [e_3 subep-of e_2] [[John love Mary] ** e_3]) ∧
(∃e_4: [[e_4 subep-of e_2] ∧ [e_3 orients e_4]] [[Mary love John] ** e_4])]
** e_2]))]
** e_1])

Notice in (9.1′d) episodes e_3 and e_4 — the episode of John's losing his job and the episode of his losing his wife. This is in contrast with Reichenbach's approach which introduces episodes only for tensed clauses.

The deindexing rule for coordination of tensed conjuncts is simpler as shown below.

∧ `(tensed)`: For Φ or Ψ tensed formulas,

$$[\Phi \wedge \Psi]_T \leftrightarrow [\Phi_T \wedge \Psi_{\Phi \cdot T}]$$

Tree transformation: $[\Phi \wedge \Psi] \cdot T = (\Psi \cdot (\Phi \cdot T))$

This rule basically treats Φ and Ψ as two separate, successive sentences. It stores the episode characterized by each conjunct in the main tense tree rather than "hiding" them in an embedded tree. Below is shown the deindexed ELF of (9.6), with its scoped LF.

(9.6′) a. Mary left John, and he misses her.

c. (decl [(past [Mary leave John]) ∧ (pres [John miss Mary])])

d. ($\exists e_1$: [e_1 same-time *Now1*]
[[Speaker tell Hearer (That
[($\exists e_2$: [e_2 before e_1] [[Mary leave John] ** e_2]) ∧
($\exists e_3$: [e_3 at-about e_1] [[John miss Mary] ** e_3])])]
** e_1])

Note that there is no explicit `orients` relation between e_2 and e_3 above. I leave deindexing of the rest of the example sentences to the reader as they are quite straightforward to trace.

9.1.2 Disjunction

We are interested in the following kinds of disjunctive sentences.

(9.7) John will eat a pizza or a hamburger.

(9.8) I will either read the newspaper or watch TV for an hour.

(9.9) John may or may not come.

(9.10) Either John ate the hamburger or Mary ate it.

(9.11) Either John is sick or he heard some bad news.

Sentence (9.7) involve disjunction in the PP-adverbial or in the object NP, with a single tensed VP. In (9.7), there is John's eating episode in the future, but its exact characterization is unknown — it could be JOHN EAT PIZZA or JOHN EAT HAMBURGER. Another way of looking at it is consider (9.7) as describing *two* episodes — a "pizza eating" episode and a "hamburger eating" episode — without indicating which of them is actual. Each of sentences (9.8)–(9.11) involves disjunction of two tensed VPs. Here, one could take these sentences as describing two possible episodes or one episode with two alternative characterizations. Thus, there arises the question of whether to factor out the same tenses as in conjunction, as different scopings give different readings. Sentences like (9.8) seem to favor factoring; so we tentatively propose to do so.

As a final point, note that there are no `orients` relations between the episodes evoked by the disjuncts. For instance, in (9.10), the episode of John's eating and the episode of Mary's eating are not related with each other in any way. One of them does not even exist. I now show the (unscoped) ULFs and (scoped) LFs for a couple of the sentences shown above. Notice that tense operators have been factored in (9.10′c).

(9.10′) a. Either John ate the hamburger or Mary ate it.
b. (decl <∨ [John <past eat> <The hamburger>$_i$] [Mary <past eat> *It*$_i$]>)
c. (decl (past (The x: [x hamburger] [[John eat x] ∨ [Mary eat x]])))

(9.11′) a. Either John is sick or he heard some bad news.
b. (decl <∨ [John$_i$ <pres sick>] [*He*$_i$ <past hear> <∃ ((attr bad) news)>]>)
c. (decl [(pres [John sick]) ∨ (past (∃x: [x ((attr bad) news)] [John hear x]))])

One caveat in developing deindexing rules for disjunction is that the episodes evoked by disjuncts are not necessarily actual and should not be used as orienting episodes while subsequent sentences are processed. This is easily taken care of in the tense tree mechanism by placing those episodes in an embedded tree, instead of the main tree. Only a super-episode whose scope is over the disjunction may be placed in the main tree. I propose the following rules for disjunction.

∨`(tenseless)`: For Φ and Ψ tenseless formulas,

$$[\Phi \vee \Psi]_{\mathrm{T}} \leftrightarrow [\Phi_{\mathrm{T}'} \vee \Psi_{\Phi\cdot\mathrm{T}'}]$$

$$\text{where } \mathrm{T}' = \otimes_{\mathtt{Last}_{\mathrm{T}}} \mapsto \mathrm{T}$$

Tree transformation: $[\Phi \vee \Psi] \cdot \mathrm{T} = \leftarrow(\Psi \cdot (\Phi \cdot (\otimes_{\mathtt{Last}_{\mathrm{T}}} \mapsto \mathrm{T})))$

∨`(tensed)`: For Φ and Ψ tensed formulas,

$$[\Phi \vee \Psi]_{\mathrm{T}} \leftrightarrow (\exists e_{\mathrm{T}}{:}\, [e_{\mathrm{T}} \text{ at-about } \mathtt{Emb}_{\mathrm{T}}]\ [[\Phi_{\mathrm{T}'} \vee \Psi_{\mathrm{T}''}] ** e_{\mathrm{T}}])$$

$$\text{where } \mathrm{T}' = \mapsto \circ \mathrm{T}, \quad \text{and} \quad \mathrm{T}'' = \mapsto \leftarrow (\Phi \cdot \mathrm{T}')$$

Tree transformation: $[\Phi \vee \Psi] \cdot \mathrm{T} = \leftarrow(\Psi \cdot (\mapsto \leftarrow (\Phi \cdot (\mapsto \circ \mathrm{T}))))$

In the tenseless-**∨** rule, the symbol '$\otimes_\eta \mathrm{T}$', for η an episode, denotes the tense tree which is just like T except that the token η (note: *not* the default, new episode e_{T}) has been added to the focal node of T. ('$\circ\mathrm{T}$' may be regarded as an abbreviation of '$\otimes_{e_{\mathrm{T}}} \mathrm{T}$'.) The tenseless-**∨** rule first embeds a new tree, shifts focus to the root node of the embedded tree, stores $\mathtt{Last}_{\mathrm{T}}$ (i.e., copies the token $\mathtt{Last}_{\mathrm{T}}$) at the new focus, and processes each disjunct within the embedded tree. Even though Φ and Ψ are tenseless, they may be headed by "pseudo-tense" operators such as `perf` and `futr`. Any such operators will be interpreted with respect to $\mathtt{Last}_{\mathrm{T}}$ which is stored at the root node. After both disjuncts are processed, the focus is shifted back to the embedding node. The result is that neither of the episodes evoked by the disjuncts (if any) will be available while subsequent sentences are processed. The tensed-**∨** rule works slightly differently. It first generates a new episode e_{T}, stores it at the current focal node, and predicates it to be at the same time as the embedding episode (which is likely to be an utterance or an attitude episode). The new episode e_{T} becomes a superordinate episode characterized by the disjunctive formula. After this, a new tree is embedded, Φ is processed within the embedded tree, and the focus is shifted back to the embedding node. Then another new tree is embedded, and Ψ is processed in an analogous way. Since episodes generated by Φ and Ψ are stored in different subtrees, they will not be co-accessible. Note that the superordinate episode e_T automatically plays the role of $\texttt{Emb}_T$ when the tense operators heading Φ and Ψ are deindexed.

I now show examples: (9.10′) below illustrates the tenseless-**V** rule, and (9.11′) illustrates the tensed-**V** rule. I repeat the scoped LFs.

(9.10′) a. Either John ate the hamburger or Mary ate it.

c. (decl (past (The x: [x hamburger] [[John eat x] $\vee$ [Mary eat x]])))

d. ($\exists e_1$: [e_1 same-time *Now1*]
    [[Speaker tell Hearer (That
        ($\exists e_2$: [e_2 before e_1]
            [(The x: [x hamburger] [[John eat x] $\vee$ [Mary eat x]]) ** e_2]))]
    ** e_1])

(9.11′) a. Either John is sick or he heard some bad news.

c. (decl [(pres [John sick]) $\vee$ (past ($\exists x$: [x ((attr bad) news)] [John hear x]))])

d. ($\exists e_1$: [e_1 same-time *Now1*]
    [[Speaker tell Hearer (That
        ($\exists e_2$: [e_2 at-about e_1]
            [[($\exists e_3$: [e_3 at-about e_2] [[John sick] ** e_3]) $\vee$
             ($\exists e_4$: [e_4 before e_2]
                [($\exists x$: [x ((attr bad) news)] [John hear x]) ** e_4])]
            ** e_2]))]
    ** e_1])

Thanks to the way the disjunction rules work, the deindexed formulas are concise and have uniform appearance in both tensed and tenseless cases. Interpretation of disjunctive formulas is straightforward with the truth conditions provided in Chapter 4.

9.1.3 Negation

Negative sentences describe a situation in which *no* events of a certain type take place or a certain type of state does not hold true. Consider the following examples.

(9.12) John has not left.

(9.13) John did not sleep for three days.

(9.14) John did not ask Mary to dance at the party. *It* made her angry.

(9.15) Mary was not happy during the party.

(9.16) Mary did not turn off the stove.

(9.17) John did not wash and shave in ten minutes.

Sentences (9.12)–(9.14) indicate non-occurrence of a certain type of event, i.e., John's "leaving," "sleeping," and "asking," during a certain time interval. (9.15) indicates that a certain state, namely, MARY HAPPY, did not hold true during the party. Sentences (9.16)–(9.17) also say that a certain type of events did not take place (though their temporal locations are not specified). Interestingly, such non-occurrences can be causal antecedents or consequents, as illustrated in (9.14). Note also that though the episodes within the scope of negation may be "nonexistent," they may be related to each other temporally or otherwise. For example, the "washing" episode orients the "shaving" episode in (9.17). However, such nonexistent events cannot orient episodes outside the negation. Again, the tree embedding technique can be used to this effect.

Let us now consider some possible deindexing rules for negation. The following is a simple version adequate for some simple negated formulas.

$\neg$: $(\neg \Phi)_{\mathrm{T}} \leftrightarrow (\neg \Phi_{\mathrm{T}'})$, where $\mathrm{T}' = \otimes_{\mathtt{Last}_\mathrm{T}} \mapsto \mathrm{T}$
Tree transformation: $(\neg\Phi) \cdot \mathrm{T} = \leftarrow(\Phi \cdot (\otimes_{\mathtt{Last}_\mathrm{T}} \mapsto \mathrm{T}))$

We assume here the negated formula $(\neg \Phi)$ is immediately embedded by '**'. This is likely to be the case if tense operators take a wider scope than the negation. The $\neg$-rule first embeds a new tree, shifts focus to its root node, stores $\mathtt{Last}_\mathrm{T}$ there, and processes Φ within the embedded tree, with the result that the nonexistent episodes (if any) generated under the scope of the negation are inaccessible for orienting relationship with episodes in the main tree. However, these episodes *will* be co-accessible; e.g., in (9.17), the "washing" and "shaving" episodes will be co-accessible (cf., "washing" orients "shaving"). The following illustrates the use of the $\neg$-rule.

(9.12′) a. John has not left.
b. (decl (pres (perf (¬ [John leave]))))
c. $(\exists e_1$: $[e_1$ same-time *Now1*]
[[Speaker tell Hearer (That
$(\exists e_2$: $[e_2$ at-about $e_1]$
$[(\exists e_3$: $[e_3$ impinges-on $e_2]$ $[(\neg$ [John leave]) ** e_3]) ** e_2]))]
** e_1])

(9.13′) a. John did not sleep for three days.
b. (decl (past ((adv-e (lasts (K1 ((num 3) day)))) (¬ [John sleep]))))

c. $(\exists e_1 \colon [e_1 \text{ same-time } \mathit{Now1}]$
$[[\texttt{Speaker} \text{ tell } \texttt{Hearer} \text{ (That}$
$(\exists e_2 \colon [e_2 \text{ before } e_1]$
$[[[e_2 \text{ lasts (K1 ((num 3) day))}] \wedge (\neg [\text{John sleep}])] \text{ ** } e_2]))]$
$\text{** } e_1])$

(9.16′) a. Mary did not turn off the stove.

b. (decl (The x: [x stove] (past (¬ [Mary turn-off x]))))

c. $(\exists e_1 \colon [e_1 \text{ same-time } \mathit{Now1}]$
$[[\texttt{Speaker} \text{ tell } \texttt{Hearer} \text{ (That}$
$(\text{The } x \colon [x \text{ stove}]$
$(\exists e_2 \colon [[e_2 \text{ before } e_1] \wedge [e_9 \text{ orients } e_2]] [(\neg [\text{Mary turn-off } x]) \text{ ** } e_2])))]$
$\text{** } e_1])$

Formula (9.12′c) says that at episode e_3, which extends to the utterance time, John leave is false, i.e., there was no episode of type John leave. (9.13′c) implies there was a three day-long episode, at which John was awake.[4] This splitting of an adverbially modified sentence into a conjunction is done while deindexing adverbials (to be discussed shortly). (9.16′c) says that during some time in the past, whose exact temporal location is to be contextually determined (note [e_9 orients e_2], where e_9 is assumed to be an episode introduced by the previous text), no event of type Mary turn off Stove took place.

Note that the above treatment of negation is in contrast with Krifka's [1989], who analyzes a negated expression as referring to "maximal events." For example, he represents (9.13), "John did not sleep for three days," as *the fusion of all events* that took place during or at the same time as the three day-long interval and are not events of John's sleeping. The difficulty with this kind of approach is that such events cannot serve as causal antecedents or consequents.

The ¬-rule needs extension, however. First, it cannot be applied if the negation is not immediately embedded by '**' because then $\texttt{Last}_\text{T}$ will be either undefined or identified with a wrong episode. In such cases, negation should be interpreted with respect to $\texttt{Emb}_\text{T}$, rather than $\texttt{Last}_\text{T}$. Also, the rule is not general enough to handle complex Φ, involving embedded clauses and/or adverbials. For instance, for Φ [John know (*That* (*past* (*perf* [Mary leave]))]), "John did (*not*) know that Mary had left," the embedded indexical operators need to be deindexed with respect to the episode that is characterized by Φ. Since the ¬-rule does not introduce an episode for Φ, the deindexing process cannot proceed at this point.

Next, as an example of Φ involving an adverbial, let us consider another (less preferred) reading of (9.13), namely, "It is not the case that John (ever) slept for three days," as

[4]There is another reading of (9.13), e.g., "It was not the case that John slept for three days." This will be discussed shortly.

indicated below:

(decl $(\neg$ (past ((adv-e (lasts (K1 ((num 3) day)))) [John sleep])))).

This says there was no situation in the past which lasted three days *and* at which JOHN SLEEP was true. Here, it is a nonexistent situation the adverbial ("three day-long") is referring to. For this we need to make those situations implicit within the scope of negation explicit. We may modify the $\neg$-rule as follows.

$\neg$ (1): $(\neg\Phi)_{\mathrm{T}} \leftrightarrow (\neg(\exists e_{\mathrm{T}}{:}[e_{\mathrm{T}} \text{ at-about } \mathtt{Last}_{\mathrm{T}}][\Phi_{\mathrm{O}\mapsto\mathrm{T}} ** e_{\mathrm{T}}]))$,
for $(\neg\Phi)$ immediately embedded by '**'

Tree transformation: $(\neg\Phi)\cdot\mathrm{T} = \leftarrow(\Phi\cdot(\mathrm{O}\mapsto\mathrm{T}))$

$(\neg\Phi)_{\mathrm{T}} \leftrightarrow (\neg(\exists e_{\mathrm{T}}{:}[e_{\mathrm{T}} \text{ at-about } \mathtt{Emb}_{\mathrm{T}}][\Phi_{\mathrm{T}\mapsto\mathrm{O}} ** e_{\mathrm{T}}]))$,
for Φ tensed and $(\neg\Phi)$ *not* immediately embedded by '**'

Tree transformation: $(\neg\Phi)\cdot\mathrm{T} = \leftarrow(\Phi\cdot(\mathrm{T}\mapsto\mathrm{O}))$

With this rule, the above, less preferred reading of (9.13) may be deindexed as follows.

(9.13″) b. (decl $(\neg$ (past ((adv-e (lasts (K1 ((num 3) day)))) [John sleep]))))
c. $(\exists e_1{:}[e_1$ same-time *Now1*]
[[Speaker tell Hearer (That
$(\neg(\exists e_2{:}[e_2$ at-about $e_1]$
$[(\exists e_3{:}[e_3$ before $e_2]$
$[[[e_3$ lasts (K1 ((num 3) day))] $\wedge$ [John sleep]] ** $e_3])$
** $e_2]))))]$
** $e_1])$

Notice the episode e_3 which is three day-long and at which JOHN SLEEP is true. Also notice that the past operator is essentially interpreted with respect to episode e_1 that was originally at the *embedding* node of the main tree.

One disadvantage of the $\neg$ (1)-rule is that it is redundant in the sense that the RHS of the rule automatically follows from the semantics of negation. The simplicity we originally had with the $\neg$-rule is lost. For instance, compare the following two possible translations for "John is not happy": the first is obtained by the $\neg$-rule, the second by the $\neg$ (1)-rule.

$(\exists e_2{:}[e_2$ at-about $e_1]$ $[(\neg$ [John happy]) ** $e_2])$
$(\exists e_2{:}[e_2$ at-about $e_1]$ $[(\neg(\exists e_3{:}[e_3$ at-about $e_2]$ [[John happy] ** $e_3]))$ ** $e_2])$

Incidentally, it may be felt that a negated sentence necessarily expresses a fact rather than a time bounded episode. That is, the logical form of the first sentence in (9.14) would be something like "There is a time interval which coincides with the party and during which there are no events of John's asking Mary to dance"; and the antecedent of *It* in the second sentence would be the *fact* that there is a time interval which coincides with the party, etc. However, the examples below indicate otherwise; i.e., *That* in the second sentences makes anaphoric reference to the alleged negative episodes of the first sentences, which could *not* possibly be factual episodes.

(9.18) A couple of times I didn't finish a project I had started. *That* happened when I was really depressed.

(9.19) Jack didn't see Mary for almost three weeks last summer. *That* was while Sue was visiting him.

(9.20) Sometimes they do not talk to each other for days. *That* usually happens when they are back from holidays.

The negation episode in (9.18) cannot be a fact as facts cannot "happen". Only time bounded episodes or events do (cf., [Vendler, 1967]). Also, in (9.19), it is more natural to take the negative episodes of the first sentence as episodes because of their timeboundedness as demonstrated in the second sentence. In (9.20), the negation episode is a generic one, but again it is a "generic episode," not a "generic fact," as a fact cannot happen.

9.1.4 Quantifiers

The interaction between tense operators and nominals, especially quantified nominals, has been frequently discussed since Enç raised the issue [1981]. This issue will be addressed here by way of discussing quantifier deindexing. Recall from Chapter 4 that quantifiers can be classed into *monotone increasing* ones, *monotone decreasing* ones, and that are neither. A quantifier '$\mathtt{Q}$' is monotone increasing if for all formulas Φ, Ψ and Υ,

$$(\mathtt{Q}\,\alpha{:}\,\Phi\Psi), (\forall\alpha{:}\,\Psi\Upsilon) \vdash (\mathtt{Q}\,\alpha{:}\,\Phi\Upsilon).$$

Examples of monotone increasing ones are $\forall$, $\exists$, `The`, `Most`, `Many`, etc., and those of monotone decreasing ones are `No`, `Few`, etc. I will discuss only monotone-increasing ones here. However, note that many (if not all) monotone decreasing quantifiers, and ones that are neither monotone increasing nor decreasing, can be paraphrased in terms of monotone increasing quantifiers plus negation. For instance, (Few $\alpha{:}\,\Phi\Psi$) can be paraphrased as $\neg$(Many $\alpha{:}\,\Phi\Psi$), and (No $\alpha{:}\,\Phi\Psi$) can be paraphrased as $\neg(\exists\,\alpha{:}\,\Phi\Psi)$. A cardinal "quantifier" such as "exactly 12" is interpreted here as involving an existential quantification of a *collection* of things, with specified cardinality; e.g., "*some* collection of exact size 12." Here the existential quantifier is again monotone increasing.

I show below some sample sentences involving monotone increasing quantifiers.

(9.21) A millionaire proposed to Mary

(9.22) A bubble burst

(9.23) A man lived in a cottage for 50 years

(9.24) *A child lived in a cottage for 50 years

(9.25) A child was born

(9.26) *A 16 year old boy was born

(9.27) *A man was born

(9.28) ?The man was born

(9.29) The man was born in South Africa

(9.30) ?My father grew up

(9.31) My father grew up in a small town

(9.32) The old man was a hunter

(9.33) ?Every man grew up

(9.34) Every secretary got promoted last year

(9.35) Every bearded man shaved

(9.36) Most passengers got off at Union Station

It seems that an existential quantifier typically insists that the restriction hold at the same time as the matrix. For instance, in (9.21), the person who proposed to Mary must have been a millionaire at the time of proposing; in (9.22), the thing must have been a bubble during the bursting event (possibly excluding the end point). This also explains why sentences (9.24) and (9.26)–(9.27) sound odd. This requirement, however, loses its force with the definite quantifier The. That is, the detachment of the time at which the restriction holds from the time at which the matrix holds becomes easier as seen in (9.28)–(9.32). In fact, (9.32) even admits ambiguity: e.g., the speaker may be talking about a man who is old at the speech time and was a hunter when he was younger, or about an old hunter of some past time (he might no longer be alive).

This detachment of the time reference of the restriction clause from that of the main clause occurs even more freely with universal quantifiers. For instance, sentence (9.33) does not sound very odd though it certainly sounds trivial. Consequently, ambiguity arises more easily with universal quantification. In (9.34), for instance, the speaker may be talking about current employees or last year's employees. We assume such ambiguity can be obtained through different scopal analyses. Note, however, as demonstrated in

(9.35) and (9.36), the restriction holds at the time of the matrix by default much as in existential quantifiers.

I now show a possible deindexing rule for quantifiers.

Quant : $(\mathtt{Q}\,\alpha\!:\Phi\Psi)_{\mathrm{T}} \leftrightarrow (\mathtt{Q}\,\alpha\!:\Phi_{\mathrm{T}}\ \Psi_{\Phi\cdot\mathrm{T}})$
Tree transformation: $(\mathtt{Q}\,\alpha\!:\Phi\Psi)\cdot\mathrm{T} = \Psi\cdot(\Phi\cdot\mathrm{T})$

This rule assumes the restriction Φ is atomic and untensed, involving no embedded clauses, conjunction or relative clauses. Note that this rule does not generate new episodes. Many of the legitimate sentences in (9.21)–(9.36) can be handled by this rule. I show deindexed ELFs for some of those sentences to illustrate the Quant-rule.

(9.21′) a. A millionaire proposed to Mary.
b. (decl (past ($\exists x$: [x millionaire] [x propose-to Mary])))
c. ($\exists e_1$: [e_1 same-time *Now1*]
[[Speaker tell Hearer (That
($\exists e_2$: [e_2 before e_1] [($\exists x$: [x millionaire] [x propose-to Mary]) ** e_2]))]
** e_1])

(9.35′) a. Every bearded man shaved.
b. (decl (past ($\forall x$: [x ((attr bearded) man)] [x shave])))
c. ($\exists e_1$: [e_1 same-time *Now1*]
[[Speaker tell Hearer (That
($\exists e_2$: [e_2 before e_1] [($\forall x$: [x ((attr bearded) man)] [x shave]) ** e_2]))]
** e_1])

(9.36′) a. Most passengers got off at Union Station
b. (decl (past (Most x: [x passenger] ((adv-e (at-loc UnionStation)) [x get-off]))))
c. ($\exists e_1$: [e_1 same-time *Now1*]
[[Speaker tell Hearer (That
($\exists e_2$: [e_2 before e_1]
[(Most x: [x passenger] [x get-off-at-UnionStation])
** e_2]))]
** e_1]),
with a rough translation of the adverbial

Deindexing of quantifiers is similar to that of conjunctions. (9.21′c) says that at situation c_2, someone is a millionaire *and* he proposes to Mary. In (9.35′c), we can intuitively see that there is not only an "overall" super-episode covering all the shavings, but also a set of individual "shaving" episodes, one for each person in question. Each "shaving" is temporally contained in its corresponding "being bearded" episode (possibly except the end point).[5] This follows from the semantics of $\forall$, and can be made explicit with a meaning postulate. Nonetheless, only the super-episode that has been generated by the wide-scoped past will be available as an orienting episode for the subsequent text. Similarly, in (9.36′c), most of the people who were passengers during e_2 got off at Union Station while he was a passenger. Again, each "being passenger" episode ends when its corresponding "getting off" episode does. Also, only the "overall" episode will be available as an orienting episode; one cannot "get at" the episode corresponding to *individual* passengers getting off at Union Station as an orienting episode. This is automatically handled by the Quant-rule (notice that the subepisodes about individual passengers are not even generated).

Unfortunately, the Quant-rule has a similar problem as the $\neg$-rule. It cannot deal with sentences involving embedded clauses, e.g., "Everyone knew that Mary was smart." Besides, it cannot handle sentences involving relative clauses. A restriction involving a relative clause is typically of form $[\Phi \wedge \Phi']$, where Φ is tenseless and Φ' is tensed (e.g., $[[x$ man$] \wedge$ (past $[x$ propose-to Mary$])]$, "man who proposed to Mary"). Since the Quant-rule does not create an episode for the restriction, indexical operators in the restriction (if any) cannot be deindexed unless the restriction is immediately embedded by '**' (cf., see the tensed Λ-rule). To remedy this, what's needed is, first, embed a tree at the current focal node, generate and store an episode for the restriction at the root node of the embedded tree, and process Φ within the embedded tree, shifting the focus back to the embedding node afterwards. Also, to handle sentences with embedded clauses, a new tree needs to be embedded for the matrix clause (the body of the assertion, often derived from a VP). I omit describing the details of the new rules; instead, show below a couple of target deindexed ELFs for quantification involving relative clauses. Note that in (9.37), the restriction is immediately embedded by '**', but not in (9.38).

(9.37) a. John married a girl who would exploit him.

b. (decl (past ($\exists x$: [[x girl] $\wedge$ (past (futr [x exploit John]))] [John marry x])))

c. ($\exists e_1$: [e_1 same-time *Now1*]
  [[Speaker tell Hearer (That
    ($\exists e_2$: [e_2 before e_1]
      [($\exists x$: [[x girl] $\wedge$ ($\exists e_3$: [e_3 at-or-before e_2]
                  [($\exists e_4$: [e_4 after e_3] [[x exploit John] ** e_4])
                  ** e_3])]
        [John marry x])
      ** e_2]))]
  ** e_1])

[5]More precisely, the end point of the "shaving" episode and the end point of the "being bearded" episode meet.

(9.38) a. John married a girl who will exploit him.

b. (decl ($\exists x$: [[x girl] $\wedge$ (pres (futr [x exploit John]))] (past [John marry x])))

c. ($\exists e_1$: [e_1 same-time *Now1*]
  [[Speaker tell Hearer (That
   ($\exists x$: ($\exists e_2$: [e_2 at-about e_1]
     [[[x girl] $\wedge$ ($\exists e_3$: [e_3 at-about e_2]
         [($\exists e_4$: [e_4 after e_3] [[x exploit John] ** e_4])
         ** e_3])]
     ** e_2])
    ($\exists e_5$: [e_5 before e_1] [[John marry x] ** e_5])))]
  ** e_1])

I should mention at this point that the Quant-rule neglects some problems concerning the time reference of the quantifier restriction, Φ. These arise from well-known problems in the time reference of noun phrases, illustrated by sentences like "Every adult was once a baby," where "being an adult" should not necessarily refer only to the past or only to the present, and should not extend over the episode of "being a baby" (see [Enç, 1981]). Another problem is that sometimes the restriction on the quantification ($\forall$ and Most, in particular) needs to be further constrained by the context. For instance, compare "Mary sent a Christmas card to every friend" and "Every boy clapped at Mary." In the former, the restriction will correctly select every one who was a friend of Mary at the time of her card-writing; in the latter, however, the restriction should not be interpreted as referring to every one who was a boy at the time of the "clapping" episode, rather it refers to every one who was a boy at the time of clapping episode *at some salient spatial location*. As a solution to this problem, Hinrichs [1988] introduces into the translation of quantified NPs a predicate R, that ranges over properties that are salient in a given context and which serve to narrow down the reference of the NP in question (see Chapter 7). Although his point is well taken, there remains the problem of how to identify the intension of such a predicate R. This requires further investigation.

9.2 EL Operators Ka and Ke

In this section, I will briefly discuss deindexing of infinitival expressions. As discussed in Chapter 3, we regard *to*-infinitives such as "to swim," "to have loved," "to be happy," etc., as kinds of actions or attributes and *for-to* infinitives such as "for Mary to dance," "for Mary to have left before John arrived," etc., as kinds of events. Initially, these constructs are translated into nominals of form (Ka Π) or (Ke Φ), where Π is a monadic predicate, and Φ is a formula. Deindexing of these constructs introduces lambda episodic variables and characterization relations, and at the same time transforms them into an expression headed by K, a more general kind operator that maps monadic predicates into kinds. This is done by the following rules.

Ke: $(\text{Ke }\Phi)_{\text{T}} \leftrightarrow (\text{K }\lambda e[\Phi_{\text{T}'} ** e])$, where $\text{T}' = \otimes_e \mapsto \text{T}$

Tree transformation: $(\text{Ke }\Phi) \cdot \text{T} = \leftarrow (\Phi \cdot (\otimes_e \mapsto \text{T}))$

Ka: $(\text{Ka }\Pi)_{\text{T}} \leftrightarrow (\text{K }\lambda a[[(\text{fst } a)\ \Pi_{\text{T}'}] ** (\text{rst } a)])$, where $\text{T}' = \otimes_{(\text{rst } a)} \mapsto \text{T}$

Tree transformation: $(\text{Ka }\Pi) \cdot \text{T} = \leftarrow (\Pi \cdot (\otimes_{(\text{rst } a)} \mapsto \text{T}))$

The Ke-rule embeds a new tree, generates a lambda episode variable, and stores it at the root node of the embedded tree. Φ is then deindexed within the embedded tree. The Ka-rule works in a similar way; though it evokes an action variable, rather than an episode variable. In the Ka-rule, recall that, for a denoting an action or attribute, (fst a) denotes the agent of the action/attribute, and (rst a) denotes the episode containing the action.

I now show some examples illustrating these rules. I omit surface speech acts.

(9.39) a. For Mary to dance is rare.
b. (pres [(Ke [Mary dance]) rare])
c. ($\exists e_1$: [e_1 before u_1]
[[(K λe[[Mary dance] ** e]) rare] ** e_1])

(9.40) a. John likes to swim.
b. (pres [John like (Ka swim)])
c. ($\exists e_1$: [e_1 at-about u_1]
[[John like (K λa[[(fst a) swim] ** (rst a)])] ** e_1])

(9.41) a. John likes to skip a class.
b. (pres [John like (Ka $\lambda x(\exists y$:[y class][x skip y]))])
c. ($\exists e_1$: [e_1 before u_1]
[[John like (K λa[[(fst a) $\lambda x(\exists y$:[y class] [x skip y])] ** (rst a)])] ** e_1])

9.3 Adverbials

I first discuss deindexing rules for episodic adverbials and action-modifying adverbials, and then rules for propositional adverbials. Rules for frequency and cardinal adverbials are discussed after that.

9.3.1 Episode and Action-modifying Adverbials

As discussed in Chapter 6, episodic adverbials are translated into (adv-e π), where π is a predicate over situations such as

(during *Yesterday*) for "during yesterday"
(lasts (K1 ((num 2) hour))) for "for 2 hours"
(at-loc <The park>) for "at the park"

and action-modifying adverbials are translated into (adv-a π), where π is a predicate over actions or attributes such as

(with-accomp Mary) for "with Mary"
(in-degree complete) for "completely"
(in-manner polite) for "politely"

The following are possible deindexing rules for formulas headed by such constructs.

adv-e: $((\text{adv-e}\ \pi)\ \Phi)_{\mathrm{T}} \leftrightarrow [\,\hat{}\pi_{\mathrm{T}} \wedge \Phi_{\pi\cdot\mathrm{T}}]$,
for π a monadic predicate and Φ a formula
Tree transformation: $((\text{adv-e}\ \pi)\ \Phi)\cdot\mathrm{T} = \Phi\cdot(\pi\cdot\mathrm{T})$

adv-a: $((\text{adv-a}\ \pi)\ \Pi)_{\mathrm{T}} \leftrightarrow \lambda x[\,\hat{}\lambda e[[x \mid e]\ \pi_{\mathrm{T}}] \wedge [x\ \Pi_{\pi\cdot\mathrm{T}}]]$,
for Π a monadic predicate and Φ a formula
Tree transformation: $((\text{adv-a}\ \pi)\ \Pi)\cdot\mathrm{T} = \Pi\cdot(\pi\cdot\mathrm{T})$

These rules first interpret predicate π relative to the given tense tree T, and then interpret sentence Φ, or predicate Π, relative to the T (possibly modified while π was deindexed). At the same time, they split the formula into a conjunction of two subformulas much as in [Dowty, 1982]: one for the adverbial itself (describing the attribute of the episode or the action), the other for the sentence or VP modified by the adverbial. " ^ " in the above rules is a function that uniformly maps predicates over situations into sentence intensions, i.e., an operator of type $(\mathcal{S}\rightarrow(\mathcal{S}\rightarrow 2))\rightarrow(\mathcal{S}\rightarrow 2)$, with the following semantics:

For $s \in \mathcal{S}$ and π a predicate over situations,
$[\![\hat{}\pi]\!]^s = [\![\pi]\!](s,s)$, i.e., $[\![\pi]\!]^{s,s}$

That is, for η an episode, $[\,\hat{}\pi * \eta] \leftrightarrow [[\eta\ \pi] * \eta]$; thus, for instance,

$[\,\hat{}(\text{during } \mathit{Yesterday}) * e_1] = [[e_1 \text{ during } \mathit{Yesterday}] * e_1]$.

Intuitively, $\hat{}\pi$ means "π holds of the current episode."

I now show examples that illustrate the above rules. Consider (9.42a) below that involves three episodic adverbials.

(9.42) a. John waited for Mary {*for two hours*} {*yesterday*} {*at the park*}.

b. (decl (The x: [x park] (past ((adv-e (at-loc x))
((adv-e (during Yesterday))
((adv-e (lasts (K1 ((num 2) hour))))
[John wait-for Mary]))))))

c. ($\exists e_1$: [e_1 same-time *Now1*]
[[Speaker tell Hearer (That
(The x: [x park]
($\exists e_2$: [e_2 before e_1]
[[$\hat{}$(at-loc x) $\wedge$
[$\hat{}$(during Yesterday) $\wedge$
[$\hat{}$(lasts (K1 ((num 2) hour))) $\wedge$ [John wait-for Mary]]]]
** e_2]))))]
** e_1])

(9.42 b) shows the preliminary, scoped LF. Applying the adv-e rule (three times), together with the Decl, Quant and Past-rules, to (9.42 b), we get the completely deindexed formula (9.42 c) (neglecting the orients relation). Notice that the information in the past tense operator and the information in the temporal and locative adverbials are interpreted *conjunctively*.

Purely indexical temporal adverbials such as *Yesterday* in the above can be even more simply handled, namely, using rule yesterday$_T$ = (yesterday-rel-to Now$_T$), meaning "yesterday relative to the utterance time," where Now$_T$ is the most recently added episode at the root node of T (see Chapter 8). For adverbial '*Then*' (Then$_T$), we tentatively propose to interpret it as Last$_T$. (This rule may have to be weakened by means of a context-charged predication.) PP-adverbials of temporal location such as *after the war* are handled in principle, except that they often involve explicit anaphoric reference to events which may or may not have been stored in the tense tree, so they presumably depend on the more general mechanisms for anaphora interpretation (using history lists, etc.). But at least we are applying the predicates implicit in these adverbials to the appropriate episodes.

Next shown is an example involving an action-modifying adverbial.

(9.43) a. John bought a rose *for Mary*.

b. (decl (past [John ((adv-a (for-benef Mary)) λy($\exists x$: [x rose][y buy x]))]))

c. ($\exists e_1$: [e_1 same-time *Now1*]
[[Speaker tell Hearer (That
($\exists e_2$: [e_2 before e_1]
[[$\hat{}\lambda e$[[John | e] for-benef Mary] $\wedge$ ($\exists x$: [x rose][John buy x])]
** e_2]))]
** e_1])

Notice [John | e] in (9.43 c), an ordered pair of an agent and an episode, denoting an action. Also note again that the information in the adverbial is interpreted conjunctively.

Once formulas are deindexed, the remaining work of interpreting the adverbials as relations involving episodes is done by meaning postulates about the " ^ " operator. The following meaning postulates introduce episodic variables into a formula, so that episode-modifying adverbials can be applied to them. As usual, these are taken to be implicitly *necessitated*, i.e., prefixed with $\Box$.

MP 9.1. For π, π' predicates,
$[\hat{}\pi \wedge \hat{}\pi'] \leftrightarrow \hat{}\lambda x[[x\ \pi] \wedge [x\ \pi']]$

MP 9.2. For η an episode, π a predicate, and Φ a formula,
$[[\hat{}\pi \wedge \Phi] ** \eta] \leftrightarrow [[[\eta\ \pi] \wedge \Phi] ** \eta]$

Applying these meaning postulates to the deindexed formulas (9.42 c) and (9.43 c), we get the following results.

(9.42′) ($\exists e_1$: [e_1 same-time *Now1*]
    [[Speaker tell Hearer (That
        (The x: [x park]
            ($\exists e_2$: [e_2 before e_1]
                [[[e_2 at-loc x] ∧
                 [e_2 during Yesterday] ∧ [e_2 lasts (K1 ((num 2) hour))] ∧
                 [John wait-for Mary]]
                ** e_2]))))]
    ** e_1])

(9.43′) ($\exists e_1$: [e_1 same-time *Now1*]
    [[Speaker tell Hearer (That
        ($\exists e_2$: [e_2 before e_1]
            [[[[John | e_2] for-benef Mary] ∧ ($\exists x$: [x rose] [John buy x])]
             ** e_2]))]
    ** e_1])

Note in (9.42′) that episode e_2 is explicitly qualified by its spatiotemporal predications, i.e., the location and time in which e_2 took place as well as its duration. In (9.43′), the ordered pair [John | e_2] is John's *action* of buying a rose for Mary. Also, note that action [John | e_2] is explicitly qualified by the description of its beneficiary, and this is part of the characterization of episode e_2.

9.3.2 Propositional Adverbials

I consider here only atomic modal predicates, such as *certain*, *probable* and *fortunate*, that are transformed into propositional operators by an application of **adv-p**. I also assume that propositional/modal operators always apply to atemporal (or, unlocated, to be precise) formulas. This means that tense operators are usually scoped below **adv-p** type adverbials. The following is a tentative deindexing rule for propositional adverbials.

adv-p: $((\text{adv-p}\ \pi)\ \Phi)_{\mathrm{T}} \leftrightarrow [(\text{That}\ \Phi_{\mathrm{T}})\ \pi]$,
for π a monadic predicate and Φ an unlocated formula

Tree transformation: $((\text{adv-p}\ \pi)\ \Phi) \cdot \mathrm{T} = \Phi \cdot (\pi \cdot \mathrm{T})$

I now show a sample sentence. (9.44 a) below involves episode- and action-modifying adverbials, in addition to a propositional adverbial; thus, it nicely illustrates the effect of the rules shown so far. Braces in (9.44 a) indicate assumed phrase structures.

(9.44) a. {John {*certainly* {{{slept *soundly*} *for eight hours*} *yesterday*}}}.
b. (decl ((adv-p certain)
(past ((adv-e (during *Yesterday*))
((adv-e (lasts (K ((num 8) hour))))
[John ((adv-a (in-manner sound)) sleep)])))))
c. ($\exists e_1$: [e_1 same-time Now_1]
[[**Speaker** tell **Hearer** (That
[(That ($\exists e_2$: [e_2 before e_1]
[[$^\wedge\lambda e$[e during (yesterday-rel-to Now_1)] $\wedge$
[$^\wedge\lambda e$[e lasts (K ((num 8) hour))] $\wedge$
[$^\wedge\lambda e$[[John | e] (in-manner sound)] $\wedge$ [John sleep]]]]
** e_2]))
certain])]
** e_1])

(Again the orienting relation is neglected.) This is easily verified. In (9.44 b), the **past** operator has been "raised" to some sentential level, right below the modal operator. Note in (9.44 c), the proposition headed by *That* is predicated to be "certain." Also note that *Yesterday* has been deindexed to (yesterday-rel-to Now_1). Applying to (9.44 c) MP 9.1 and MP 9.2 introduced earlier, we get (9.44′) shown below. Here, the ordered pair [John | e_2] is John's *action* (or activity) of sleeping. Again, e_2 is explicitly qualified by its temporal/durational predications and by the manner of the action of which it is a constituent.

(9.44′) $(\exists e_1{:}\,[e_1$ same-time $Now_1]$
[[Speaker tell Hearer (That
[(That $(\exists e_2{:}\,[e_2$ before $e_1]$
[[[e_2 during (yesterday-rel-to Now_1)] ∧
[e_2 lasts (K ((num 8) hour))] ∧
[[John | e_2] (in-manner sound)] ∧ [John sleep]]
** e_2]))
certain])]
** e_1])

9.3.3 Frequency and Cardinal Adverbials

A sentence modified by a frequency adverbial such as *regularly, frequently* and *repeatedly* describes a composite of two or more episodes of the same type that occur in sequence. Frequency adverbials tell how episodes are distributed within a sequence, i.e., their relative distance from each other, while cardinal adverbials express the number of occurrences of episodes of a certain type. As discussed in Chapter 6, frequency adverbials are translated into (adv-f π), e.g., (adv-f regular), (adv-f frequent), (adv-f repetitive), etc.; cardinal ones are translated into (adv-n π), where π is a predicate over collections of episodes, e.g., $\pi \in$ {((num 1) episode), ((num 2) episode), ...}. Note that an episode modified by a cardinal adverbial does not necessarily have more than one component episode.

I now show possible deindexing rules for such adverbials.

adv-f: $((\text{adv-f}\ \pi)\ \Phi)_T \leftrightarrow [\hat{}((\text{attr}\ \pi)\ \text{multi-component-ep}) \wedge (\text{mult}\ \Phi_T)]$,
for π a monadic predicate and Φ a formula

Tree transformation: $((\text{adv-f}\ \pi)\ \Phi) \cdot T = \Phi \cdot T$

adv-n: $((\text{adv-n}\ \pi)\ \Phi)_T \leftrightarrow [\hat{}(\text{composite-of}\ (\pi\ \text{episode})) \wedge (\text{mult}\ \Phi_T)]$,
for π a monadic predicate and Φ a formula

Tree transformation: $((\text{adv-n}\ \pi)\ \Phi) \cdot T = \Phi \cdot T$

A *multi-component episode* is an episode whose temporal projection is a multi-interval (thus, for instance, a sequence of two or more episodes may form a multi-component episode). composite-of in adv-n-rule is a predicate modifier that uniformly maps predicates over collections of episodes into episodes. If the collection contains more than one episode, its composite episodes may form a multi-component-episode. mult is a function that maps sentence intensions into sentence intensions, i.e., an operator of type $(\mathcal{S}\to 2)\to(\mathcal{S}\to 2)$, and defined as follows.

For η an episode, π a predicates, and Φ a formula,
[(mult Φ) ** η] $\leftrightarrow$ ($\forall e$: [e component-of η] [Φ ** e]),

where for $s, s', s'' \in \mathcal{S}$, [$s$ component-of s''] iff $s \sqsubseteq s''$, *Clocktime*(s) is an interval, and for all s' such that $s \sqsubseteq s' \sqsubseteq s''$, if *Clocktime*($s'$) is an interval then $s = s'$.

Here is a sample sentence involving a frequency adverbial.

(9.45) a. John called Mary regularly for two weeks.
b. (decl (past ((adv-e (lasts (K1 ((num 2) week))))
((adv-f regular) [John call Mary]))))
c. ($\exists e_1$: [e_1 same-time *Now1*]
[[Speaker tell Hearer (That
($\exists e_2$: [e_2 before e_1]
[[^(lasts (K1 ((num 2) week)))
[^((attr regular) multi-component-ep) $\wedge$ (mult [John call Mary])]]
** e_2]))]
** e_1])

With meaning postulates applied to the nonindexical ELF (9.45 c), it yields (using arbitrary variable indices):

(9.45′) ($\exists e_1$: [e_1 same-time *Now1*]
[[Speaker tell Hearer (That
($\exists e_2$: [e_2 before e_1]
[[[e_2 lasts (K1 ((num 2) week))] $\wedge$
[e_2 ((attr regular) multi-component-ep)] $\wedge$
(mult [John call Mary])]
** e_2]))]
** e_1])

The following is an example with a cardinal adverbial.

(9.46) a. John dated Mary three times in two months.
b. (past ((adv-e (spans (K1 ((num 2) month))))
((adv-n ((num 3) episode)) [John date Mary])))
c. ($\exists e_1$: [e_1 same-time *Now1*]
[[Speaker tell Hearer (That
($\exists e_2$: [e_2 before e_1]
[[^(spans (K1 ((num 2) month)))
[^(composite-of ((num 3) episode)) $\wedge$ (mult [John date Mary])]]
** e_2]))]
** e_1])

Again, by meaning postulates, we get from (9.46c) the following:

(9.46′) $(\exists e_1{:}[e_1$ same-time *Now1*]
[[Speaker tell Hearer (That
$(\exists e_2{:}[e_2$ before $e_1]$
[[[e_2 spans (K1 ((num 2) month))] ∧
[e_2 (composite-of ((num 3) episode))] ∧ (mult [John date Mary])]
** e_2]))]
** e_1])

We have also experimented with rules for frequency adverbs such as *often*, but in general these encounter difficult problems in the semantics of *generic* sentences [Schubert and Pelletier, 1989]. We have also been working on adverbials other than adverbs and PPs, but it is largely beyond of the scope of this paper.

9.4 Beyond Sentence Pairs: Extending Tense Trees

As a final topic, I discuss possible extensions of the deindexing mechanism to handle discourse involving shifts in temporal perspective. We saw in Chapter 7 that the tense tree mechanism, and particularly the way in which it automatically supplies orienting relations, is well suited for longer narratives, including ones with tense shifts — recall, for instance, (7.1)–(7.4). This is not to say that the tense tree mechanism obviates the need for larger-scale discourse structures. For instance, as has been pointed out by Webber [1987a; 1988], many subnarratives introduced by a past perfect sentence may continue in simple past. The following is one of Webber's examples:

(9.47) a. I was$_{\{e_1\}}$ at Mary's house yesterday.
b. We talked$_{\{e_2\}}$ about her sister Jane.
c. She had$_{\{e_3\}}$ spent$_{\{e_4\}}$ five weeks in Alaska with two friends.
d. Together, they *climbed*$_{\{e_5\}}$ Mt. McKinley.
e. Mary asked$_{\{e_6\}}$ whether I would want to go to Alaska some time.

Note the shift to simple past in *d*, though as Webber points out, past perfect *could* have been used. The abandonment of the past perfect in favor of simple past signals the temporary abandonment of a perspective anchored in the main narrative - thus bringing readers "closer" to the scene (a *zoom-in* effect). In such cases, the tense tree mechanism, unaided by a notion of higher-level discourse segment structure, would derive incorrect temporal relations such as [e_5 orients e_6] or [e_6 right-after e_5].

I now show possible deindexing rules for perspective shifts, assuming for now that such shifts are independently identifiable, so that they can be incorporated into the indexical

LFs. *new-pers* is a sentence *operator* initiating a perspective shift for its operand, and *prev-pers* is a *sentence* (with otherwise no content) which gets back to the previous perspective. $\texttt{Recent}_T$ is the episode most recently stored in the subtree immediately embedded by the focal node of T.

New-pers: $(\text{new-pers } \Phi)_T \leftrightarrow [\Phi_{\mapsto T} \wedge [\texttt{Recent}_T \text{ orients } \texttt{Recent}_{T'}]]$ where $T' = \Phi \cdot (\mapsto T)$

Tree transformation: $(\text{new-pers } \Phi) \cdot T = \Phi \cdot (\mapsto T)$

Prev-pers: $\text{prev-pers}_T \leftrightarrow \top \text{ (True)}$

Tree transformation: $\text{prev-pers} \cdot T = \not\leftarrow T$

When *new-pers* is encountered, a new tree is created and embedded at the focal node (it will be a node at which *utterance* episodes are stored), the focus is moved to the root node of the new tree, and the next sentence is processed in that context. In contrast with other operators, *new-pers* causes an overall focus shift to the new tree, rather than returning the focus to the original root. Note that the predication $[\texttt{Recent}_T \text{ orients } \texttt{Recent}_{T'}]$ connects an episode of the new sentence with an episode of the previous sentence. *prev-pers* produces a trivial *True*, but it returns the focus to the embedding tree, simultaneously blocking the link between the embedding and the embedded tree (as emphasized by use of $\not\leftarrow$ instead of $\leftarrow$). Note the similarity of this technique to the caching method proposed by Webber [1988].

I now illustrate how tense trees get modified over perspective changes, using (9.47) as example. I repeat (9.47d,e) below as (9.47′d,e), augmenting them with perspective changes, and show snapshots of the tense trees at the points marked (Figure 9.1). In the trees, $u_1, \ldots, u_5$ are utterance episodes for sentences a, ..., e, respectively.

(9.47′) d. $\uparrow_{T_1}$ (*new-pers* Together, they climbed$_{\{e_5\}}$ Mt. McKinley.) $\uparrow_{T_2}$
prev-pers $\uparrow_{T_3}$
e. Mary asked$_{\{e_6\}}$ whether I would want to go to Alaska some time. $\uparrow_{T_4}$

Notice the blocked links to the embedded tree in T_3 and T_4. Also, note that $\texttt{Recent}_{T_1}$ = e4 and $\texttt{Recent}_{T_2}$ = e5. So, by **New-pers**, we get [e4 orients e5], which can be later particularized to [e5 during e4]. It is fairly obvious that the placement of *new-pers* and *prev-pers* operators is fully determined by discourse segment boundaries (though not in general coinciding with them). So, as long as the higher-level discourse segment structure is known, our perspective rules are easily applied. In that sense, the higher-level structure supplements the "fine structure" in a crucial way.

However, this leaves us with a serious problem: deindexing and the context change it induces is supposed to be *independent* of "plausible inferencing"; in fact, it is intended

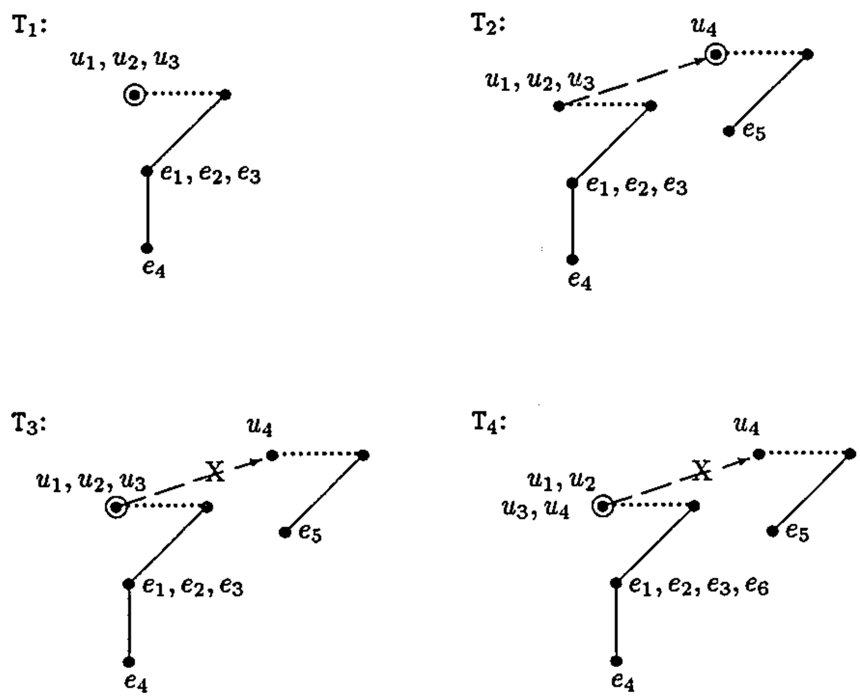

Figure 9.1: Snapshots of Tense Trees (with Perspective Shifts)

to *set the stage* for the latter. Yet the determination of higher-level discourse structure—and hence of perspective shifts—is unquestionably a matter of plausible inference. For example, if *past perfect* is followed by *past*, this could signal either a new perspective within the current segment (see (9.47 c, d)), or the closing of the current subsegment with no perspective shift. If *past* is followed by *past*, we may have either a continuation of the current perspective and segment (see (9.48 a, b) below), or a perspective shift with opening of a new segment (see (9.48 b, c)), or closing of the current segment, with resumption of the previous perspective (see (9.48 c, d)).

(9.48) a. Mary found that her favorite vase was broken.
b. She was upset.
c. She bought it at a special antique auction, and
d. she was afraid she wouldn't be able to find anything that beautiful again.

Only plausible inference can resolve these ambiguities. This inference process will interact with resolution of anaphora and introduction of new individuals, identification of spatial and temporal frames, the presence of modal/cognition/ perception verbs, and most of all will depend on world knowledge. In (9.48), for instance, one may have to rely on the knowledge that one normally would not buy broken things, or that one does not buy things one already owns.

As approaches to this general difficulty, we may think of the following two strategies: (A) Make a best initial guess about presence or absence of *new-pers/prev-pres*, based on *surface* (syntactic) cues and then use failure-driven backtracking if the resulting interpretation is incoherent. A serious disadvantage would be lack of integration with other forms of disambiguation. (B) Change the interpretation of $\texttt{Last}_\texttt{T}$, in effect providing multiple alternative referents for the first argument of *orients*. In particular, we might use

$$\texttt{Last}_\texttt{T} = \{e_i \mid e_i \text{ is the last-stored episode at the focus of T, or was stored in the subtree rooted at the focus of T } \textit{after} \text{ the last-stored episode at the focus of T}\}.$$

Subsequent processing would resemble anaphora disambiguation. In the course of further interpreting the deindexed LF, plausible inference would particularize the schematic orienting relation to a temporal (or causal, etc.) relation involving just two episodes. The result would then be used to make certain structural changes to the tense tree (*after* LF deindexing).

For instance, suppose such a schematic orienting relation is computed for a simple past sentence following a past perfect sentence (like (9.47 c,d)). Suppose further that the most coherent interpretation of the second sentence (i.e., (9.47 d)) is one that disambiguates the orienting relation as a simple temporal inclusion relation between the successively reported events. One might then move the event token for the second event (reported in simple past) from its position at the past node to the rightmost position at the past perfect node, just as if the second event had been reported in the past perfect. (One might in addition record a perspective shift, if this is still considered useful.) In other words, we would "repair" the distortion of the tense tree brought about by the speaker's "lazy" use of simple past in place of past perfect. Then we would continue as before.

In both strategies we have assumed a general coherence-seeking plausible inference process. While it is clear that the attainment of coherence entails delineation of discourse segment structure and of all relevant temporal relations, it remains unclear in which direction the information flows. Are there independent principles of discourse and temporal structure operating *above* the level of syntax and LF, *guiding* the achievement of full understanding, or are higher-level discourse and temporal relations a mere byproduct of full understanding? As noted in Chapter 7, Webber [1987a] proposed independent temporal focusing principles similar to those in [Grosz and Sidner, 1986] for discourse. Song and Cohen [1991] added heuristics for making the extraction of temporal constraints deterministic. However, since the sentences considered were relatively simple, and the heuristics are subject to counterexamples, their work still open the question about independent structural principles above the level of syntax and LF. The recent work by Lascarides *et al.* [Lascarides and Asher, 1991; Lascarides and Oberlander, 1992] based on combined discourse principles and world knowledge (as discussed in Chapter 8) seems to be a promising approach, and suggests that there are indeed higher-level structural principles of work; i.e., the informational flows *both* ways. However, a final verdict must await a more complete integration with theories of logical form and plausible inference.

9.5 Conclusions

In this chapter I proposed rules for deindexing some complex expressions, in particular, those involving logical operators, infinitives, and adverbials. Though most of the proposed rules are still tentative, they seem to indicate the potential of tense trees as the fine structure of discourse. Future work will focus on extension of the rules so as to cover relative clauses, clausal adverbials and generic tenses.

Chapter 10

Application to General NLU

In Part I of this thesis, I developed Episodic Logic (EL) as a SR/KR for NLU. I also emphasized that episodic logical form (ELF) representations allow for very direct, effective inference; at that point, I made no attempt to illustrate this claim with a nontrivial example. Having outlined in Part II of the thesis the process by which logical forms can be derived, I now return to the inference theme. The main concern in this chapter is to provide a nontrivial illustration of story inference, so as to demonstrate the practical potential of EL as a basis for inference in general NLU systems. I first discuss what kinds of inferences are involved in "understanding" natural language, and then show how some of these inferences are carried out in EPILOG, by way of a sample run.

10.1 Inference and Understanding

In Section 6.1, I outlined our view of the general understanding process (see Figure 6.1). However, my discussion so far has left several gaps with respect to that view, most notably stage V, i.e., ampliative inference. It is a major remaining research issue to specify precisely what happens in this stage. Here I provide only the barest outline, before zeroing in on one of its several aspects.

Ampliative inference is thought of as involving the following (interleaved) processes. First, inference rules are applied to the deindexed translation of the input, in combination with stored knowledge (meaning postulates, probabilistic conditionals, and other general and specific knowledge); this may generate, among other things, new *predictions* and *explanations*, and these may in turn trigger further inferences. This *chaining* is the aspect I will focus on in this chapter.

The second kind of process we see as part of ampliative inference is one that can be thought of as "implicit question answering": a text (or discourse) *raises certain questions*, and new inputs are preferentially interpreted so as to answer these implicit questions. We identify "raising a question" with inferring a prediction or explanation (more or less certainly) by inference chaining. The question raised is "answered" when the next input sentence, or one of its (more or less certain) consequences via inference chaining, is found to be supported by a positive or negative answer to the question. Here is a summary statement of the positive case, which I will illustrate with a simple example.

Implicit Question-Answering Principle A. Suppose Φ is an uncertain inference obtained (through inference chaining) from the text just seen, and Φ supports an inference Ψ obtained from a new input clause. Then Φ is strongly confirmed, and so are the intermediate conclusions in the inference chain which led to Φ, and from Φ to Ψ. The following simple example will make this more concrete.

(10.1) a. John dropped the glass on the floor.
b. It broke.

Inference chaining based on (10.1 a), along with axioms about "dropping," "glasses," and about fragile objects striking hard surfaces would quickly lead to a rather probable prediction that the glass broke. In our terminology, therefore, the question of *whether the glass broke* is raised. Now in (10.1 b), as long as the pronoun is unresolved, we cannot take this as implying that *the glass* broke. However, it does imply that *something* broke.[1] And this inference (playing the role of Ψ in the *Implicit Question-Answering Principle*) is indeed supported by a positive answer to the "question," "Did the glass break?". Thereby the predicted breakage of the glass is strongly confirmed. Note also that a tentative inference from the orienting relation computed for (10.1 a) and (10.1 b), namely, that the breaking was right after, and caused by, the dropping will be similarly confirmed (assuming that such an immediately following causal consequence was predicted from (10.1 a)). Even more importantly, the pronoun will be resolved, since the (probabilistic) "proof" of Ψ ("something broke") from (10.1 a) requires that "something" that broke be identified with the glass. (In theorem proving terms, the variable in the goal of the proof of "Something broke" is *unified* with the constant or other term denoting the glass.)

So far this will sound rather familiar. What is happening, in effect, is that inferences from successive sentences are being "matched" and unified. This is quite similar to what would happen in an ACT-based (or script or MOP-based, etc.) understanding system, where the predictions implicit in an evoked script are matched against subsequent inputs (e.g., MARGIE [Schank *et al.*, 1975] and SAM [Cullingford, 1981]). Also, this view of interpretation is closely related to the abduction approaches of Charniak and Goldman

[1] We might say this follows by existential generalization of the parameter 'it'; or we might take the view that the translation of the pronoun is existential in the first place, amounting to "some salient neuter entity," or "some neuter entity on the focus list in the current context," or something similar.

[Charniak, 1988] [Charniak and Goldman, 1988, 1989] and Hobbs *et al.* [1988, 1990], in which a new input is interpreted so that it is derivable from what is already known with a minimal set of supplementary assumptions.

However, the other half of our proposed principle is that the *denial* of a prior inference (a "negative answer" to the question raised) can play the same role in determining the interpretation of new material as its affirmation. In this respect our proposal seems quite different from previous ones. Here is the statement of the second part, followed by a variant of the previous illustration.

Implicit Question-Answering Principle B. Suppose $\neg\Phi$ is the denial of an uncertain inference Φ obtained (through inference chaining) from the text just seen, and $\neg\Phi$ supports an inference Ψ obtained from a new input clause. Then $\neg\Phi$ is strongly confirmed, and so is the denial of the conjunction of all the intermediate conclusions in the inference chain which led up to Φ, and the (affirmation of the) intermediate conclusions in the inference from $\neg\Phi$ to Ψ. A suitable illustration is (10.1 a) plus the denial of (10.1 b):

(10.2) a. John dropped the glass on the floor.
b. It didn't break.

In this case, it is the *denial* of a predication from (10.2 a) that the glass broke, i.e., that the glass *didn't* break, that supports the inference from (10.2 b) that *something* didn't break. Furthermore, by the same unification process as before, the pronoun is again resolved to the glass. By contrast, approaches like those of Charniak and Goldman and Hobbs *et al.* which insist on interpreting new inputs as logically supported by prior inputs (and background knowledge) would get the wrong interpretation here. In particular, since (10.2 a) certainly supports the conclusion that *the floor* didn't break, the pronoun would be resolved to refer to the floor.[2]

The underlying idea in the *Implicit Question-Answering Principles* is that the interpretation of sentences is determined not just by world knowledge, but by narrative or discourse conventions. A narrative *does* raise questions in the reader's or hearer's mind; and the narrator is under some (mild) obligation to answer them at least implicitly. The key point, though, is that she is free to answer them either positively or negatively, regardless of how unexpected the answer may be. In fact, a story in which all expectations are confirmed would be utterly uninteresting.

As a further (sketchy) illustration, consider

[2]It might be countered that the resolution of the pronoun in both (10.1) and (10.2) is the result of "centering," where the verb object in (a) is the preferred center. However, this is disconfirmed by "John dropped the brick on glass. It didn't break."

(10.3) a. Mary heard steps behind her.
b. She began to run.

A spontaneous explanatory inference from (10.3 a) is likely to be that there was someone behind Mary, quite close to her (and she knew this). In turn, this leads to the (possibly very tentative) conclusion that Mary may believe herself in danger, and may try to get away from the person behind her. (Of course, prior context may disable this inference, and in that case (10.3 b) may be quite differently interpreted than for (10.3) in isolation. Also, the inference of danger from (10.3 a) seems to have something to do with expectations based on what typically happens in stories, as opposed to world-experience. But that is not the issue here.)

Now (10.3 b) also leads to the inference of possible explanations, though intuitively there are several *alternative* explanations for the truth of the sentence, taken in isolation: "she" may be in a hurry, may be trying to get away from someone or something near her, or may be exercising. (These seem like the most probable explanations.) Once again, we can assume that the pronoun is interpreted existentially, although in this case context provides a unique referent; i.e., some contextually salient female (who can only be Mary here) is in a hurry, or trying to get away from someone or something near her, or exercising. And once again, we find that this disjunctive conclusion is supported by a positive answer to the "question" raised by (10.3 a), as to whether Mary is trying to get away from someone behind her. (The logical principle here is that $\Phi_1 \models \Phi_1 \vee \Phi_2 \vee \Phi_3$.) Thus by the *Implicit Question-Answering Principles*, the positive answer is confirmed, and as a byproduct, the identity between "she" and "Mary" is further confirmed and in addition, the indefinite (existentially quantified) individual behind Mary inferred from (10.3 a) is unified with the indefinite "someone or something" conjectured in the disjunctive explanation for (10.3 b).

Of course, such inferential connections will not always be found. For instance, "John greeted Mary. Mary was startled" (used to illustrate an axiom of causal connection in Chapter 8) seems to acquire its coherence simply from the way the sentences are sequenced, not from any "implicit question answering" — that was the point of the example.

The method of interpretation based on implicit question answering may be a potential alternative to the method of Lascarides *et al.* discussed in Chapter 8, or perhaps unifiable with that approach. Let us once again — very briefly — consider an example from that discussion:

(a) Guy experienced a lovely evening last night.
(b) He had a fantastic meal.
(c) He ate salmon.
(d) He devoured lots of cheese.
(e) He won a dancing competition.

In the proposed approach, meaning postulates and world knowledge applied to (a) would lead to the conclusion that Guy was a participant in some situations or events last night which gave him pleasure. This "raises the question" of whether indeed there were such events. When (b) is processed, the application of MPs and world knowledge would supply the conclusion that the meal gave Guy pleasure, and also (given the meaning of "meal") that he ate some foods. The first of these conclusions partially answers the question raised by (a), unifying the "fantastic meal" with one of the pleasurable events inferred from (a). The second conclusion raises the question of whether indeed Guy ate some foods, and it is clear that (c) and (d) will provide answers to this question. Thus, (c) and (d) are linked to (b). Finally, (e) (and inferences based on it) are unlikely to answer any questions raised by (b), (c) or (d); but it *does* supply a further answer to the question raised by (a) about the occurrence of pleasurable events, assuming that winning a competition can be inferred to be pleasurable.

This is of course extremely sketchy, and leaves open numerous questions, such as how distant a clause answering a question may be from the clause raising it, and under what conditions a question raised by a clause admits answers from multiple clauses following it (thus leading to an elaboration-like discourse structure). But the sketch seems sufficiently plausible to be worthy of further investigation. Such further investigation should also clarify the relationships to the approach of Lascarides *et al.*, and the possibility of amalgamating these approaches.

I leave the discussion of "implicit question-answering" here. Admittedly, the discussion has barely scratched the surface of this aspect of ampliative inference, an aspect that constitutes a large part of the "coherent interpretation" problem. However, developing the ideas sketched above into a convincing and computationally explicit theory would certainly require another thesis. The emphasis in the present thesis is on representation, on the relation between different levels of logical form, and on what inferences are possible in principle. How a coherent interpretation is selected from myriad possibilities goes well beyond these goals.

I now return to the first aspect of ampliative inference, (spontaneous) inference chaining. As already indicated, the links in inference chaining are often provided by meaning postulates (such as that dropping an object entails that it falls), by causal axioms (such as that a fragile object striking a hard surface will probably break *as a result*), and by explanatory axioms (such as that a person starting to run may be in a hurry, may be trying to get away from someone or something, or may be exercising).

However, stories are often not just about physical events, but also about what goes on in people's minds, i.e., about mental events and processes. Now it seems that the easiest and most natural way to think about someone else's thinking is to try to *simulate* their thought processes, rather than reasoning purely axiomatically. The point is this: to *simulate* someone's thinking only requires that one *have* (and be able to "run") a mental apparatus similar to theirs. But to reason axiomatically about someone's thinking, one needs a detailed *theory* of their mental apparatus — a requirement extremely unlikely to be met. Therefore, it makes sense to assume that a story understander can make inferences about mental processes by simulation. For instance, to infer what someone else will infer upon learning Φ, we "pretend" to learn Φ ourselves (while also "pretending" to believe certain relevant propositions that the other individual can be assumed to believe) and simply wait for our spontaneous inference processes to run their course, delivering various conclusions. We then ascribe these same conclusions to the simulated thinker — though presumably only after checking whether the knowledge that participated in delivering the conclusions is likely to be shared with the simulated thinker.

So, we want to assume some special predicate for expressing the relation between "input information" (which we "pretend" to receive) and "output information" (the conclusions we spontaneously reach) in a mental simulation. Specifically, let

$[x \text{ would-infer-from } y\ z]$

mean that if I run my "mental simulator" with input information y, and with background information consisting of the beliefs I explicitly ascribe to x plus what I take to be shared general knowledge, then I spontaneously reach conclusion z. This predicate, then, would be evaluated for particular x, y, z not by derivation from other knowledge, but by "direct query" to the simulator. Its status is much like that of a perception predicate such as "I see such-and-such an object before me." This also would not be something I would *prove*, but rather something I would obtain by "asking" my visual perception system. Given the ability to evaluate such predicate instances, it becomes (logically) trivial to make an inference about other people's thinking. We can use a general axiom like the following.

Simulative Inference.

$$(\exists x\ (\exists y\ (\exists z\ (\exists e\ [[[x \text{ learn } y] ** e] \land [[x \text{ would-infer-from } y\ z] * (\text{time-of } e)]]))))$$
$$\rightarrow_{.8, x, y, z, e}\ (\exists e' {:} [e \text{ cause-of } e']\ [[x \text{ infer } z] ** e'])$$

Note that the `would-infer-from` predication, though it is just a matter of evaluation by simulation, does need to be temporally qualified, since the simulation involves the use of knowledge about the beliefs of x, and those beliefs will in general be time-dependent. In other words, the simulation needs to be run using x's presumed beliefs *at the time x learns* y (some proposition).

Though the above rule needs to be reformulated more carefully and with some generalization, e.g., to be able to handle nested beliefs, the point I want to make is that all kinds of axioms needed for narrative understanding can be uniformly represented in EL with probabilistic conditionals, including axioms involving predicates evaluated "introspectively." Though simulative reasoning has not so far been implemented, I will have a little more to say about it in the extended example in the next section.

10.2 EL and EPILOG in Narrative Understanding

EL has been put to the test in EPILOG [Schaeffer *et al.*, 1991], a prototype implementation of EL. The starting point of EPILOG was an existing, resolution-based semantic net system in ordinary first order logic called ECONET [de Haan and Schubert, 1986; Miller *et al.*, 1987]. ECONET supports efficient deduction (both general and specialized) and fast, selective access to knowledge relevant to a particular set of concepts and topics, with a sophisticated agenda-driven control structure for goal chaining, with goals ranked according to estimated difficulty and with accessing of new knowledge for use in a proof via concept and topic hierarchies. Its aim is not so much theorem proving power *per se*, but the ability to get at the *relevant* knowledge in a large knowledge base.

Several of these techniques developed for ECONET have been incorporated into EPILOG, a hybrid reasoning system combining efficient storage and access mechanism, forward and backward chaining, an agenda-driven control structure, and multiple "specialists" for taxonomies, temporal reasoning, etc. EPILOG handles the types of questions that were originally handled by ECONET, e.g., "Did anyone have some cake?" or "Does grandmother live in a shoe?" In addition, EPILOG handles new examples involving '**', λ-abstraction, etc., and performs many of the inferences alluded to so far including ones based on meaning postulates. Furthermore, based on its use of an expressively much richer, NL-like representation (i.e., EL), and its more powerful inference techniques (RI and GC with multiple instantiation, plus natural deduction), EPILOG is capable of making more subtle inferences and doing it more efficiently than ECONET. EPILOG is already making some quite complex inferences and answering questions based on logically represented simple narratives or telegraphic messages (cf., [Namioka *et al.*, 1991, 1992]).

10.2.1 Understanding a LRRH Story Fragment

As a way of testing the practical potential of EL for NLU, we ran on EPILOG a small fragment of the *Little Red Riding Hood* story shown below.

> *In the forest, Little Red Riding Hood met a wolf. The wolf would have very much liked to eat her, but he dared not do so on account of some woodcutters nearby.*

Fully processing this fragment requires extensive reasoning including inferences based on meaning postulates, predictive inferences, explanatory inferences and simulative inferences. For example, to understand the third sentence, one should be able to explain why the wolf decided against eating *Little Red Riding Hood* (hereafter, LRRH), and how the presence of woodcutters nearby affected the wolf's decision. So, one has to know that when some agent *dares* not do something, he must think it possible that his attempt to do it would result in something unpleasant to himself; then one has to simulate his reasoning process to guess what unpleasant consequences he anticipates.

Depending on the degree of sophistication of the knowledge possessed, people may explain the wolf's decision in various ways. Correspondingly, depending on the kind of knowledge provided, the inference machinery should be able to produce various lines of reasoning; this includes the following, relatively simple line of reasoning

- Attacking a child is extremely wicked.
- Trying to eat a living creature involves attacking it, and such an attack is conspicuous and likely to be noticed by nearby people.
- Doing something extremely wicked is likely to bring severe punishment, if noticed by anyone.
- So, if the wolf tries to eat LRRH, the nearby woodcutters may notice it, and he is likely to be severely punished for it.

or, the more sophisticated version

- When a predatory animal eats a non-predatory creature of comparable size while the creature is conscious, the predator attacks it as a preparation for eating it.
- The wolf would attack LRRH before eating her.
- Attacking a person is a conspicuous action, and is likely to be noticed by nearby people.
- If people notice a predatory animal attacking a person, they will most probably want to rescue the person from the animal.
- To rescue a person from a predatory animal, one may kill it.
- Thus, the woodcutters may kill the wolf.

Upon reaching a conclusion that it is possible that the wolf might be killed or severely punished, the inference machinery may attribute its own ability to infer that conclusion to the wolf (this could in principle be done using the simulative inference axiom shown earlier). Then it is easily explained why the wolf decided against eating LRRH right then and there.

EPILOG was run on a set of facts and general knowledge pertaining to this fragment with a new "input" that the wolf tries to eat LRRH. This was done on the basis that to simulate the wolf's reasoning, the story understander would at a certain point pursue the consequences of assuming that the wolf *does* try to eat LRRH. The results were quite gratifying: EPILOG computed the inferences we had obtained by hand-simulation that account for the wolf's decision not to eat LRRH right away when he first met her. I now outline the forward inference chain that was computed by EPILOG. The assumption was made that there were woodcutters nearby when the wolf met LRRH, and then inference chaining was triggered for the hypothesis that the wolf *would* try to eat LRRH right then and there. In the following I show that part of the reasoning process reaching the conclusion "The wolf may be severely punished." The inferences I show were generated by EPILOG in nearly the same order. For simplicity of exposition, various details, including most of the inferred temporal relationships, are omitted. (N.B. Since the current version of implementation does not use controlled variables in the probabilistic conditionals, there is a slight difference in the logical form used in EPILOG. But as the sample story does not have any sentences involving proportion problems, this is of no consequence.) The control structure is designed to systematically combine each new clause with relevant meaning postulates and other general knowledge. All of the inferences are based on the explicit, formalized rules of inference introduced earlier. Simulative and narrative inferences are not yet made, but are not needed in this portion of the reasoning process. After listing meaning postulates and world knowledge, I show the simplified logical translation of the story and a trace of the reasoning process.

Meaning Postulates

M1. *To walk, to attack someone, to try to do something, to die, etc., are types of actions.*

For Π an action predicate:
□ [(Ka Π) action-type]

An "action predicate" is an expression $(\pi\ \tau_1, \ldots, \tau_{n-1})$, where π is an n-adic atomic action predicate, $n \geq 1$, and $\tau_1, \ldots, \tau_{n-1}$ are terms.

M2 (A meaning postulate regarding actions/attributes). For Π an action predicate:

$$\square\ (\forall x\ (\forall e\ [[[x\ \Pi]\ **\ e]\ \leftrightarrow\ [[x \mid e]\ \text{instance-of}\ (\text{Ka}\ \Pi)]]))$$

For example,

[[John eat] ** E1] ↔[[John | E1] instance-of (Ka eat)].

Note that [John | E1] *is* an action, not just an arbitrary individual-episode pair, so that John *is* the agent of that action.

M3. *If there is a collection of things of some type, then there is a thing of that type which belongs to that collection (we regard collections as non-empty by definition).*

For Π a monadic predicate:
□ $(\forall x{:}[x\ (\text{plur}\ \Pi)]\ (\exists y{:}[y\ \text{in}\ x][y\ \Pi]))$

plur is a function that maps a predicate applicable to things into a predicate applicable to collections of things.

World Knowledge

K1. *For a creature to attack a child is extremely wicked.*

$(\exists x{:}[x\ \text{creature}]\ (\exists y{:}[y\ \text{child}]\ (\exists e\ [[x\ \text{attack}\ y]\ {**}\ e])))$
$\rightarrow_{.9,e}\ [[x \mid e]\ ((\text{-ly extreme})\ \text{wicked})]$

K2. *Trying to eat any living creature involves attacking it.*

$(\forall x{:}[[x\ \text{alive}] \wedge [x\ \text{creature}]]$
$[(\text{Ka}\ (\text{try}\ (\text{Ka}\ (\text{eat}\ x))))\ \text{involve}\ (\text{Ka}\ (\text{attack}\ x))])$

K3. *If one type of action involves another, then any creature doing an instance of the first will do an instance of the second during it.*

$(\forall a1{:}[a1\ \text{action-type}]$
$(\forall a2{:}[[a2\ \text{action-type}] \wedge [a1\ \text{involve}\ a2]]$
$(\forall x{:}[x\ \text{creature}]$
$(\forall e1{:}[[x \mid e1]\ \text{instance-of}\ a1]$
$(\exists e2{:}[e2\ \text{during}\ e1]\ [[x \mid e2]\ \text{instance-of}\ a2])))))$

K4. *For a sizable creature to attack a sizable thing is conspicuous (relative to a human observer).*

$(\exists x{:}[x\ \text{person}]$
$(\exists y{:}[[y\ \text{creature}] \wedge \neg [y\ \text{tiny-rel-to}\ x]]$
$(\exists z{:}[[z\ \text{creature}] \wedge \neg [z\ \text{tiny-rel-to}\ y]]$
$(\exists e\ [[y\ \text{attack}\ z]\ {**}\ e]))))$
$\rightarrow_{.9,x,e}\ [[y \mid e]\ \text{conspicuous-to}\ x]$

By contrast, for an ant to attack something would not be conspicuous to a human.

K5. *If a creature performs a conspicuous action within plain sight of a person, that person is likely to notice that action.*

$(\exists x{:}[x\ \text{creature}]\ (\exists y{:}[y\ \text{person}]$
$(\exists e1{:}[[x\ \text{within-plain-sight-of}\ y]\ {**}\ e1]$
$(\exists e2{:}[e2\ \text{during}\ e1]\ [[x \mid e2]\ \text{conspicuous-to}\ y]))))$
$\rightarrow_{.6,x,y,e1,e2}\ (\exists e3{:}[e3\ \text{during}\ e2]\ [[y\ \text{notice}\ [x \mid e2]]\ {**}\ e3])$

K6. *Doing something extremely wicked may bring severe punishment from some group of people, if noticed by anyone.*

$(\exists x\colon [x \text{ creature}]\ (\exists e1\colon [[x \mid e1]\ ((\text{-ly extreme})\ \text{wicked})]$
$\quad (\exists y\colon [y \text{ person}]\ (\exists e2\ [[y \text{ notice } [x \mid e1]] ** e2]))))$
$\rightarrow_{.3,x,y,e1,e2}\ (\exists z\colon [z\ (\text{plur person})]$
$\quad (\exists e3\colon [e2 \text{ cause-of } e3]\ [[z\ ((\text{-ly severe})\ (\text{punish } x))] ** e3]))$

K7. *A human is not tiny relative to a wolf, and vice versa.*

$(\forall x\colon [x \text{ human}]\ (\forall y\colon [y \text{ wolf}]\ [\neg[x \text{ tiny-rel-to } y] \wedge \neg[y \text{ tiny-rel-to } x]]))$

K8. *If a creature is near a person and not tiny relative to the person, it is probably within plain sight of the person.* (This could be improved by assuming that we are dealing with a *daytime* episode in an open setting.)

$(\exists x\colon [x \text{ person}]\ (\exists y\colon [[y \text{ creature}] \wedge \neg\, [y \text{ tiny-rel-to } x]]$
$\quad (\exists e1\ [[y \text{ near } x] ** e1])))$
$\rightarrow_{.6,e1}\ (\exists e2\colon [e2 \text{ same-time } e1]\ [[y \text{ within-plain-sight-of } x] ** e2])$

K9. *Woodcutters are humans.*

$(\forall x\colon [x \text{ woodcutter}]\ [x \text{ human}])$

Story Fragment

Let us now work out the possible consequences of the wolf's trying to eat LRRH. (We then would attribute this reasoning to the wolf if we were handling simulative inference.) The relevant assumptions and story facts are as follows. Surface speech acts are omitted in the logical form. I use the convention of having variables in lower case, and constants in upper case.

\# *The wolf tries to eat Little Red Riding Hood.*

(∃e1: [*Now1* during e1] (The x1: [x1 wolf][[x1 try (Ka (eat LRRH))] ** e1]))

By skolemization {E1/e1} and reference determination {W/x1}:

S1. [*now* during E1]
S2. [W wolf]
S3. [[W try (Ka (eat LRRH))] ** E1]

\# *Little Red Riding Hood is a girl and alive.*

S4. [LRRH girl]
S5. [LRRH alive]

There are woodcutters nearby.

(∃y1: [y1 (plur woodcutter)] (∀x: [x in y1] (∃e2: [E1 during e2] [[W near x] ** e2])))

By Skolemizing {C1/y1}:

S6. [C1 (plur woodcutter)]

S7. (∀x: [x in C1] (∃e2: [E1 during e2] [[W near x] ** e2]))

Assume the following type-hierarchical knowledge is available (at least indirectly, via a type "specialist"):

S8. [W creature]

S9. [LRRH child]

S10. [LRRH human]

S11. [LRRH creature]

Reasoning Process

- Note that simple time inferences such as

  [E1 during E2] ∧ [E2 during E3] ∧ [E3 same-time E4] ⊢ [E1 during E4]

  will be taken for granted during the inference process.

- In the following,

  RI [A; B] {Subst C/v; Imm-Skol C′/v′}

  indicates that the subsequent inference(s) has been made via RI of rule B by premise(s) A, with variable substitution C/v, and an existential variable v′ in the inferred formula has been immediately skolemized as C′.

RI [S5, S11; K2] {Subst LRRH/x}:

1. [(Ka (try (Ka (eat LRRH)))) involve (Ka (attack LRRH))]

   "Trying to eat LRRH involves attacking her."

RI [; M1] {Subst (try (Ka (eat LRRH)))/Π}:

2. [(Ka (try (Ka (eat LRRH)))) action-type]

   "Trying to eat LRRH is an action type."

RI [; M1] {Subst (attack LRRH)/Π}:

3. [(Ka (attack LRRH)) action-type]

   "Attacking LRRH is an action type."

RI [S3; M2] {Subst W/x, E1/e, (try (Ka (eat LRRH)))/Π}:

4. [[W | E1] instance-of (Ka (try (Ka (eat LRRH))))]

   "The wolf's trying to eat LRRH is an instance of someone's trying to eat LRRH."

RI [2, 3, 1, S8, 4; K3]

{Subst (Ka (try (Ka (eat LRRH))))/a1, (Ka (attack LRRH))/a2, W/x, E1/e1; Imm-Skol E2/e2}:

5. [E2 during E1]
6. [[W | E2] instance-of (Ka (attack LRRH))]

RI [6; M2] {Subst W/x, E2/e, (attack LRRH)/Π}:

7. [[W attack LRRH] ** E2]
   "The wolf attacks LRRH."[3]

RI [S8, S9, 7; K1] {Subst W/x, LRRH/y, E2/e}:

8. [[W | E2] ((-ly extreme) wicked)].[9]

\# Up to here:

`The wolf attacks LRRH, and that's extremely wicked.`

RI [S10, S2; K7] {Subst LRRH/x, W/y}:

9. ¬[LRRH tiny-rel-to W]
   "LRRH is not tiny relative to the wolf."

RI [S6; M3] {Subst C1/x, woodcutter/Π; Imm-Skol C2/y}:

10. [C2 in C1]
11. [C2 woodcutter]
    "There is a woodcutter."

---

[3]Inferences 5 and 7 could be obtained in one step rather than six by using the following knowledge K2′ instead of K2.

K2′. When a creature tries to eat a creature that is alive, he attacks it during that episode.
($\exists x$: [x creature] ($\exists y$: [[y alive] $\wedge$ [y creature]] ($\exists e1$ [[x try (Ka (eat y))] ** $e1$])))
$\rightarrow$ ($\exists e2$: [$e2$ during $e1$] [x attack y] ** $e2$])

Specifically,

RI [S8, S5, S11, S3; K2′] {Subst W/x, LRRH/y, E1/e1; Imm-Skol E2/e2}:

5. [E2 during E1]
7. [[W attack LRRH] ** E2]

However, our aim is to obtain the desired inferences in narrative understanding from *any* reasonable, intuitively natural way of formulating the relevant world knowledge. K2 is probably more natural than K2′, and more importantly, was written down *prior* to detailed consideration of the reasoning process it was intended to support. If we are going to have a robust system whose knowledge base and range of understanding is easily expanded, we cannot afford to "tailor" the syntactic form of the axioms to the inference chains we choose as examples.

RI [11; K9] {Subst C2/x}:

12. [C2 human]
"The woodcutter is a human."

With type-hierarchical knowledge, we get from 12:

13. [C2 person]

RI [12, S2; K7] {Subst C2/x, W/y}:

14. ¬[W tiny-rel-to C2]
"The wolf is not tiny relative to the woodcutter."

RI [13, S8, 14, S11, 9, 7; K4] {Subst C2/x, W/y, LRRH/z, E2/e}:

15. [[W | E2] conspicuous-to C2]$^{.9}$

\# Up to here:
`The wolf's attack is conspicuous to the woodcutter.`

RI [10; S7] {Subst C2/x; Imm-Skol E3/e}:

16. [E1 during E3]
17. [[W near C2] ** E3]
"The wolf is near the woodcutter (when he tries to eat LRRH)."

RI [13, S8, 14, 17; K8] {Subst C2/x, W/y, E3/e1; Imm-Skol E4/e2}:

18. [E4 same-time E3]
19. [[W within-plain-sight-of C2] ** E4]$^{.6}$
"The wolf is likely to be within plain sight of the woodcutter."

RI [S8, 13, 19, (5, 16, 18), 15; K5] {Subst W/x, C2/y, E4/e1, E2/e2;
Imm-Skol E5/e3}:

20. [E5 during E2]
21. [[C2 notice [W | E2]] ** E5]$^{.324}$

\# Up to here:
`The woodcutter may notice the wolf's attacking` LRRH.

RI [S8, 8, 13, 21; K6] {Subst W/x, E2/e1, C2/y, E5/e2; Imm-Skol C3/z, E6/e3}:

22. [C3 (plur person)]
23. [E5 cause-of E6]
24. [[C3 ((-ly severe) (punish W))] ** E6]$^{.087}$

\# `The wolf may be severely punished by some group of people.`

This inference chain could be extended to provide an explanation for the wolf's decision *not* to try to eat LRRH at that point in the story. First, rule K3 would be slightly augmented so as to express the fact that if one action involves another, and that other action has certain consequences, then these are also consequences of the first action. Rule K5 would be similarly augmented to make the "noticing episode" e3 a causal consequence of the episode e2 (or action [x | e2]) noticed. *The "punishing episode,"* E6, *in conclusion 24 would then be inferred to be a consequence of the wolf's attempt to eat* LRRH. Given that being severely punished is very bad, and that agents generally refrain from actions that they think may have very bad consequences for them, we would have an explanation for the wolf's restraint. Note, however, that this requires application of simulative inference axioms of the kind shown in the previous section, i.e., we must attribute the above inference chain to the wolf, and draw further conclusions from this attribution.

10.3 Summary

I briefly discussed possible ways of making ampliative and simulative inferences that are essential to understanding a narrative. I emphasized that axioms even for the latter kinds of inferences can be represented in EL as probabilistic conditionals, showing a sample simulative inference axiom. I also mentioned that most kinds of inferences alluded to throughout this thesis are actually carried out in EPILOG, a prototype implementation of EL, and gave an extensive example illustrating EPILOG's inference process. I hope to have shown through this example the practical potential of EL for general NLU systems and that despite its rich syntax — or perhaps thanks to it — EL allows for effective, compact (in terms of number of steps) and arbitrarily subtle reasoning. As well, there are reasons for optimism about the *efficiency* with which probable consequences of given inputs will be deduced, and answers to questions found, even when the relevant facts and axioms are embedded within a much larger knowledge base. In the first place, as I have pointed out, the EPILOG inference strategy (in QA mode) is derived from that of the earlier ECONET system, which proved very robust in the face of increased KB size and indeed was designed with this as goal [de Haan and Schubert, 1986]. EPILOG has been extensively (and successfully) tested on the ECONET examples. As well, like the earlier system, EPILOG is supported by an array of extremely efficient taxonomic, partonomic, temporal, set-theoretic, and other specialists. Furthermore, given that an application of an EL inference rule often packs many "ordinary" inference steps into one, the combinatorics of search are much less severe than in the earlier system. Finally, as an "anecdotal" indicator of EPILOG efficiency, it is worth mentioning that EPILOG solved the well-known "Steamroller" challenge problem [Stickel, 1985] (in about 2 1/2 minutes on a Sun 360), even though it is designed for drawing intuitively "obvious" conclusions, not difficult ones requiring lengthy proofs.

Chapter 11

Conclusions and Future Work

11.1 What's been Achieved?

At the beginning of this thesis, I argued that it is time to "put it all together," combining what we have learned in various subareas of language understanding, toward building a general NLU system, and called for creative synthesis of all aspects of NLU. From that perspective, I enumerated what I thought to be desiderata for a general NLU system, namely,

- It should have an *expressive* SR/KR.
- The SR/KR should be easily *computable* from English surface structure.
- The SR/KR should be direct, transparent, conceptually modular, and formally *interpretable*.
- The SR/KR should allow the system to *reason* efficiently and soundly about agents' beliefs, desires, goals, plans, etc.
- The SR/KR should meet the "interlocking" needs of all these.

Thus, the main concern of this thesis research has been to develop such an adequate SR and KR for a general NLU system. Given this extremely ambitious scope, I have necessarily had to trade off theoretical rigor and depth against comprehensiveness. In particular, from the practical view point of building working NLU systems, I believe it would be a mistake to cut back on expressive power for the sake of theoretical tractability. Thus, the strategy taken in this research was a top-down one, using a target representation with more or less full coverage of English from the outset, and subjecting this representation to continual revision in the light of the interlocking needs of grammar, computation of logical form, formal semantics, and inference. Future work is expected to lead to further refinement and deepening of the theoretical basis of EL.

The main outcome of this research is then two-fold: the development of EL as very expressive SR/KR and the development of a new method of deindexing using a new type of dynamic context structure called tense trees. The EL representations obtained via the deindexing mechanism do largely meet the desiderata set out above. EL combines ideas from Montague grammar, situation semantics, DRT, and natural language interpretation as understood in AI, and adds a number of new ideas concerning the semantics of situations, actions, propositions and facts, times, quantification and tense and aspect. The deindexed episodic logical formulas obtained via tense trees can be used for inference by methods similar to those familiar in AI. EL has been implemented and tested on realistic, though small, text samples. The results so far is encouraging, suggesting that it is indeed possible to grapple simultaneously with a wide spectrum of problems in natural language understanding.

I now list as specific contributions of this thesis the following features of EL and experience with its implementation in EPILOG.

(a) EL is an *expressive* SR/KR— it allows the content of most English sentences and most world knowledge to be represented in an intuitively comprehensible and formally analyzable manner. It makes implicit time and situation dependencies explicit through the use of episodic variables, and admits unbound anaphoric variables and the representation of generic conditionals, as well as restricted quantifiers, modal operators, and nominalization operators. Most significantly, all these have been brought together for the first time in a logic for narrative understanding.

(b) The representation of phrase structure is modular and transparent, as is the mapping from phrase structure to EL. The mapping handles many combinations of tense, aspect and adverbials. Although we do not have a complete grammar, and a full-fledged version of an ELF generator is not available yet, the compositional approach of EL has proven practical by a start-up implementation that has successfully handled dialogues in a planning system in the TRAINS domain [Allen and Schubert, 1991].

(c) In the process of developing EL, a new temporal deindexing mechanism has emerged that consists of a new type of dynamic context structure called tense trees and a small set of recursive rules. The mechanism not only interprets English tense and aspect, and their interactions with negation, quantifiers, and time adverbials, but is capable of automatically identifying orienting and referent episodes.

(d) The rules of inference in EL, RI and GC, though not yet completely analyzed, are explicit and amenable to formal analysis. Being probabilistic, they allow evidence for explanations or predictions to be *weighed*, much as in expert systems. Furthermore, these rules have been implemented in EPILOG, which already makes quite complicated inferences.

(e) EL allows in implementation all types of linguistic and domain knowledge to be strictly separated from parsing and inference control structure, allowing the former to be expanded and revised independently of the latter.

(f) Experience with processing of actual story fragments and question-answering indicates that the logical framework of EL is epistemologically adequate for story understanding.[1]

11.2 What's Left?

I hope to have provided evidence that EL can provide a comprehensive and unusually clean foundation for story understanding. However, building a complete framework for a general NLU obviously cannot be achieved within one thesis project, and much work remains to be done on all aspects of the logical approach proposed. I hereby list some pointers to future work.

Probability handling: One of the most important remaining problem is the principled handling of probabilities. The state of the art in probabilistic inference (e.g., [Pearl, 1988; Bacchus, 1990]) is not such as to provide concrete technical tools for a logic as general as EL. For instance, the probabilistic constructs and inferences have not been fully formalized yet, though in implementation a "noncircularity principle" has been successfully used which prevents the same knowledge from being used twice to "boost" the probability of a particular conclusion. This is done by keeping track of the support set in a probabilistic inference process. Apart from this, independence assumptions are used where there are no known dependencies, and lower probabilities are manipulated in accord with the laws of probability.

[1]This claim about epistemological adequacy may come as something of a surprise to those who have concentrated their research on "higher-level" knowledge structures. Whatever happened to scripts, plans, TAUs, TOPs, MOPs, etc.? Are these higher-level knowledge structures not essential to story comprehension? There is no doubt that they are. However, there does not seem to be any sharp divisions between any of them. The more focused the successive stages of a script are on an ultimate goal, the more it resembles a plan. The more abstract its level of description, the more it resembles a TAU or a TOP, and so on. Furthermore, there is no particular obstacle to encoding all of them as axiomatic knowledge in EL. For example, the M-BORROW MOP [Dyer, 1983, 207] can be cast as a set of generic conditionals along the following lines. If some person x wants to have some object y temporarily, which he knows to be in the possession of some person z, he may well ask z to lend him y and this may induce z to do so, fulfilling x's goal. If some person x has some object y on loan from some person z, then x is obligated to return y to z, and z will probably want him to do so; etc. I consider the taxonomy of scripts, plans, MOPs, etc., and their elaborate subcategorization, more of a potential guide to control structure — what knowledge is likely to be useful when — than a guide to representation.

Further development of EL semantics: There are still many uncertainties and gaps in the EL semantic theory — as in every extant situation theory. The formal semantics of various constructs, e.g., nominalization, questions, etc., needs to be augmented with detailed algebraic proposals, and axiomatizations justifiable in terms of the underlying algebras. Perhaps most importantly, the existence of model structures of the type described needs to be formally proved (where these models are nontrivial, conforming with the intuitions that motivated the conditions we have assumed).

Further development of EL inference rules: As discussed in Chapter 5, the current versions of RI and GC are not quite general enough. Ways of extending/modifying these rules are currently being pursued.

Computing ELF-translations: A great deal of grammar development remains to be done (including semantic rules), and further deindexing rules are needed, most importantly for relative clauses and various clausal adverbials. Some results on grammar and LF-translation that are not included in this thesis are reported in [Hwang and Schubert, 1992b]. The grammar is also being expanded to cover texts occurring in the TRAINS project. As for deindexing, further investigation is required to properly analyze generic tenses and the interaction between perfect aspect and various temporal adverbials. Also, refinement of the aspectual class system is in order. Some preliminary rules for deindexing clausal adverbials such as "Until Ψ, Φ," "When Ψ, Φ," etc., have been developed and look promising. However, the interaction between temporal perspective shifts and discourse segments requires further investigation as mentioned in Chapter 9. Also, expanding the deindexing mechanism to store event nominals at the tense tree nodes is needed. Finally, for quantified nominals, an algorithm needs to be developed to get the right spatiotemporal frame for restriction clauses. This would require incorporating various techniques developed in research in discourse including centering/focusing techniques.

Implicit question-answering and referent merging: A major research undertaking will be the detailed development of a computational theory of disambiguation, along the lines discussed in Chapter 10 under the "ampliative inference" heading, and more particularly "implicit question answering." As discussed there, our hypothesis is that choice of a "coherent" interpretation, among the many interpretations possible in principle, is guided by an attempt to interpret new inputs in such a way that they answer questions "raised" by the prior discourse; and in this process, new ambiguous referents are merged (unified) with prior ones as a byproduct.

Narrative and simulative inference: A full set of (probabilistic) axioms needs to be developed for the `orients` relation so that "context-discharged" interpretations of this relation in narratives will be correctly computed. Likewise, a full set of axioms for simulative inference, similar to the one given in Chapter 10, need to be developed.

Further, the simulative mechanism which allows a predicate like [x would-infer-from y z] to be evaluated by direct introspection needs to be developed.

Inference control: In the implementation, *inference control* is a serious problem during understanding and question-answering. This problem will be magnified if we attempt to incorporate implicit question-answering and simulative inference into EPILOG. Right now EPILOG stops a particular chain of inferencing when the probability or the "interestingness" of the inferred formulas becomes too low. However, the "interestingness" computation remains unsatisfactory. This is a very subtle problem, since how interesting objects, or inferences about an object, are is strongly interrelated and strongly context-dependent. For instance, that some insignificant object (a wine glass, a shoe, etc.) is in some particular place (e.g., a counter, a kitchen, etc.) may be of insufficient interest to deserve any thought in most contexts, but in a murder mystery it may strike us as being of great interest and significance (e.g., as a clue to the crime).

I have listed some of the imminent issues on our agenda toward the goal of building a narrative understanding system. Also, as we enlarge our purview to include discourse, we would need to concern ourselves with expanding our EL context representation to deal with discourse structures. And, obviously, a huge body of domain knowledge needs to be compiled to deal with significant portions of actual stories.

11.3 An Epilog

I believe there is cause for optimism about the possibility of constructing theoretical and computational frameworks for full, general NLU. Our efforts in developing EL in that direction have led to a rather well-integrated conception of syntax, LF, knowledge representation, context, and inference, and of the interfaces linking these into an organic whole. What is important about the research on EL is its expressive logical syntax, its ease of derivation from syntactic analyses, and in the way the logical forms lend themselves for inference. The conception is not yet complete or fully "debugged," but it is sufficiently far along to have provided a basis for diverse start-up implementations [Schaeffer *et al.*, 1991; Hwang, 1992]. Unlike most past implemented NLU and inference systems, these implementations strenuously avoid cutting corners in syntax, LF-computation, and most of all, knowledge representation and inference. Thus, there is reason to regard the theoretical framework and the implementations as a solid and extensible basis for further work toward the ultimate goal of general NLU.

Bibliography

[Aczel, 1990] Aczel, P. 1990. Replacement systems and the axiomatization of situation theory. *Situation Theory and its Applications*, V. 1, edited by R. Cooper, K. Mukai and J. Perry, CSLI, Stanford, CA.

[Allen and Schubert, 1991] Allen, J. and Schubert, L. K. 1991. The TRAINS project. TR 382, Dept. of Comp. Sci., U. of Rochester, Rochester, NY.

[Allen, 1984] Allen, J. F. 1984. Towards a general theory of action and time. *Artificial Intelligence* 23(2):123–154.

[Allen, 1987] Allen, J. 1987. *Natural Language Understanding.* Benjamin/Cummings Publ. Co., Reading, MA.

[Alshawi and van Eijck, 1989] Alshawi, H. and van Eijck, J. 1989. Logical forms in the Core Language Engine. In *Proc. 27th Annual Meeting of the ACL*, Vancouver, Canada. 25–32.

[Alshawi, 1990] Alshawi, H. 1990. Resolving quasi logical forms. *Computational Linguistics* 16(3):133–144.

[Andrews, 1981] Andrews, P. B. 1981. Theorem proving via general matings. *J. of the ACM* 28(2):193–214.

[Ayuso, 1989] Ayuso, D. M. 1989. Discourse entities in JANUS. In *Proc. 27th Annual Meeting of the ACL*, Vancouver, Canada. 243–250.

[Bacchus, 1988] Bacchus, F. 1988. *Representing and Reasoning with Probabilistic Knowledge.* Ph.D. Dissertation, U. of Alberta, Edmonton, Canada.

[Bacchus, 1990] Bacchus, F. 1990. *Representing and Reasoning with Probabilistic Knowlege: A Logical Approach to Probabilities.* MIT Press, Cambridge, MA.

[Ballard and Stumberger, 1986] Ballard, B. W. and Stumberger, D. E. 1986. Semantic acquisition in TELI: A transportable, user-customized natural language processor. In *Proc. 24th Annual Meeting of the ACL*, New York, NY. 20–29.

[Barwise and Cooper, 1981] Barwise, J. and Cooper, R. 1981. Generalized quantifiers and natural language. *Linguistics and Philosophy* 4(2):159–219.

[Barwise and Perry, 1983] Barwise, J. and Perry, J. 1983. *Situations and Attitudes.* MIT Press (Bradford Book), Cambridge, MA.

[Barwise, 1989] Barwise, J. 1989. *The Situation in Logic.* CSLI, Stanford, CA.

[Bennett and Partee, 1978] Bennett, M. and Partee, B. 1978. *Toward the Logic of Tense and Aspect in English.* Inidiana U. Linguistics Club.

[Bibel, 1979] Bibel, W. 1979. Tautology testing with a generalized matrix reduction method. *Theor. Comput. Sci.* 8:31–44.

[Brachman *et al.*, 1983] Brachman, R. J.; Fikes, R. E.; and Levesque, H. J. 1983. Krypton: a functional approach to knowledge representation. *IEEE Computer* 16:67–73.

[Carlson, 1982] Carlson, G. N. 1982. Generic terms and generic sentences. *J. of Philosophical Logic* 11:145–181.

[Charniak and Goldman, 1988] Charniak, E. and Goldman, R. 1988. A logic for semantic interpretation. In *Proc. 26th Annual Meeting of the ACL*, Buffalo, NY. 87–94.

[Charniak and Goldman, 1989a] Charniak, E. and Goldman, R. 1989a. Plan recognition in stories and life. In *Proc. 5th Workshop on Uncertainty in Artificial Intelligence*, Windsor, Canada. 54–59.

[Charniak and Goldman, 1989b] Charniak, E. and Goldman, R. 1989b. A semantics for probabilistic quantifier-free first-order languages, with particular application to story understanding. In *Proc. 11th IJCAI*, Detroit, MI. 1074–79.

[Charniak and Shimony, 1990] Charniak, E. and Shimony, S. E. 1990. Probabilistic semantics for cost based abduction. In *Proc. AAAI-90.* Boston, MA. 106–111.

[Charniak, 1988] Charniak, E. 1988. Motivation analysis, abductive unification, and non-monotonic equality. *Artificial Intelligence* 34:275–295.

[Chierchia and Turner, 1988] Chierchia, G. and Turner, R. 1988. Semantics and property theory. *Linguistics and Philosophy* 261–302.

[Cooper, 1985] Cooper, R. 1985. Aspectual classes in situation semantics. Technical Report CSLI-84-14C, CSLI, Stanford, CA.

[Cooper, 1987] Cooper, R. 1987. Meaning representation in Montague grammar and situation semantics. *Computational Intelligence* 3:35–44.

[Cullingford, 1981] Cullingford, R. 1981. Sam. In Schank, R. C. and Riesbeck, C. K., editors 1981, *Inside Computer Understanding.* Lawrence Erlbaum. pages 75-119.

[Dahl *et al.*, 1987] Dahl, D. A.; Palmer, M. S.; and Passonneau, R. J. 1987. Nominalizations in PUNDIT. In *Proc. 25th Annual Meeting of the ACL*, Stanford, CA. 131–139.

[Dahlgren *et al.*, 1989] Dahlgren, K.; McDowell, J.; and E. P. Stabler, Jr. 1989. Knowledge representation for commonsense reasoning with text. *Computational Linguistics* 15(3):149–170.

[Dalrymple, 1988] Dalrymple, M. 1988. The interpretation of tense and aspect in English. In *Proc. 26th Annual Meeting of the ACL*, Buffalo, NY. 68–74.

[Davidson, 1967] Davidson, D. 1967. The logical form of action sentences. In Rescher, N., editor 1967, *The Logic of Decision and Action.* U. of Pittsburgh Press. Reprinted in D. Davidson and G. Harman, editors, *The Logic of Grammar*, pages 235-245, Dickenson Publ., Encino, CA, 1975.

[Davidson, 1970] Davidson, D. 1970. Events as particulars. *Noûs* 4:25–32.

[de Haan and Schubert, 1986] de Haan, J. and Schubert, L. K. 1986. Inference in a topically organized semantic net. In *Proc. AAAI-86.* Philadelphia, PA. 334–338.

[Devlin, 1991] Devlin, K. 1991. *Logic and Information.* Cambridge U. Press, Cambridge, UK.

[Dowty, 1979] Dowty, D. 1979. *Word Meaning and Montague Grammar.* Reidel, Dordrecht, Holland.

[Dowty, 1982] Dowty, D. 1982. Tense, time adverbs and compositional semantic theory. *Linguistics and Philosophy* 5:23–55.

[Dowty, 1986] Dowty, D. 1986. The effect of aspectual classes on the temporal structure of discourse: semantics or pragmatics? *Linguistics and Philosophy* 9(1):37–61.

[Dyer, 1983] Dyer, M. G. 1983. *In-Depth Understanding.* MIT Press, Cambridge, MA.

[Enç, 1981] Enç, M. 1981. *An Analysis of Nouns as Indexicals.* Ph.D. Dissertation, U. of Wisconsin, Madison. Also available from Indiana U. Linguistics Club.

[Fenstad *et al.*, 1987] Fenstad, J. E.; Halvorsen, P.-K.; Langholm, T.; and van Benthem, J. 1987. *Situations, Language and Logic.* D. Reidel Publ., Boston, MA.

[Fodor and Sag, 1982] Fodor, J. D. and Sag, I. A. 1982. Referential and quantificational indefinites. *Linguistics and Philosophy* 5:355–398.

[Gawron *et al.*, 1982] Gawron, J. M.; King, J.; Lamping, J.; Loebner, E.; Paulson, E. A.; Pullum, G. K.; Sag, I. A.; and Wasow, T. 1982. The GPSG linguistic system. In *Proc. 20th Annual Meeting of the ACL*, Toronto, Canada. 74–81.

[Gazdar *et al.*, 1985] Gazdar, G.; Klein, E.; Pullum, G.; and Sag, I. 1985. *Generalized Phrase Structure Grammar.* Harvard U. Press, Cambridge, MA.

[Ginzburg, 1991] Ginzburg, J. 1991. Questions without answers, Wh-phrases without scope: A semantics for direct Wh-questions and their responses, Volume 2 of *Situation Theory and its Applications.* CSLI, Stanford, CA. 363–404.

[Grätzer, 1971] Grätzer, G. 1971. *Lattice Theory: First Concepts and Distributive Lattices.* W. H. Freeman, San Francisco, CA.

[Grosz and Sidner, 1986] Grosz, B. J. and Sidner, C. L. 1986. Attention, intentions, and the structure of discourse. *Computational Linguistics* 12:175–204.

[Harman, 1972] Harman, G. 1972. Logical form. *Foundations of Language* 9:38–65. Reprinted in D. Davidson and G. Harman, editors, *The Logic of Grammar*, pages 289-307, Dickenson Publ., Encino, CA, 1975.

[Harper and Charniak, 1986] Harper, M. P. and Charniak, E. 1986. Time and tense in English. In *Proc. 24th Annual Meeting of the ACL*, New York, NY. 3–9.

[Heim, 1982] Heim, I. 1982. *The Semantics of Definite and Indefinite Noun Phrases.* Ph.D. Dissertation, U. of Mass.

[Hemingway, 1952] Hemingway, E. 1952. *The Old Man and the Sea.* Canadian School Edition (1964), Charles Scriber's Sons, Toronto, Canada.

[Hinrichs, 1985] Hinrichs, E. W. 1985. *A Compositional Semantics for Aktionsarten and NP Reference in English.* Ph.D. Dissertation, Ohio State U.

[Hinrichs, 1986] Hinrichs, E. 1986. Temporal anaphora in discourses of English. *Linguistics and Philosophy* 9(1):63–82.

[Hinrichs, 1987] Hinrichs, E. W. 1987. The syntax and semantics of the JANUS semantic interpretation language. In R. M. Weischedel *et al.*, *Research and Development in Natural Language Understanding as Part of the Strategic Computing Program, Annual Report, 21 December 1985–22 December 1986.*

[Hinrichs, 1988] Hinrichs, E. W. 1988. Tense, quantifiers, and contexts. *Computational Linguistics* 14(2):3–14.

[Hirst, 1988] Hirst, G. 1988. Semantic interpretations and ambiguity. *Artificial Intelligence* 34:131–177.

[Hirst, 1991] Hirst, G. 1991. Existence assumptions in knowledge representation. *Artificial Intelligence* 49:199–242.

[Hobbs *et al.*, 1986] Hobbs, J. R.; Croft, W.; Davies, T.; Edwards, D.; and Laws, K. 1986. Commonsense metaphysics and lexical semantics. In *Proc. 24th Annual Meeting of the ACL*. New York, NY. 231–240.

[Hobbs *et al.*, 1987] Hobbs, J. R.; Croft, W.; Davies, T.; Edwards, D.; and Laws, K. 1987. Commonsense metaphysics and lexical semantics. *Computational Linguistics* 13(3-4):241–250.

[Hobbs *et al.*, 1988] Hobbs, J. R.; Stickel, M.; Martin, P.; and Edwards, D. 1988. Interpretation as abduction. In *Proc. 26th Annual Meeting of the ACL*, Buffalo, NY. 95–103.

[Hobbs *et al.*, 1990] Hobbs, J. R.; Stickel, M.; Appelt, D.; and Martin, P. 1990. Interpretation as abduction, Technical Note 499. SRI International, Menlo Park, CA.

[Hobbs, 1979] Hobbs, J. R. 1979. Coherence and coreference. *Cognitive Science* 3:67–90.

[Hobbs, 1985a] Hobbs, J. R. 1985a. On the coherence and structure of discourse. Technical Report CSLI-85-37, CSLI, Stanford, CA.

[Hobbs, 1985b] Hobbs. J. R. 1985b. Ontological promiscuity. In *Proc. 23rd Annual Meeting of the ACL*. Chicago, IL. 61–69.

[Hornstein, 1977] Hornstein, N. 1977. Towards a theory of tense. *Linguistic Inquiry* 3:521–557.

[Hovy, 1990] Hovy, E. H. 1990. Pragmatics and natural language generation. *Artificial Intelligence* 43:153–197.

[Hurum and Schubert, 1986] Hurum, S. and Schubert, L. K. 1986. Two types of quantifier scoping. In *Proc. 6th Canadian Conf. on Artificial Intelligence (AI-86)*. Montreal, Canada. 39–43.

[Hurum, 1987] Hurum, S. 1987. *Quantifier scoping in initial logical translations of English sentences.* Master's thesis, U. of Alberta, Edmonton, Canada.

[Hwang and Schubert, 1992a] Hwang, C. H. and Schubert, L. K. 1992a. Episodic Logic: A situational logic for natural language processing. Presented at *3rd Conf. on Situation Theory and its Applications*, November 18-21, Oiso, Japan; To appear in *Situation Theory and its Applications*, V. 3, edited by P. Aczel, D. Israel, Y. Katagiri, and S. Peters, CSLI, Stanford, CA.

[Hwang and Schubert, 1992b] Hwang, C. H. and Schubert, L. K. 1992b. From parse trees to preliminary logical forms: Sample NP rules. Technical report, Dept. of Comp. Sci., U. of Rochester, Rochester, NY. In preparation for Boeing Co., Seattle, WA.

[Hwang and Schubert, 1992c] Hwang, C. H. and Schubert, L. K. 1992c. Tense trees as the "fine structure" of discourse. In *Proc. 30th Annual Meeting of the ACL.* Newark, DE. 232–240. (Also in *Working Notes: AAAI Fall Symposium on Discourse Structure in Natural Language and Generation*, pages 33-38, November 15-17, 1991, Asilomar, CA.)

[Hwang, 1992] Hwang, C. H. 1992. The deindexer in TRAINS-90. Technical report, Dept. of Comp. Sci., U. of Rochester, Rochester, NY. To appear.

[Jacobs, 1987] Jacobs, P. S. 1987. Knowledge-intensive natural language generation. *Artificial Intelligence* 33:325–378.

[Kadmon, 1987] Kadmon, N. 1987. *On Unique and Non-unique Reference and Asymmetric Quantification.* Ph.D. Dissertation, U. of Mass.

[Kamp, 1971] Kamp, H. 1971. Formal properties of 'now'. *Theoria* 37:227–273.

[Kamp, 1981] Kamp, H. 1981. A theory of truth and semantic representation. In Groenendijk, J.; Janssen, T.; and Stokhof, M., editors 1981, *Formal Methods in the Study of Language.* Mathematical Centre-tracts, U. Amsterdam, Amsterdam, Netherlands.

[Kolodner, 1981] Kolodner, J. L. 1981. Organization and retrieval in a conceptual memory for events or con54, where are you? In *Proc. 7th IJCAI*, Vancouver, Canada. 227–233.

[Kratzer, 1989] Kratzer, A. 1989. An investigation of the lumps of thought. *Linguistics and Philosophy* 12:607–653.

[Krifka, 1989] Krifka, M. 1989. Nominal references, temporal constitution and quantification in event semantics. In Bartsch, R.; van Benthem, J.; and van Emde Boas, P., editors 1989, *Semantics and Contextual Expressions.* Foris Publ., Dordrecht, Holland. 75–115.

[Krifka, 1991] Krifka, M. 1991. Focus and the interpretation of generic sentences. MS.

[Kuhn, 1989] Kuhn, S.T. 1989. Tense and time. In Gabbay, D. and Guenthner, F., editors 1989, *Handbook of Philosophical Logic, Vol. 4: Topics in the Philosophy of Language.* Reidel Publ., Dordrecht, Holland. 513–552.

[Kyburg, 1983] Kyburg, H. E. 1983. The reference class. *Philosophy of Science* 50:374–97.

[Lascarides and Asher, 1991] Lascarides, A. and Asher, N. 1991. Discourse relations and defeasible knowledge. In *Proc. 29th Annual Meeting of the ACL.* Berkeley, CA. 55–62.

[Lascarides and Oberlander, 1992] Lascarides, A. and Oberlander, J. 1992. Temporal coherence and defeasible knowledge. *Theoretical Linguistics* 8. To appear.

[Lascarides *et al.*, 1992] Lascarides, A.; Asher, N.; and Oberlander, J. 1992. Inferring discourse relations in context. In *Proc. 30th Annual Meeting of the ACL*. Newark, DE. 1–8.

[Leech, 1987] Leech, G. 1987. *Meaning and the English Verb (2nd ed)*. Longman, London, UK.

[Lehnert *et al.*, 1983] Lehnert, W. G.; Dyer, M.; Johnson, P.; Yang, C.; and Harley, S. 1983. BORIS – an experiment in in-depth understanding of narratives. *Artificial Intelligence* 20:15–62.

[Link, 1983] Link, G. 1983. The logical analysis of plurals and mass terms: A lattice-theoretical approach. In Bäuerle, R.; Schwarze, C.; and von Stechow, A., editors 1983, *Meaning, Use, and Interpretation of Language*. Walter de Gruyter, Berlin, Germany. 302–323.

[Link, 1987] Link, G. 1987. Generalized quantifiers and plurals. In Gärdenfors, P., editor 1987, *Generalized Quantifiers: Linguistic and Logical Approaches*. Reidel Publ. Co., Dordrecht, Holland. 151–180.

[McCarthy, 1979] McCarthy, J. 1979. First order theories of individual concepts and propositions. In Hayes, J. E.; Michie, D.; and Mikulich, L. I., editors 1979, *Machine Intelligence*, 9. Halsted Press, New York, NY. 129–147. Also in Expert Systems in the Micro-Electronic Age, D. Michie (ed.), Edinburgh U. Press, 1979.

[McCawley, 1981] McCawley, J. D. 1981. Notes on the English present perfect. *Australian J. of Linguistics* 1:81–90.

[McCoard, 1978] McCoard, R. 1978. *The English Perfect: Tense Choice and Pragmatic Inferences*. North-Holland Pub. Co., Amsterdam, Holland.

[Miller *et al.*, 1987] Miller, S.; de Haan, J.; and Schubert, L. K. 1987. *The User's Guide to ECoNet* (Prepared for Boeing Co. under Purchase Contract W-278258). Edmonton, Canada.

[Moens and Steedman, 1988] Moens, M. and Steedman, M. 1988. Temporal ontology and temporal reference. *Computational Linguistics* 14(2):15–28.

[Montague, 1973] Montague, R. 1973. The proper treatment of quantification in ordinary English. In Hintikka, J.; Maravcsik, J.; and Suppes, P., editors 1973, *Approaches to Natural Language: Proc. of the 1970 Stanford Workshop on Grammar and Semantics*. Reidel, Dordrecht, Holland.

[Nakashima *et al.*, 1988] Nakashima, H.; Suzuki, H.; Halvorsen, P.-K.; and Peters, S. 1988. Towards a computational interpretation of situation theory. In *Proc. Internat. Conf. on Fifth Generation Computer Systems (FGCS-88)*, Tokyo, Japan.

[Namioka *et al.*, 1991] Namioka, A.; Hwang, C. H.; Schaeffer, S.; and Russo, P. 1991. Using the inference tool EPILOG for a message processing application. In *Proc. 4th Internat. Conf. on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems (IEA/AIE-91)*, Kauai, Hawaii. 763–769.

[Namioka *et al.*, 1992] Namioka, A.; Hwang, C. H.; and Schaeffer, S. 1992. Using the inference tool EPILOG for a message processing application. *Internat. J. of Expert Systems*. In print.

[Palmer and Finin, 1990] Palmer, M. and Finin, T. 1990. Workshop on the evaluation of natural language processing systems. *Computational Linguistics* 16(3):175–181.

[Partee, 1973] Partee, B. 1973. Some structural analogies between tenses and pronouns in English. *J. of Philosophy* 70:601–609.

[Partee, 1984] Partee, B. 1984. Nominal and temporal anaphora. *Linguistics and Philosophy* 7:243–286.

[Passonneau, 1987] Passonneau, R. J. 1987. Situations and intervals. In *Proc. 25th Annual Meeting of the ACL*. Stanford, CA. 16–24.

[Passonneau, 1988] Passonneau, R. J. 1988. A computational model of the semantics of tense and aspect. *Computational Linguistics* 14(2):44–60.

[Pearl, 1988] Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufman, San Mateo, CA.

[Pollard and Sag., 1987] Pollard, C. and Sag., I. A. 1987. *Information-based Syntax and Semantics: Vol. 1. Fundamentals*. CSLI, Stanford, CA.

[Prior, 1967] Prior, A. 1967. *Past, Present and Future*. Clarendon Press, Oxford, UK.

[Reichenbach, 1947] Reichenbach, H. 1947. *Elements of Symbolic Logic*. Macmillan, New York, NY.

[Richards and Heny, 1982] Richards, B. and Heny, F. 1982. Tense, aspect, and time adverbials. *Linguistics and Philosophy* 5:59–154.

[Roberts, 1973] Roberts, D. D. 1973. *The existential graphs of Charles S. Peirce*. Mouton, The Hague.

[Root, 1986] Root, R. 1986. *The Semantics of Anaphora in Discourse*. Ph.D. Dissertation, U. of Texas, Austin, TX.

[Rosenschein and Shieber, 1982] Rosenschein, S. J. and Shieber, S. M. 1982. Translating English into logical form. In *Proc. 20th Annual Meeting of the ACL*, Toronto, Canada. 1–8.

[Schaeffer *et al.*, 1991] Schaeffer, S.; Hwang, C. H.; de Haan, J.; and Schubert, L. K. 1991. *The User's Guide to* EPILOG (Prepared for Boeing Co. under Purchase Contract W-278258). Edmonton, Canada.

[Schank and Leake, 1989] Schank, R. C. and Leake, D. B. 1989. Creativity and learning in a case-based explainer. *Artificial Intelligence* 40:353–385.

[Schank *et al.*, 1975] Schank, R. C.; Goldman, N. M.; Rieger III, C. J.; and Riesbeck, C. K. 1975. *Conceptual Information Processing.* North Holland, New York, NY.

[Schubert and Hwang, 1989] Schubert, L. K. and Hwang, C. H. 1989. An Episodic knowledge representation for narrative texts. In *Proc. 1st Inter. Conf. on Principles of Knowledge Representation and Reasoning (KR '89)*, Toronto, Canada. 444–458.

[Schubert and Hwang, 1990a] Schubert, L. K. and Hwang, C. H. 1990a. An Episodic knowledge representation for narrative texts. TR345, Dept. of Comp. Sci., U. of Rochester, Rochester, NY.

[Schubert and Hwang, 1990b] Schubert, L. K. and Hwang, C. H. 1990b. Picking reference events from tense trees: A formal, implementable theory of English tense-aspect semantics. In *Proc. Speech and Natural Language, DARPA Workshop*, Hidden Valley, PA. 34–41.

[Schubert and Pelletier, 1982] Schubert, L. K. and Pelletier, F. J. 1982. From English to logic: context free computation of 'conventional' logical translations. *American J. of Comp. Ling.* 8:26–44. Also in *Readings in Natural Language Processing*, B. Grosz, K. Jones and B. Webber, eds., 293-311, Morgan Kaufman, Los Altos, CA, 1986.

[Schubert and Pelletier, 1989] Schubert, L. K. and Pelletier, F. J. 1989. Generically speaking, or, using discourse representation theory to interpret generics. In Chierchia, G.; Partee, B.; and Turner, R., editors 1989, *Property Theory, Type Theory, and Semantics, V.2: Semantic Issues.* Kluwer Academic Publ., Boston, MA. 193–268.

[Schubert, 1992a] Schubert, L. K. 1992a. Formal foundations of Episodic Logic. Technical report, Dept. of Comp. Sci., U. of Rochester, Rochester, NY. To appear.

[Schubert, 1992b] Schubert, L. K. 1992b. Making SAPIR robust: Some proposals. Report prepared for Boeing Computer Services.

[Schuster, 1988] Schuster, E. 1988. Anaphoric reference to events and actions: A representation and its advantages. In *Proc. COLING-88.* 602–607.

[Schuster, 1992] Schuster, E. 1992. *Pronominal Reference to Events and Actions: Computational Foundations.* Ph.D. Dissertation, U. of Pennsylvania, Philadelphia, PA.

[Smith, 1978] Smith, C. 1978. The syntax and interpretations of temporal expressions in English. *Linguistics and Philosophy* 2:43–99.

[Smith, 1991] Smith, C. 1991. *The Parameter of Aspect.* Kluwer Academic Publ., Dordrecht, Holland.

[Song and Cohen, 1991] Song, F. and Cohen, R. 1991. Tense interpretation in the context of narrative. In *Proc. AAAI-91.* Anaheim, CA. 131–136.

[Song, 1991] Song, F. 1991. *A Processing Model for Temporal Analysis and its Application to Plan Recognition.* Ph.D. Dissertation, U. of Waterloo, Waterloo, Canada.

[Steedman, 1982] Steedman, M. J. 1982. Reference to past time. In *Speech, Place and Action.* John Wiley and Sons, New York, NY. 125–157.

[Stickel, 1985] Stickel, M. E. 1985. Schubert's steamroller problem: formulations and solutions. *J. of Automated Reasoning* 2:89–101.

[ter Meulen, 1986] ter Meulen, A. 1986. Locating events. In Groenendijk, J.; de Jongh, D.; and Stokhof, M., editors 1986, *Foundations of Pragmatics and Lexical Semantics.* Foris Publ., Dordrecht, Holland. 27–40.

[ter Meulen, 1988] ter Meulen, A. 1988. Structuring domains for events. MS, Dept. of Linguisitics, U. of Washington, Seattle, WA.

[ter Meulen, 1991] ter Meulen, A. 1991. Dynamic event-structures: Modelling the content of tense and aspect. Presented at *3rd Conf. on Situation Theory and its Applications,* November 18-21, Oiso, Kanagawa, Japan.

[Traugott, 1986] Traugott, J. 1986. Nested resolution. In *8th International Conference on Automated Deduction (CADE-8): Proceedings.* Springer-Verlag, Oxford, England. 394–402.

[Vendler, 1967] Vendler, Z. 1967. Causal relations. *J. of Philosophy* 64:704–713.

[Webber, 1978] Webber, B. L. 1978. *A Formal Approach to Discourse Anaphora.* Ph.D. Dissertation, Harvard U., Cambridge, MA.

[Webber, 1983] Webber, B. L. 1983. So what can we talk about now? In Brady, M. and Berwick, R., editors 1983, *Computational Models of Discourse.* MIT Press, Cambridge, MA. 331–371. Also in *Readings in Natural Language Processing,* B. Grosz, K. Jones and B. Webber, eds., 395-414, Morgan Kaufman, Los Altos, CA, 1986.

[Webber, 1987a] Webber, B. L. 1987a. The interpretation of tense in discourse. In *Proc. 25th Annual Meeting of the ACL*, Stanford, CA. 147–154.

[Webber, 1987b] Webber, B. L. 1987b. Position paper: Event reference. In *TINLAP-3*, Las Cruces, NM. 158–162.

[Webber, 1988] Webber, B. L. 1988. Tense as discourse anaphor. *Computational Linguistics* 14(2):61–73.

[Webber, 1991] Webber, B. L. 1991. Structure and ostension in the interpretation of discourse deixis. *Language and Cognitive Processes* 6(2):107–135.

[Wilensky *et al.*, 1988] Wilensky, R.; Chin, D. N.; Luria, M.; Martin, J.; Mayfield, J.; and Wu, D. 1988. The Berkeley UNIX Consultant Project. *Computational Linguistics* 14:35–84.