

Episodic Logic Meets Little Red Riding Hood: A Comprehensive, Natural Representation for Language Understanding

LENHART K. SCHUBERT

University of Rochester

Rochester, New York 14627-0226

and

CHUNG HEE HWANG

Microelectronics & Computer Technology Corporation (MCC)

3500 West Balcones Center Drive, Austin, TX 78759-5398

(schubert@cs.rochester.edu; hwang@mcc.com)

Abstract. We describe a comprehensive framework for narrative understanding based on *Episodic Logic (EL)*. This situational logic was developed and implemented as a semantic representation and commonsense knowledge representation that would serve the full range of interpretive and inferential needs of general NLU. The most distinctive feature of **EL** is its natural language-like expressiveness. It allows for generalized quantifiers, lambda abstraction, sentence and predicate modifiers, sentence and predicate reification, intensional predicates (corresponding to wanting, believing, making, etc.), unreliable generalizations, and perhaps most importantly, explicit situational variables (denoting episodes, events, states of affairs, etc.) linked to arbitrary formulas that describe them. These allow episodes to be explicitly related in terms of part-whole, temporal and causal relations. Episodic logical form is easily computed from surface syntax and lends itself to effective inference.

The Centrality of Representation in NLP

Language understanding is an organic phenomenon, and the various stages or facets of the language understanding process—parsing, computing a representation, making inferences, etc.—should not be considered in isolation from each other. For instance, 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, and on 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 formal interpretability of the representation allows us to examine in detail whether it captures meanings as intended, and whether the proposed inference rules are semantically justifiable.

These considerations point to the centrality of the issue of 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 our work 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 excerpt from the story of *Little Red Riding Hood*.¹

¹In our later discussion of test scenarios, the wording is slightly different, as we were rather haphazardly using several children’s books. One source was (Perrault, 1961).

The wolf would have very much liked to eat her, but dared not do so on account of some wood-cutters who were in the forest. He asked her where she was going. The poor child, not knowing that it was dangerous to stop and listen to a wolf, said:

“I am going to see my grandmother.”

This excerpt exemplifies the following interesting syntactic/semantic phenomena: it involves modality “dare” that indicates eating Little Red Riding Hood would have been a substantial risk for the wolf in that particular circumstance; it involves causal relationships — both an explicit one (“on account of”) and an implicit one (“not knowing . . .”); it contains a relative clause (“who were in the forest”); it contains an indirect wh-question (“*where* she was going”); it is tensed as well as involving perfect and progressive aspects; it involves a possible fact (“*that* it was dangerous. . .”) as object of the attitude “know”; it involves a gerundive (“not knowing”) and infinitives (“to stop and listen”) whose interpretation is arguably a reified property; it involves the attitude of the narrator (“the *poor* child”); and it involves a *purpose* clause “to see my grandmother.”

Most NL researchers have shied away from fanciful narratives such as fairy tales in recent years. For all their “childishness”, these pose particularly difficult problems in representation and inference. This is not so much because of anthropomorphic animals, magic, and other departures from realism, but because of their open-ended content and their focus on the activities, goals and attitudes of human (or human-like) characters. To us, this makes fairy tales and other fiction particularly useful as a crucible for testing the adequacy of representational and interpretive techniques for NLU. However, we can attest that even simple task-oriented dialogs pose severe representational and interpretive challenges (Allen and Schubert, 1993; Traum et al., 1996).

To provide some sense of what makes our approach distinctive, we should briefly comment on the more usual approaches to semantic representation and knowledge representation in NLU. Typically, the representations employed are either informal or restricted to variants of first-order logic (FOL). In the informal approach (e.g., (Kolodner, 1981; Schank and Leake, 1989)), representations are proposed that typically include standard logical and AI devices such as predication, boolean connectives, slot-and-filler structures, inheritance hierarchies, etc., but also freely add further constructs to deal with beliefs, actions, goals, plans, etc., providing only informal, intuitive explanations of these constructs. The advantage of the informal approach is that the practitioner can quickly accommodate a rich variety of concepts and ideas in the representation and proceed with the investigation or modelling of some specific phenomena without being detained very much by intricate foundational and mathematical questions. The price, of course, is uncertainty as to 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.

The FOL approach (e.g., (Dahlgren et al., 1989; Hirst, 1988; Wilensky et al., 1988)) restricts itself to predicates and functions whose arguments are individuals (e.g., *loves(John, Mary)*, *mother-of(John)*), boolean connectives ($\wedge, \vee, \neg, \rightarrow, \dots$), equality, and some quantifiers (e.g., \forall, \exists), or syntactic forms that can in principle be reduced to those of FOL. The advantage is that FOL is well-understood syntactically and semantically. But it also has the disadvantage that very little real language is easily expressible in it. For instance, it does not (in any direct way) allow for beliefs, intensional verbs (such as *needing* something), modifiers (such as *very politely* or *possibly*), complex quantifiers (such as “every butterfly along the way”), habituals (such as “she often visited her grandmother”), and many other quite ordinary locutions (see the subsection on Nonstandard Constructs below).

In some approaches to semantic representation, the emphasis is more on mimicking the surface form of certain kinds of NL phrases than on matching the full expressive power of NL. To the extent that the proposed logical forms can readily be paraphrased in FOL, such approaches can still be classified as FOL approaches. For example, McAllester and Givan propose a form of FOL with quantification expressed in the manner of Montague (Montague, 1973) through *generalized quantifiers* (expressions corresponding to noun phrases such as “every man who walks”); but rather than trying to achieve greater expressiveness, their interest is in a subset of their language that allows polynomial-time satisfiability testing (McAllester and Givan, 1992). In a similar vein, Ali and Shapiro’s ANALOG representation (Ali and Shapiro, 1993) renders complex noun phrases (with determiners *all*, *some*, or *any*) as *structured variables* in semantic networks. While some devices are also offered for dealing with so-called branching quantifiers and donkey anaphora (for further remarks on the latter see the next section), the representation is for the most part easily translatable to

FOL.

A few systems do use significantly extended versions of FOL as a representation language. Iwanska’s UNO language (Iwańska, 1993; Iwańska, 1997) allows for some types of modifiers through the use of functions (e.g., a *speed* function to express walking speed in “John walks fast”). It also encodes complex quantified noun phrases as second-order predicates (in essence, Montague-style generalized quantifiers), and provides ways of combining first- and second-order predicates (not just sentences) with *and*, *or* and *not*. Alshawi and Eijk’s Core Language Engine (CLE) (Alshawi and van Eijck, 1989) allows for, among other things, event variables, generalized quantifiers, collective and measure terms, natural kinds, and comparatives and superlatives. TACITUS (Hobbs et al., 1987) allows for event variables, sets, scales, time, spaces and dimension, material, causal connection, etc. But where they go beyond FOL, the latter two systems tend to be unclear about semantics. Also, UNO, CLE and TACITUS still fall short of comprehensive expressiveness; for instance, they lack 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. (UNO is perhaps most nearly NL-like in form, but since it is an attribute-value logic, it needs to introduce numerous supplementary functions such as role functions for n -place predicates, a speed-of-walking function for “walk fast”, a mental-attribute function for “bright student”, etc.).

A recent trend in NLP has been to try to circumvent many of the syntactic and semantic complexities of written and spoken language by aiming to extract only certain predetermined kinds of information from narrowly focused classes of texts or discourse. While such an approach can achieve high *quantitative* returns (large amounts of data extracted from large corpora), it necessarily compromises the *quality* of understanding. We believe that achieving deeper understanding is an important and realistic goal. In fact, some things are made easier by aiming higher. For instance, computing logical forms is easier if the target representation is NL-like in expressiveness, rather than being some restricted frame-like language. System builders constrained to use restricted languages usually find themselves resorting to “illegal” (and semantically unanalyzed) add-ons, in the effort to deal with real language. As well, certain inferences are made easy by an expressive language which would be at best roundabout, if possible at all, in less expressive representations, such as inferences based on generalized quantifiers like “most” or modifiers like “almost”. In fact, we think that the inferences we have obtained experimentally for story fragments are quite beyond the capabilities of virtually all extant knowledge representations.

In the following, we report the results of our effort to develop a comprehensive representation for a general NLU system, and describe our conception of the language understanding process based on that representation. *Episodic Logic* (EL) is a highly expressive knowledge representation well-adapted to the interpretive and inferential needs of general NLU. EL serves simultaneously as the semantic representation and knowledge representation, i.e., it is capable of representing both the explicit content of discourse and the linguistic and world knowledge needed to understand it. EL is designed to be easily derivable from surface syntax, to capture the semantic nuances of NL text, and to facilitate needed inferences.

In the next section, we briefly introduce EL—its syntax, semantics and inference rules; in the subsequent section, we discuss the NLU process within our framework, with emphasis on inference and understanding. Then, in a further section, we illustrate our NLP strategy— from semantic representation to knowledge representation and to the reasoning process. In the penultimate section, we describe the EPILOG implementation and our work on some prototype NLU systems and on story fragments. In the concluding section we summarize the distinctive features of EL and outline future research.

Introduction to EL

EL is an “NL-like” logical representation whose syntax echoes the kinds of constructs that are available in all natural languages. The adjective *episodic* alludes to the use of explicit terms denoting events and other episodes, and to the fact that narrative texts focus on *time-bounded* eventualities (such as someone being hungry or having a meal), rather than on “timeless” ones (such as wolves being animals, or wage earners having to pay taxes). Our overview begins with a simple example, then enumerates the most important nonstandard constructs, and provides a sketchy outline of semantics and inference.

Basic Sentential Syntax

The following example serves to convey the “flavor” of EL.

- (1) a. Little Red Riding Hood chased a butterfly
 b. [LRRH ⟨past chase⟩ ⟨∃ butterfly⟩]
 c. (past (∃x: [x butterfly] [LRRH chase x]))
 d. (∃ e₁: [e₁ before Now₁] [(∃x: [x butterfly] [LRRH chase x]) ** e₁])

(1b) is an unscoped logical form (ULF). In particular, the angle brackets ⟨ ⟩ indicate that the tense operator ‘past’ and restricted quantifier ‘∃ butterfly’ are still to be moved leftward until they have an entire sentential formula in their scope. The result of this scoping, which also introduces a variable for the quantifier, is the LF shown in (1c). Note that the meaning of this LF is still context-dependent, since ‘past’ is an indexical operator, i.e., its meaning depends on *when* the given sentence was uttered. The process of deindexing (removing context dependence) associates an existentially quantified episodic variable e₁ (an episode of LRRH chasing a butterfly) with the sentence in the scope of ‘past’, and relates the episode explicitly to the time of utterance (denoted by Now₁, a new time constant). The result is the deindexed episodic logical form (ELF) shown in (1d). Note that we use square brackets and infix syntax (with the predicate in second place) in sentential forms like [LRRH chase x]. This is a “prettified” variant of the underlying prefix form, ((chase x) LRRH), predicating the property of “chasing x” of LRRH.² The sentential infix syntax greatly aids readability for complex formulas. The general form of restricted quantification is (Qα:ΦΨ), where Q is a quantifier such as ∃, ∀, Most, or Few, α is a variable, and restriction Φ and matrix Ψ are arbitrarily complex formulas. (∀α:ΦΨ) and (∃α:ΦΨ) are equivalent to (∀α) [Φ → Ψ] and (∃α) [Φ ∧ Ψ], respectively. (However, for nonstandard quantifiers such as Many, Most, and Few, there are no such reductions from the restricted form to an unrestricted one.)

The most unconventional feature of the ELF in (1d) is of course the ‘**’ operator associating an episode with a sentence. It can be read as *characterizes*, i.e., in this case LRRH’s chasing a butterfly characterizes episode e₁; or to put it as we did above, e₁ is an episode *of* LRRH chasing a butterfly. The ability to associate episodes with sentences is crucial not only for making tense information explicit as in (1d), but also (as illustrated below) for capturing the content of locative, temporal and other adverbials, for enabling anaphoric reference to events, and for making causal relations among events explicit.

We also have a related episodic operator ‘*’, where [Φ * η] means “Φ is true *in* (or, describes some part or aspect of) η.” Note that [Φ ** η] implies [Φ * η]; for instance, if e is an episode *of* the sun setting, then the sun sets *in* episode e. (The converse does not in general hold: it may be that the sun sets *in* a certain episode, say one where John drives from New York to Chicago, but that does not make the drive an episode *of* the sun setting.) Whereas the operator ‘**’ is introduced by English sentences as above, ‘*’ is typically introduced by *meaning postulates*, i.e., general axioms about the meanings of classes of predicates or particular predicates. For instance, suitable meaning postulates about ‘chase’, when applied to

$$[(\exists x: [x \text{ butterfly}] [LRRH \text{ chase } x]) ** e_1]$$

might lead to the conclusions

$$[(\exists x: [x \text{ butterfly}] [LRRH \text{ see } x]) * e_1], \text{ and} \\ [(\exists x: [x \text{ butterfly}] [x \text{ move}]) * e_1].$$

Note that in any episode of LRRH chasing a butterfly, she surely saw a butterfly, and a butterfly surely moved. Another way to say this would be to introduce *subepisodes* e₂, e₃ of e₁ *characterized* by LRRH seeing a butterfly, and a butterfly moving, respectively. (We could use *anaphoric variables*, briefly discussed in the next subsection, to refer to the *same* butterfly in all three episodes.) Note that ‘**’ and ‘*’ are modal operators as they are not truth-functional, i.e., they do not in general allow substitution for their sentential argument of another sentence with the same truth value. For example, [[LRRH sing] * e₁] does not entail

$$[[[LRRH \text{ sing}] \wedge [[Granny \text{ ill}] \vee \neg [Granny \text{ ill}]]] * e_1].$$

²Observe that in this underlying prefix form we are applying the predicate to one argument at a time. In this so-called “curried” form of predication (associated with the names of Curry, Schoenfeld and Church), an *n*-place predicate is interpreted as a function that can be applied to a single argument to give an (*n* − 1)-place predicate (and finally a truth value, when *n* = 1).

Anaphoric Variables

Consider the following two successive sentences and their logical forms.

- (3) a. Little Red Riding Hood chased every butterfly (that she saw)
 b. (past ($\forall x: [x \text{ butterfly}] [\text{LRRH chase } x]$))
 c. ($\exists e_1: [e_1 \text{ before Now}_1] [(\forall x: [x \text{ butterfly}] [\text{LRRH chase } x]) ** e_1]$)
- (4) a. *It* made her tired
 b. (past [*It* (make2 tired) LRRH])
 c. ($\exists e_2: [e_2 \text{ before Now}_2] [[e_1 \text{ (make2 tired) LRRH}] ** e_2]$))

As before, (3b) and (4b) are preliminary, indexical LFs (omitting the relative clause for simplicity), and (3c) and (4c) are deindexed ELFs. In (4b,c), ‘make2’ is a “2-fold predicate modifier”, mapping the 1-place predicate ‘tired’ into the 2-place predicate (make2 tired). (Note the distinction between a 2-fold predicate modifiers and ordinary ones such as ‘very’, which produce 1-place predicates.) Observe that *It* in (4b) has been resolved to e_1 in (4c), so that e_1 now occurs *outside* the scope of its \exists -quantifier in (3c). Such *anaphoric variables* are allowed in EL, thanks to a “parameter” mechanism that does much the same work as *dynamic binding* in DRT (Kamp, 1981) or dynamic predicate logic (Groenendijk and Stokhof, 1991). Intuitively, we can think of the existential quantification in (3c) as binding some value (viz., an episode of LRRH chasing every butterfly she saw) to the variable e_1 , where this binding “persists” to (4c). The effect is that the conjunction of (3c) and (4c) is interpreted as if the \exists -quantifier binding e_1 had maximally wide scope.

The parameter mechanism is also crucial for dealing with “donkey anaphora” (Geach, 1962) in sentences like “Every man who owns a donkey feeds it”, or “If I have a quarter I’ll put it in the parking meter”. Sentences of this sort provided much of the impetus behind the development of DRT and dynamic predicate logic. The difficulty in the examples lies in the fact that an existentially quantified variable (for “a donkey” in the first sentence and “a quarter” in the second) is referred to by a pronoun (“it”) lying *outside* the scope of the quantifier. The situation is thus much as in the pair of (c)-sentences above, and is handled by the parameter mechanism in much the same way. Semantically, the second sentence (from (Schubert and Pelletier, 1989)) in our treatment is logically equivalent to “If I have a quarter, then I have a quarter that I will put in the parking meter”. This differs from the standard DRT treatment in not asserting that I will put *all* quarters that I have into the meter; for some further discussion of these issues see, e.g., (Chierchia, 1995; Schubert, to appear).

Attitudes

We think that the objects of attitudes such as *believing*, *telling*, *hoping*, etc., are propositions. These are abstract individuals formed from sentence intensions (which in our case are truth-valued partial functions on situations) by applying the operator **That**, as illustrated below.

- (5) a. Mother told Little Red Riding Hood that Grandmother was ill
 b. (past [Mother tell LRRH (That (past [Granny ill]))])
 c. ($\exists e_1: [e_1 \text{ before Now}_1]$
 [[Mother tell LRRH (That
 ($\exists e_2: [e_2 \text{ at-or-before } e_1] [[\text{Granny ill}] ** e_2]$))]
 ** e_1]])

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.

Actions

Actions are distinguished from events or episodes in that they have well-defined *agents*—thus, one may perform an action, but *not* perform an episode or event; likewise, there are intentional actions, wicked actions, etc., but not “intentional events” or “wicked events.” In EL, actions are represented as agent-event

pairs; i.e., to specify a particular action is to specify both the agent of the action and the event brought about through the action. Here is an example.

- (6) a. The wolf gobbled up Grandmother
 b. *It* was a very wicked deed
- (7) a. (past [Wolf gobble-up Granny])
 b. (past [*It* ((attr (very wicked)) deed)])
- (8) a. $(\exists e_1: [e_1 \text{ before Now}_1] [[\text{Wolf gobble-up Granny}] ** e_1])$
 b. $(\exists e_2: [e_2 \text{ before Now}_2] [[[\text{Wolf} \mid e_1] ((\text{attr (very wicked)) deed})] ** e_2])$

Notice that *It* in (7b) is resolved to the ordered pair [*Wolf* | e_1] in (8b), namely, the wolf’s action of gobbling up Grandmother. ‘|’ is a pairing function applicable to individuals and tuples. (As in Lisp and Prolog, an individual paired with an n -tuple gives an $(n + 1)$ -tuple headed by the individual.)

Kinds of Actions and Events

Our approach here is inspired by (Carlson, 1982) and (Chierchia, 1985). We start with a basic kind-forming operator *K* applicable to predicates like ‘dog’ (e.g., (*K dog*) represents dog-kind, whose instances are dogs), and then posit analogous operators for forming *kinds of actions* and *kinds of events*. For example, “to visit Grandmother” is a kind of action, and (9a) says that Little Red Riding Hood likes to do that kind of action. (Compare with “Little Red Riding Hood likes animals”, which is about her liking a particular kind of *thing*, viz., animals.) On the other hand, “for Little Red Riding Hood to talk to a stranger” is a kind of event, and (10a) asserts that this kind of event is not unusual. (Compare with “Gray wolves are not unusual”, which makes a generic claim about the kind of thing, gray wolves. To be more accurate, we should perhaps use *gpres* (generic present) tense below.)

- (9) a. Little Red Riding Hood likes to visit Grandmother
 b. (pres [LRRH like (*Ka* (visit Granny))])
 c. $(\exists e_1: [e_1 \text{ at-about Now}_1] [[\text{LRRH like } (\text{Ka (visit Granny)})] ** e_1])$
- (10) a. For Little Red Riding Hood to talk to a stranger is not unusual
 b. (pres $(\neg [(\text{Ke } (\exists x: [x \text{ stranger}] [\text{LRRH talk-to } x])) \text{ unusual}])$)
 c. $(\exists e_2: [e_2 \text{ at-about Now}_2] [(\neg [(\text{Ke } (\exists x: [x \text{ stranger}] [\text{LRRH talk-to } x])) \text{ unusual}]) ** e_2])$

In these representations, *Ka* maps 1-place action predicates into kinds of actions, and *Ke* maps sentences into kinds of events. *Ka*- or *Ke*-constructs can be equivalently written as constructs headed by the *K* operator.³

Probabilistic Conditionals

We use probabilistic conditionals of form $\Phi \rightarrow_{p, \alpha_1, \alpha_2, \dots, \alpha_n} \Psi$, where $\alpha_1, \alpha_2, \dots, \alpha_n$ are *controlled variables*, to represent extensionally interpretable generic statements. Intuitively, the meaning is that *at least* a fraction p of the tuples of values of $\alpha_1, \dots, \alpha_n$ that satisfy Φ also satisfy Ψ . Let us consider the following example. (Here *non* is a predicate modifier with the property that $[x (\text{non } \pi)]$, for π a monadic predicate, entails $\neg [x \pi]$.)

- (11) a. If a predatory animal finds a nonpredatory creature of modest size,
 he may attack it.
 b. $(\exists e_1 (\exists x: [x ((\text{attr predatory}) \text{ animal})]$
 $(\exists y: [[y ((\text{attr non predatory}) \text{ creature})] \wedge$
 $(\neg [y \text{ big-rel-to } x]) \wedge (\neg [y \text{ tiny-rel-to } x])]$
 $[x \text{ find } y])]) ** e_1)$

³In particular, (*Ke* Φ) is equivalent to (*K* $\lambda e[\Phi ** e]$), i.e., the kind of event that is characterized by Φ ; and (*Ka* π) is equivalent to (*K* $\lambda a[[\text{fst } a] \pi] ** [\text{rst } a]$), i.e., the kind of action such that the event this action brings about ($(\text{rst } a)$, the second element of a) is characterized by the sentence $[(\text{fst } a) \pi]$, where $(\text{fst } a)$ is the agent of the action (the first element of a).

$$\rightarrow_{.2, e_1} (\exists e_2: [e_1 \text{ immed-cause-of } e_2] [[x \text{ attack } y] ** e_2])$$

This formula says that in at least 20% of the situations e_1 in which the antecedent is true, the consequent will also be true. This statistical probability becomes the epistemic probability of the consequent, when we detach the consequent for some true instance of the antecedent. For instance, given that “A wolf found a rabbit” (and background knowledge to the effect that a wolf is a predatory animal, that a rabbit is a nonpredatory animal that is neither big nor tiny relative to a wolf), the above conditional allows us to conclude that the wolf may have attacked the rabbit, with minimum epistemic probability (degree of confidence) .2. The fact that only e_1 (and not x and y) are controlled in the conditional means that in a situation where multiple predators (e.g., a pack of wolves) encounter multiple potential prey (e.g., a herd of deer), we do not predict an attack by each predator on each prey, just *some* predator-prey attack.⁴ Probabilistic conditionals are very convenient in representing generic world knowledge, and as will be seen in a later section, are used extensively in our implementation.⁵

Adverbials

We focus on verb phrase adverbials here such as temporal, locative and manner adverbials, since these are the most common. We interpret such adverbials in a two-stage process. First, in forming the preliminary (indexical) LF, we map adverbials that intuitively modify actions (including manner adverbials such as “politely”) into predicate operators applied to the interpretation of the verb phrase; and we map adverbials that intuitively modify episodes (including temporal and locative adverbials such as “in the forest”) into sentential operators applied to the interpretation of the sentence. In the second, deindexing stage, we recast the action-modifying operators as explicit predications about actions, and episode-modifying operators as explicit predications about episodes. This is made possible by the introduction of explicit episodic variables in the deindexing process that produces the final ELF. For example, consider the following sentence, involving two adverbials.

- (12) a. The wolf politely greeted Little Red Riding Hood in the forest
 b. (past (The x : [x wolf] (The y : [y forest]
 ((adv-e (in-loc y)) [x ((adv-a (in-manner polite))
 (greet LRRH))))))
 c. ($\exists e_1$: [e_1 before Now₁]
 [[[e_1 in-loc Forest₁] \wedge [[Wolf | e_1] (in-manner polite)] \wedge
 [Wolf greet LRRH]]
 ** e_1])

In (12a), “in the forest” modifies the *episode* described by “the wolf greet Little Red Riding Hood,” or, more specifically, its spatial location; “politely,” on the other hand, modifies the *action* of the wolf’s greeting Little Red Riding Hood (by specifying the manner in which the action was performed). In the indexical LF, an episode-modifying adverbial assumes the form (*adv-e* π), where π is a predicate over episodes, and an action-modifying adverbial takes the form (*adv-a* π), where π is a 1-place predicate over actions (more generally, over “attributes”, allowing for static verb phrases). That is, *adv-e* is an operator that uniformly maps 1-place predicates into sentence modifiers, and *adv-a* is an operator that maps 1-place predicates into predicate modifiers. Additional examples are (*adv-e* (during Yesterday)) for “yesterday,” (*adv-e* (lasts-for (K ((num 1) hour)))) for “for an hour,” and (*adv-a* (with-accomp LRRH)) for “with Little Red Riding Hood.”

Note that the scope of **** in (12c) extends leftward over a conjunction of three formulas, so that e_1 is asserted to be an episode of the wolf greeting LRRH politely in the forest. Certain general axioms allow us to narrow the scope of **** to exclude atemporal formulas like [e_1 in-loc Forest₁] and [[Wolf | e_1]

⁴Controlled variables thus allow us to address the “proportion problem” (Kadmon, 1987). If we made x an additional controlled variable in the example, we would be quantifying over individual predators, even in situations where several predators simultaneously find some nonpredatory creature(s).

⁵The topic of generic sentences is a complex one (e.g., (Carlson and Pelletier, 1995)), and “genericity” cannot in general be equated with statistical preponderance. For instance, Carlson’s sentence “Alligators die before they are 20 weeks old” seems false even if survival chances for baby alligators are very poor. Nonetheless, statistical generalizations are very useful, and cover more ground than might be thought (see (Cohen, 1997) for some interesting points).

(in-manner polite)]; in the present case, this brings out the fact that an episode of the wolf greeting LRRH politely in the forest is necessarily an episode of the wolf greeting LRRH. With this scope-narrowing, and after Skolemizing e_1 to E_1 and separating conjuncts, we obtain

- a. $[E_1 \text{ before Now}_1]$
- b. $[E_1 \text{ in-loc Forest}_1]$
- c. $[[\text{Wolf} \mid E_1] \text{ (in-manner polite)}]$
- d. $[[\text{Wolf greet LRRH}] ** E_1]$

This makes plain how adverbials in EL ultimately provide conjunctive information about the described episode.⁶ It is also worth noting that this ultimate format is quite similar to a Davidsonian one (Davidson, 1967). However, while Davidson introduced event variables as “extra arguments” of verbs, our approach (following (Reichenbach, 1947) and the more recent situation-theoretic tradition) associates episodic variables with arbitrarily complex sentences. This has the important advantage that it allows us to make formal sense of such notions as “a three-day episode of the wolf not eating anything” (involving negation), “an episode of John drinking and driving” (involving conjunction), and “the lengthy process of each graduate ascending to the podium and taking receipt of his or her diploma” (involving quantification). In other words, we contend that not only atomic predications, but arbitrary sentences, can be used to characterize episodes; several of our examples (e.g., (10c), (12c)) have illustrated this point.

Some linguistic phenomena whose representation in EL is still somewhat up in the air are quantifier modifiers (as in “very few”, or “all but five”), comparatives (such as “ran as fast as his feet would carry him”, or “the better to hear you with”), and Wh-questions (“Why are your ears so large?”). Jumping ahead a little, we should mention that our EPILOG system (the computational system for EL; (Miller et al., 1991)) is able to answer many yes-no and Wh-questions, for instance “Who met whom?”, expressed as

$$(W? x: [x \text{ person}] (W? y: [y \text{ person}] (\exists e [[x \text{ meet } y] * e]])).$$

However, the analysis of questions we currently favor calls for some modification of this format, involving the use of intension and extension operators.⁷ Despite these open issues, we believe that EL is the most expressive knowledge and semantic representation yet to be brought to bear on the problem of NL understanding.

Ontology and Glimpses of Semantics

Our syntactic examples involved not just ordinary individuals such as people and butterflies, but also events (episodes), collections, actions, propositions, kinds of things, and kinds of events and actions. Correspondingly, the semantics of EL is based on an ontology of *possible individuals* \mathcal{D} that includes all of these sorts of things (and some others), as shown in Fig. 1. As (Hobbs, 1985) argues, it is better to expand one’s ontology to allow more kinds of entities than to complicate the logical form of sentences, or the interpretive process. *Possible* individuals are meant to include not only real or actual individuals but also imaginary or fictitious ones, such as those denoted by the phrases “Sherlock Holmes, the fictitious detective” and “the cancelled lecture” (from (Hirst, 1991)).

The most distinctive aspect of our semantics concerns the ontology of *possible situations* \mathcal{S} (the lower left portion of Fig. 1), their special subclasses, and their part-of structure. We use the term “situation” rather than “episode” when discussing denotational semantics, in deference to custom in situation semantics and also to avoid the implication that we are limiting ourselves to time-bounded situations. We discuss situations and their subclasses in greater detail below.

Disjointly from \mathcal{S} , we have not only ordinary individuals of our experience, but also propositions \mathcal{P} , possible facts \mathcal{F} (which as mentioned are consistent propositions), kinds of individuals \mathcal{K} (including kinds

⁶See (Hwang and Schubert, 1994) for an extensive discussion of our treatment of temporal and other kinds of adverbials. However, since writing that paper we have made some significant adjustments in our conception of the connection between episodes and sentences, leading to a simpler ELF for sentences with adverbials.

⁷In the above example, the embedded question would be prefixed by an extension operator “ \forall ”; and a question like “Which book did each child read?” would involve both an intension and an extension operator:

$$\wedge (\forall x: [x \text{ child}] \forall (W? y: [y \text{ book}] (\exists e [[x \text{ read } y] * e]])).$$

We cannot discuss question semantics here, except to mention that we view questions as functions on episodes, where the value of a question at an episode, if defined, is a full and true answer to the question (and thus is a sentence intension). Our reliance on *full* answers in the semantics is something of a departure from more standard approaches (e.g., (Chierchia, 1993)).

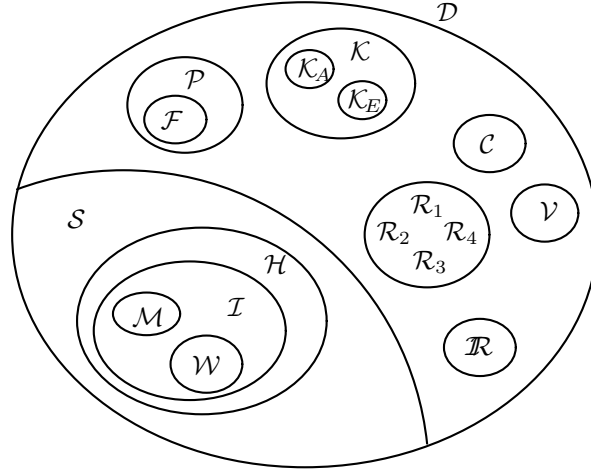


Fig. 1. Ontology of Basic Individuals

of ordinary individuals, kinds of actions \mathcal{K}_A , and kinds of episodes, or situations, \mathcal{K}_E), the real numbers \mathcal{R} (augmented with $-\infty$ and $+\infty$), and n -dimensional regions \mathcal{R}_n ($1 \leq n \leq 4$), containing subsets of \mathcal{R}^n . \mathcal{R}_4 contains space-time trajectories that may not be connected. These are important since we regard situations as occupying times and places, or, more generally, spatiotemporal trajectories (regions). Finally, there are collections \mathcal{C} and n -vectors (i.e., tuples) \mathcal{V} , $n = 2, 3, \dots$, of all of these.

Situations, Times, and Worlds

Possible situations subsume what we might ordinarily (informally) call specific events, states of affairs, and circumstances or eventualities. Unlike situations in situation semantics (which are deemed to be *real*), possible situations in EL are “partial possible worlds,” in that predicate symbols are assigned partial extensions (argument values for which they are true) and antiextensions (argument values where they are false) relative to them. Indeed, we get from arbitrary situations to possible worlds by maximizing over space, time, and information: among the possible situations \mathcal{S} are the informationally maximal *exhaustive situations* \mathcal{H} , and among the exhaustive situations are the spatially maximal *possible times* \mathcal{I} (intervals), conceived of as “everything that happened or was the case over a particular clock-time interval”; possible times in turn include the spatiotemporally maximal *possible worlds* \mathcal{W} and the spatially maximal, temporally minimal *moments of time* \mathcal{M} . Thus the usual indices of semantic evaluation, worlds and times, are here collapsed into one, viz., situations.

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.⁸ Note that *actions* or *activities* are not included in \mathcal{S} . Actions are regarded as events paired with their agents, as illustrated in the earlier subsection on Actions.

Part-of Structure of Situations and Persistence of Information

The notion of “maximizing” along spatiotemporal and information dimensions presupposes a *part-of* ordering among situations. The particular part-of structure we assume is motivated by certain intuitively warranted entailment patterns (or “truth persistence” patterns). These can be observed when we evaluate sentences relative to ever-more-inclusive situations, or conversely, relative to ever-smaller subsituations. We briefly illustrate some of the phenomena at issue, using mnemonic abbreviations for certain sentences and for episodes characterized by those sentences:

[WalkToCottage * EntireWalk], [WalkThroughForest * ForestWalk].

⁸ *Clock times* are distinguished from *times* in the episodic sense: clock times are formally modelled as multi-intervals on the real numbers, and as such have no “information content.”)

Here “WalkToCottage” stands for the (formalized, tenseless) *sentence* “LRRH walks to Grandmother’s cottage”, and “EntireWalk” stands for an *episode* characterized by that sentence, i.e., it is an episode of LRRH walking to Grandmother’s cottage. “WalkThroughForest” stands for “LRRH walks through the forest” and correspondingly “ForestWalk” is an episode of LRRH walking through the forest, specifically the *part* of “EntireWalk” that is located in the forest on LRRH’s way to Grandmother’s cottage. We will also use “MeetWolf” to abbreviate “LRRH meets the wolf”, and “Alone” to abbreviate “LRRH is alone”.

Given this, we note the following entailments and nonentailments. (Here ‘ \models ’ and ‘ \models ’ mean “entails” and “is entailed by”, and the crossed-off versions deny the corresponding entailments.)

- a. $[\text{MeetWolf} * \text{ForestWalk}] \models_{\neq} [\text{MeetWolf} * \text{EntireWalk}]$
- b. $[\text{Alone} * \text{ForestWalk}] \not\models_{\neq} [\text{Alone} * \text{EntireWalk}]$
- c. $[(\neg \text{MeetWolf}) * \text{ForestWalk}] \not\models_{\neq} [(\neg \text{MeetWolf}) * \text{EntireWalk}]$
- d. $[(\neg \text{Alone}) * \text{ForestWalk}] \not\models_{\neq} [(\neg \text{Alone}) * \text{EntireWalk}]$

Here (a) illustrates what we call *outward persistence* of *telic* formulas; i.e., if a telic formula Φ – one that describes an inherently time-bounded, culminating episode such as meeting or greeting someone, walking to Grandmother’s cottage, etc. – is true in a temporal segment of a larger episode, then it is also true in the larger episode. The converse does not hold, i.e., knowing only that LRRH meets the wolf in her walk to Grandmother’s cottage, we cannot say that she meets him in her walk through the forest – the meeting might take place in another part of the walk. Thus we do not have *inward persistence* for telic formulas. We have the opposite situation for the *atelic* formula in (b) (asserting that LRRH is alone) – an inherently “homogeneous” description, not entailing a particular culmination or termination within the described episode. Clearly if LRRH is alone in her walk to Grandmother’s cottage, she is alone in any temporal segment of that walk, in particular in her walk through the forest, while the converse need not hold. So for atelic formulas, we have inward, but not in general outward persistence. (As just indicated, inward persistence is sometimes called *homogeneity* in the literature.) In (c) and (d), we see that negated formulas behave like atelic ones, whether the original formula was telic or atelic. In both cases, we have inward, but not outward persistence.⁹

Situations can be part of one another in both a temporal sense and in an informational sense. For example, LRRH’s walk through the forest is a temporal segment of her walk to Grandmother’s cottage. On the other hand, if LRRH was alone and carefree in her walk to Grandmother’s cottage, then the situation of her being alone, and that of her being carefree, are *coextensive* (simultaneous) parts of the “more informed”, cumulative situation of her being alone and carefree in that walk.

The coextensive subepisode ordering is written $s \preceq s'$ and relates a “less informed” situation s to a “more informed” coextensive situation s' , i.e., one with the same temporal location but with more situational content. This basic ordering supports full persistence of information: whatever is true (false) in s is also true (false) in s' . We call this form of persistence *upward persistence* (imagining more informed situations as being “higher” in the \preceq -ordering). There is also a *subsegment* relation \trianglelefteq , where $s \trianglelefteq s'$ means that s is a temporal segment of s' (or a multi-segment, consisting of multiple disjoint segments). Only telic and atemporal (eternal) sentences are guaranteed to have persistent extensions through the \trianglelefteq ordering. For instance, if [LRRH meet Wolf] or [5 integer] is true in s , then it is also true in s' , for $s \trianglelefteq s'$. This is what we called “outward persistence” above. But for an atelic sentence like [LRRH alone], its truth in s does not guarantee its truth in s' , for $s \trianglelefteq s'$.

We can combine \preceq and \trianglelefteq by forming the transitive closure of their disjunction, i.e., the transitive closure of

$$\{\langle s, s' \rangle \mid s \preceq s' \text{ or } s \trianglelefteq s'\}.$$

We write the resulting relation as \sqsubseteq , and refer to this as the (general) *subepisode* relation. In this partial ordering, a subepisode can be both informationally and temporally “smaller” than the situations of which it is a part. Note that telic and atemporal sentences are outward/upward persistent in the \sqsubseteq -ordering. Atelic sentences are inward persistent in the \trianglelefteq -ordering, but not in general inward/downward persistent in the

⁹In expanding out (d) to “LRRH is not alone in her walk through the forest”, we have to be careful not to misunderstand the negation as having wide scope; i.e., the intended reading is that LRRH is *unaccompanied* in her walk through the forest, rather than that “It is false that LRRH is alone in her walk through the forest”. In fact, for the wide-scope negative reading, the entailments are obviously reversed: $\neg[\text{Alone} * \text{ForestWalk}] \models_{\neq} \neg[\text{Alone} * \text{EntireWalk}]$.

\sqsubseteq -ordering.

A transitive, reflexive relation $Actual \subset \mathcal{D} \times \mathcal{S}$ determines what individuals are actual with respect to a given situation. The *Actual* relation extends \sqsubseteq , since we would like to regard any part of a situation as 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. We assume that an individual that is nonactual with respect to a given situation is also nonactual with respect to any more inclusive situation.

Interpretations

A *model* $\mathcal{M} = \{\mathcal{D}, I\}$ in EL consists of a domain of individuals \mathcal{D} (structured as outlined above, with various additional constraints) and an interpretation function I that partially interprets individual constants and variables,¹⁰ function and predicate constants, predicate modifiers, and several other kinds of atoms (again subject to various constraints).

The most important aspect of any type of situation semantics is the semantics of predication, and how this provides the basis for truth-in-situations relative to a model \mathcal{M} . For our purposes, there are two alternative ways we could conceptualize predicate interpretations: as determining *characterizations* of certain basic situations in terms of atomic predications, or as directly determining *truth/falsity in* situations, for atomic predications. In the first approach we would say, for example, that $I(\text{sneeze})(d)(s) = 1, 0$, or is undefined respectively if s is a situation (episode) of individual d sneezing, a situation of d not sneezing, or neither. The notion of truth/falsity *in* a situation would then be derivative – for instance, if $I(\text{sneeze})(d)(s) = 1$, we would say it is true *in all situations s' more inclusive than s* (i.e., $s \sqsubseteq s'$) that d sneezes. Note that this would assure outward persistence of telic sentences.

Though the notion of interpretation in the first approach is intuitively very natural, we opt for the second approach, since this simplifies (and makes more nearly “conventional”) the connection between interpretations and truth. Thus we say, for example, that $I(\text{sneeze})(d)(s) = 1, 0$, or is undefined respectively if s is a situation (episode) *in which* individual d sneezes, one in which d doesn’t sneeze, or one where d ’s sneezing or not sneezing is not determinate. Here s need no longer be an episode *of* d sneezing in order for $I(\text{sneeze})(d)(s)$ to be 1 – rather, s might consist of many subepisodes, only one of which happens to be an episode of d sneezing. So in this case, persistence properties are presumed to be “built into” the interpretations of predicates. To begin with, we assure upward persistence of all predications (whether telic or atelic) by assuming that if a predication is true in s , then it is also true in s' for $s \preceq s'$. In addition, to ensure outward persistence of a telic predication like [Mary sneeze], we would assume that whenever $I(\text{sneeze})(d)(s) = 1$ holds, $I(\text{sneeze})(d)(s') = 1$ holds for any more inclusive situation s' (i.e., for $s \sqsubseteq s'$). In the same way we assume that the inward persistence of atelic predications and of negated predications is built into the interpretations of the relevant predicates. For instance, if we have $I(\text{alone})(d)(s) = 1$ (individual d is alone in situation s), this no longer means that s is a situation *of* d being alone, but only that it contains such a situation as a *coextensive part*; in other words, we assume that there is an $s' \preceq s$ which is a situation of d being alone and all of whose temporal segments are also situations in which d is alone (for all $s'' \preceq s'$, $I(\text{alone})(d)(s'') = 1$). Similarly, if we have $I(\text{alone})(d)(s) = 0$ (individual d is not alone in situation s), then we assume that there will also be a coextensive part $s' \preceq s$ (intuitively, that part or aspect of s which is the situation *of* d not being alone) all of whose temporal segments are also situations in which d is not alone.

Note that we have assumed above that we can apply the interpretation of a monadic predicate successively to an individual and a situation to obtain a truth value in $\{0, 1\}$. So the interpretations of ‘meet’ and ‘alone’, for example, are “curried” partial functions of type $\mathcal{D} \rightarrow (\mathcal{S} \rightarrow 2)$ (writing 2 for $\{0, 1\}$).¹¹ Upon applying such a function to an individual, we obtain a *sentence intension* – a partial 0, 1-valued functions on situations. In the same way, we interpret 2-place predicates as elements of $\mathcal{D} \rightarrow (\mathcal{D} \rightarrow (\mathcal{S} \rightarrow 2))$; etc. For instance, **greet** denotes an element of $\mathcal{D} \rightarrow (\mathcal{D} \rightarrow (\mathcal{S} \rightarrow 2))$, (**greet Mary**) denotes an element of $\mathcal{D} \rightarrow (\mathcal{S} \rightarrow 2)$, and ((**greet Mary**) John) (also written in “flattened,” infix form as [John greet Mary]) denotes an element of $\mathcal{S} \rightarrow 2$.

With this approach to predicate interpretation, it is the notion of characterization that becomes derivative. In other words, we need to specify the semantics of the ‘**’ operator in terms of truth *in* situations, since

¹⁰I.e., we do not separate variable assignments from interpretations.

¹¹Refer to footnote 2.

predicate interpretations no longer provide characterizations of situations in any direct way. (In a sense the indirect characterization of ‘**’-semantics becomes necessary anyway as soon as we consider characterization of situations by logically complex sentences.) We indicate in the next subsection how we do this.

Before proceeding, we should say a few words about the interpretation of atoms other than predicates. Two examples of nonstandard constructs for which our ample ontology provides direct interpretations are those for nominalizing (reifying) actions and sentences. In particular, if π is an action predicate (e.g., $\pi =$ (greet Mary)), with a denotation in $\mathcal{D} \rightarrow (\mathcal{S} \rightarrow 2)$, then $(\text{Ka } \pi)$ denotes an element of \mathcal{K}_A (a kind of action — in the example, the action of greeting Mary). Similarly if Φ is a sentence, then $(\text{That } \Phi)$ denotes an element of \mathcal{P} (i.e., a proposition). The abstract individuals obtained in this way can be “talked about” in EL like any others.

Truth Conditions

As was seen above, sentences are assigned denotations of type $\mathcal{S} \rightarrow 2$, i.e., a sentence may be true, false or truth-valueless in a given (possible) situation. The sentences which are true or false in a situation can be thought of defining its “information content.” A well-known advantage of this type of partial semantics is that it avoids the assumption of “omniscience” in the logic of knowledge, belief, and other attitudes; i.e., believers are not presumed to believe all the consequences of their beliefs.

Let’s indicate briefly how we arrive at truth values of sentences in situations, relative to a model $\mathcal{M} = \{\mathcal{D}, I\}$. First, given our “truth-based” (rather than “characterization-based”) semantics of predication, the truth conditions for an atomic sentence (where π is an n -place predicate and τ_1, \dots, τ_n are terms) are simply

$$\llbracket \pi(\tau_1) \dots (\tau_n) \rrbracket_{\mathcal{M}}^s = \left\{ \begin{array}{c} 1 \\ 0 \end{array} \right\} \text{ iff } I(\pi)(\llbracket \tau_1 \rrbracket_{\mathcal{M}}) \dots (\llbracket \tau_n \rrbracket_{\mathcal{M}})(s) = \left\{ \begin{array}{c} 1 \\ 0 \end{array} \right\},$$

where s is an arbitrary situation in \mathcal{S} and the $\llbracket \tau_i \rrbracket_{\mathcal{M}}$ are the denotations of the τ_i in model \mathcal{M} . (We omit the semantics of terms, except to mention that terms may have undefined denotations, but are rigidly interpreted, i.e., their values, if any, are independent of particular situations.)

As was seen in section on Basic Sentential Syntax, the ‘*’ operator allows the truth of an EL sentence relative to a situation to be expressed within EL itself, and this is what enables us to explicitly describe events, circumstances, etc., through sentences that hold in them. That ‘*’ does indeed correspond to truth in a situation can be seen from its semantics, which says that (for Φ a sentence, η a term, and s a situation in \mathcal{S}),

$$\begin{aligned} \llbracket \Phi * \eta \rrbracket_{\mathcal{M}}^s &= 1 \text{ iff Actual}(\llbracket \eta \rrbracket_{\mathcal{M}}, s) \text{ and } \llbracket \Phi \rrbracket_{\mathcal{M}}^{\llbracket \eta \rrbracket_{\mathcal{M}}} = 1; \\ &= 0 \text{ iff Nonactual}(\llbracket \eta \rrbracket_{\mathcal{M}}, s) \text{ or } \llbracket \Phi \rrbracket_{\mathcal{M}}^{\llbracket \eta \rrbracket_{\mathcal{M}}} = 0. \end{aligned}$$

The requirement that $\llbracket \eta \rrbracket_{\mathcal{M}}$ must be actual in order for Φ to be true in it makes ‘*’ (and indirectly, ‘**’) a *factive* operator; i.e., if $[\Phi * \eta]$ holds then η , and hence a subepisode of type Φ , must in fact have occurred. For instance, though we can in principle talk about a fictitious episode E , as soon as we assert $[[\text{Mary sneeze}] * E]$ we are committed to the reality of E and the actual occurrence of a subepisode of Mary sneezing in E .

The meaning of $[\Phi ** \eta]$ (“ Φ characterizes η ”) is similar to that of $[\Phi * \eta]$ but requires that η as a whole, rather than just some part of it, be of type Φ . Instead of giving a direct truth-conditional definition we treat ‘**’ as syntactically defined as follows. The definition says that η is either a minimal episode in which Φ holds, or it is comprised of temporal segments in all of which Φ holds, but none of which have coextensive proper parts in which Φ holds.

$$\begin{aligned} [\Phi ** \eta] &\equiv_{def} [[\Phi * \eta] \wedge (\forall e : [e \sqsubset \eta] \neg [\Phi * e])] \vee \\ &\quad (\forall e : [e \leq \eta] [[\Phi * e] \wedge (\forall e' : [e' \prec e] \neg [\Phi * e'])]) \end{aligned}$$

(For conciseness we have used the metalinguistic ordering relations ‘ \sqsubset ’, ‘ \leq ’, and ‘ \prec ’ here, where in our implementation we would use object language predicates like ‘proper-subep-of’, ‘subsegment-of’, and ‘proper-coexten-subep-of’.) For telic formulas, the definition simplifies to the first disjunct, and for atelic ones to the second. Many formulas – though not all – can be classified as telic or atelic. Without going into details, we assume that atomic predicates are dichotomized in this way, and that certain operators produce a telic or atelic result. For instance, activity predicates such as ‘walk’ and ‘sit’ are atelic, but when we

include a destination adverbial such as “to Grandmother’s cottage” or a duration adverbial such as “for an hour” in sentences based on atelic predicates, the result is telic. This is because a modified sentence such as “LRRH walked to Grandmother’s cottage” implies a culminated action, whereas “LRRH walked” does not.¹² On the other hand, application of the *progressive* operator ‘*prog*’ to a telic sentence (formula) produces an atelic result; e.g., “LRRH was walking to Grandmother’s cottage” is atelic. Probabilistic conditionals (which as explained are used for certain kinds of generic sentences) likewise produce an atelic result. A conjunction of a telic and atelic sentence is telic. Negation produces an atelic result, as does application of the ‘*’ and ‘**’ operators. In fact formulas of form $[\Phi * \eta]$ or $[\Phi ** \eta]$ are atemporal – they are true at all situations where they have a truth value, or false at all such situations.¹³

The semantics of logical connectives have a rather familiar look (*modulo* partiality), and we mention only two examples:

$$\begin{aligned} \llbracket \neg \Phi \rrbracket_{\mathcal{M}}^s &= \left\{ \begin{array}{c} 1 \\ 0 \end{array} \right\} \text{ iff } \llbracket \Phi \rrbracket_{\mathcal{M}}^s = \left\{ \begin{array}{c} 0 \\ 1 \end{array} \right\}; \\ \llbracket \Phi \vee \Psi \rrbracket_{\mathcal{M}}^s &= 1 \text{ iff } \llbracket \Phi \rrbracket_{\mathcal{M}}^s = 1 \text{ or } \llbracket \Psi \rrbracket_{\mathcal{M}}^s = 1; \\ &= 0 \text{ iff } \llbracket \Phi \rrbracket_{\mathcal{M}}^s = 0 \text{ and } \llbracket \Psi \rrbracket_{\mathcal{M}}^s = 0. \end{aligned}$$

We omit the truth conditions for conjunction (\wedge) and the material conditional (\rightarrow) since these involve some small complications to allow for the possible presence of anaphoric connections. These in turn depend on our slightly unconventional approach to \exists - and **The**-quantification. $(\exists \alpha: \Phi \Psi)$ has the expected semantics (intuitively, “Some value of α satisfying Φ satisfies Ψ ”) only if α does not have a value under the current interpretation I . If α does have a value, $(\exists \alpha: \Phi \Psi)$ is equivalent to $[\Phi \wedge \Psi]$. Analogous remarks apply to **(The** $\alpha: \Phi \Psi)$. Consequently, certain \exists - or **The**-quantified variables of a formula, called its *parameters*, can have their values “externally” supplied, and this allows us to deal with anaphora in the DRT-like manner we previously illustrated.¹⁴

Our semantics of \forall , **Most**, **Many**, etc., together with the semantics of ‘**’, leads to a conception of episodes with quantified characterizations as the *join* of a set of subepisodes of the type quantified over. For instance, in the sentences, “On her way to Grandmother’s cottage, Little Red Riding Hood chased every butterfly she saw. *That* took up half the trip,” the quantified episode consists of the join of chasing subepisodes, which may be separated by breaks of various lengths; that is what makes it possible for the second sentence (about the proportion of the time taken up by butterfly-chasing) to be true even if the time-stretch from the first to the last butterfly-chase covers the entire trip. Still, the truth conditions for \exists , \forall and standard connectives do not differ radically from “standard” ones (e.g., as in (Barwise, 1989; Devlin, 1991)).

We leave matters of ontology and formal semantics here, and proceed to our inference rules. (For further details on semantics, see (Hwang and Schubert, 1993b; Hwang, 1992), with the *caveat* that we have significantly altered our semantics for ‘*’ and ‘**’.¹⁵)

Inference Rules

We should begin by mentioning certain normalizing rules that we apply whenever possible:

¹²More precisely, we take the *basic* reading of “LRRH walked” to be non-culminated. There is also a culminated reading, tantamount to “LRRH took a walk”, but we take this reading to involve tacit application of a certain “episode-bounding” operator.

¹³Formulas like $[[\text{LRRH greet } w] \vee \neg [\text{LRRH silent}]]$ or $(\forall x: [x \text{ person}] [[x \text{ die}] \vee \neg [x \text{ die}]])$ are neither telic nor atelic. Still, if they are true in a given situation, they will satisfy the first or second disjunct of our definition for some part η of that situation. Which disjunct is satisfied depends on whether or not there is a telic fact among the atomic facts in virtue of which the formula is true in the given situation.

¹⁴It may also allow us to account for the dual existential/referential character of indefinites (cf. (Fodor and Sag, 1982)).

¹⁵The main change is that we have abandoned the notion that the situations characterized by NL sentences support the truth of just those sentences (and certain equivalent ones) and are atomic (have no parts). Rather, we regard those situations as potentially having an arbitrarily fine-grained part-structure and as supporting arbitrarily large amounts of “information”. This seems like a much more naturalistic notion of situations. For instance, we can now say that an extended episode such as an episode of LRRH being alone can have many (temporally smaller) parts, and all those parts are also episodes of LRRH being alone. (Previously we had to stipulate that an extended episode of LRRH being alone entails the existence of shorter episodes of LRRH being alone, at all times during the given episode.) We can also say now that *in* an episode of LRRH greeting the wolf, LRRH and the wolf are near of each other, instead of having to say that such a being-near episode exists at the same time as the greeting episode (but not as a part of it).

- Minimize the scope of negation. For instance, change

$$\neg(\forall x: [x \text{ person}] [x \text{ afraid-of Wolf}])$$
 to

$$(\exists x: [x \text{ person}] \neg[x \text{ afraid-of Wolf}]).$$
- Skolemize top-level existential variables (i.e., replace them by new constants). For instance, change

$$(\exists x: [x \text{ person}] \neg[x \text{ afraid-of Wolf}])$$
 to

$$[[C \text{ person}] \wedge \neg[C \text{ afraid-of Wolf}]].$$
- Separate top-level conjuncts. For instance, change

$$[[C \text{ person}] \wedge \neg[C \text{ afraid-of Wolf}]]$$
 to

$$[C \text{ person}], \neg[C \text{ afraid-of Wolf}].$$
- For formulas involving the atomic sentences \top (truth) or \perp (falsity), apply a set of simplifying rules. For instance, change $\neg\top$ to \perp , $\neg\perp$ to \top , $\Phi \vee \top$ to \top , $\Phi \wedge \top$ to Φ , $\Phi \vee \perp$ to Φ , $\Phi \rightarrow \perp$ to $\neg\Phi$, $(\forall \alpha : \perp \Phi)$ to \top , etc.

The main inference rules of EL are **Rule Instantiation (RI)** and **Goal Chaining (GC)**. These are generalizations of what are commonly referred to as “forward chaining” and “backward chaining” in AI terminology. RI is heavily used in input-driven inference, i.e., the process of elaborating the meaning and discourse significance of a new input sentence, in the light of meaning postulates, world knowledge, and prior discourse context. GC predominates in goal-driven inference, such as would occur during question-answering or discourse planning. It is also used in support of input-driven inference, typically to satisfy antecedents of input-triggered rules. We first illustrate the use of RI and GC and then state them more precisely.

Rule Instantiation (RI)

RI consists of two variant rules, each of which allows arbitrarily many minor premises to be matched against arbitrarily deeply embedded subformulas of a complex major premise. Though there are no formal constraints on the syntactic forms of the premises, the major premise will usually be an implicative and/or quantified formula. Such formulas are often called “rules” in the AI literature, hence the term “rule instantiation”. This creates some ambiguity in speaking of “rules”, since these may be inference rules or general formulas (rule-like knowledge), so in the following explanation of RI, we adhere to the “major/minor premise” terminology.

RI is related to the well-known rules of *modus ponens*, *modus tollens*, and *resolution*,¹⁶ but besides allowing for matching of arbitrarily many, arbitrarily deeply embedded subformulas, it can also instantiate probabilistic conditionals. The following (non-probabilistic) example illustrates the main features of RI. We first state the sample inference in English, and then in its logical format.

$$\frac{\begin{array}{l} \text{Every dress or hood that LRRH wears is pretty;} \\ \text{LRRH wears a certain cap or hood H} \end{array}}{\text{Therefore, if H is a dress or not a cap, it is pretty}}$$

$$\frac{(\forall x: [[[x \text{ dress}] \vee [x \text{ hood}]] \wedge [\text{LRRH wears } x]] [x \text{ pretty}]); \quad [H \text{ cap}] \vee [H \text{ hood}], [\text{LRRH wears H}]}{[[H \text{ dress}] \vee \neg [H \text{ cap}]] \rightarrow [H \text{ pretty}]}$$

The inference is obtained by two matching operations and several substitutions, as follows:

1. We match part of the first minor premise, namely its disjunct $[H \text{ hood}]$, against the embedded clause $[x \text{ hood}]$ of the major premise, recording the substitution $\{H/x\}$ (a substitution of a constant for a universally quantified variable).
2. We apply the substitution to the major premise, obtaining

$$[[[H \text{ dress}] \vee [H \text{ hood}]] \wedge [\text{LRRH wears H}]] \rightarrow [H \text{ pretty}].$$

Note that in the process of substituting for the universal variable, the restrictor and matrix of the universal formula become the antecedent and consequent respectively of a conditional formula. We will refer to this formula as the *converted major premise*.

¹⁶In particular see the embedded form of resolution employed in (Traugott, 1986); however, RI avoids Skolemization.

3. We now form the negation of the minor premise we used, after replacing the matched portion by \perp (falsity): $\neg([\mathbf{H \ cap}] \vee \perp)$, which is the same as $\neg[\mathbf{H \ cap}]$. We call this the *converted minor premise*.
4. We substitute the converted minor premise for the matched portion $[\mathbf{H \ hood}]$ of the converted major premise, obtaining

$$[[[\mathbf{H \ dress}] \vee \neg [\mathbf{H \ cap}]] \wedge [\mathbf{LRRH \ wears \ H}]] \rightarrow [\mathbf{H \ pretty}].$$
 We refer to this formula as the *intermediate result*. It is in fact a valid inference, but we are only half-way, since we also want to use the second minor premise.
5. Proceeding as in step (1), we match the second minor premise, $[\mathbf{LRRH \ wears \ H}]$, against the embedded clause $[\mathbf{LRRH \ wears \ H}]$.
6. Since no substitution is required, the analogue of step (2) is trivial, and the converted intermediate result is the same as the intermediate result.
7. Again we form the negation of the minor premise we used, with the matched portion replaced by \perp : $\neg\perp$, which is \top (truth). This is the new converted minor premise.
8. We substitute the converted minor premise (\top) for the matched portion $[\mathbf{LRRH \ wears \ H}]$ of the (converted) intermediate result, obtaining

$$[[[\mathbf{H \ dress}] \vee \neg [\mathbf{H \ cap}]] \wedge \top] \rightarrow [\mathbf{H \ pretty}].$$
 This simplifies to

$$[[\mathbf{H \ dress}] \vee \neg [\mathbf{H \ cap}]] \rightarrow [\mathbf{H \ pretty}],$$
 which is the inference delivered by RI.

One constraint tacitly observed in the above procedure is that in matching a part of a minor premise against a part of the major premise, these parts must occur *positively* and *negatively* in their respective formulas. A formula occurs positively in another if it is embedded within zero or more operators that create a positive embedding environment, and within an even number of operators that create a negative embedding environment. For instance, consider

$$[[\mathbf{H \ dress}] \vee \neg [\mathbf{H \ cap}]] \rightarrow [\mathbf{H \ pretty}].$$

The subformula $[\mathbf{H \ dress}]$ occurs negatively, since it lies within the scope of ‘ \vee ’ (which creates a positive embedding environment) and within the antecedent of ‘ \rightarrow ’ (which creates a negative embedding environment for its antecedent and a positive environment for its consequent). Similarly $\neg[\mathbf{H \ cap}]$ occurs negatively, while $[\mathbf{H \ cap}]$ occurs positively, since it is embedded by two operators that create a negative embedding environment, namely the conditional antecedent and the negation. $[\mathbf{H \ pretty}]$ occurs positively, since it lies in the consequent of the conditional, which is a positive environment. Additional relevant operators are conjunction (\wedge) and \exists -quantification, both of which create only positive environments, and \forall -quantification, which creates a negative environment in the restrictor and a positive environment in the matrix clause. Premises of RI may also involve probabilistic conditionals and quantifiers like **Most**, but we postpone discussion of probabilistic inference.

Steps (1-4) above for obtaining a conclusion from a major premise and a single minor premise can be concisely summarized as follows, writing $MAJ^-(\Phi)$ for a major premise with a negative occurrence of subformula Φ , and $MIN^+(\Phi')$ for a minor premise with a positive occurrence of subformula Φ' , where Φ and Φ' are matchable (unifiable):

$$\frac{MAJ^-(\Phi), \ MIN^+(\Phi')}{MAJ_\sigma^-(\neg(MIN_\sigma^+(\perp)))}$$

Here σ is the substitution that unifies (matches) Φ and Φ' . Steps (1-4) correspond to (1) forming unifying substitution σ , (2) forming the converted major premise $MAJ_\sigma^-(\Phi_\sigma)$, (3) forming the converted minor premise $\neg(MIN_\sigma^+(\perp))$, and (4) substituting the converted minor premise for the matched subformula Φ_σ in the converted major premise to obtain the conclusion, $MAJ_\sigma^-(\neg(MIN_\sigma^+(\perp)))$.

One point needing further clarification is the mechanics of matching (unification). A variable in a major or minor premise is matchable (i.e., we may substitute a term for it) if it is bound by a positively occurring universal quantifier or negatively occurring existential quantifier. For instance, substitution of w for x

is legal in a positively embedded subformula ($\forall x: [x P] [x Q]$), yielding $[[w P] \rightarrow [w Q]]$, and the same substitution is legal in a negatively embedded subformula ($\exists x: [x P] [x Q]$), yielding $[[w P] \wedge [w Q]]$.

The variant of RI we have been discussing turns out to be *sound* (yielding only true conclusions from true premises) if the matched subformula Φ' in the minor premise contains no unmatchable free variables which are bound in $MIN^+(\Phi')$ as a whole. So in particular, the rule is sound if $MIN^+(\Phi')$ contains only constants and top-level universal (hence matchable) variables. In certain cases where the condition for soundness is violated, we can apply another variant of RI which interchanges the roles of the major and minor premises in the conclusion, as follows:

$$\frac{MAJ^-(\Phi), MIN^+(\Phi')}{MIN^+(MAJ^-(\Phi))}.$$

This variant is sound if the matched subformula Φ in the major premise contains no unmatchable free variables which are bound in $MAJ^-(\Phi)$ as a whole.

Suppose, for instance, that our disjunctive minor premise in steps (1-4) above had been existentially quantified:

$$(\exists y [[y \text{ cap}] \vee [y \text{ hood}]])$$

In this case if we attempt to match $[y \text{ hood}]$ against $[x \text{ hood}]$ in the major premise, we observe that y is free in $[y \text{ hood}]$ but is unmatchable since it is bound by a positively occurring existential quantifier at the top level. Thus we are not allowed to apply the first variant of RI. However, we can apply the second variant, and the reader can verify that the result at step (4) is

$$(\exists y [[y \text{ cap}] \vee [[LRRH \text{ wears } y] \rightarrow [y \text{ pretty}]]]).$$

Goal Chaining (GC)

GC is a pair of very general goal reduction rules, analogous to the two variants of RI; however, instead of deriving a conclusion from a major premise and one or more minor ones, we derive a subgoal from a major premise, possibly some minor premises, and a given goal. Chaining from consequents to antecedents of quantified or unquantified conditionals is a special case. An example that closely parallels the RI-example would be the following. Note that the goal is to prove that H is pretty, which we write as $?[H \text{ pretty}]$:

$$\begin{array}{c} \text{Every dress or hood that LRRH wears is pretty;} \\ \text{LRRH wears a certain cap or hood H} \\ \text{Goal: Is H pretty?} \\ \hline \text{Subgoal: Is H a dress or not a cap?} \end{array}$$

$$\begin{array}{c} (\forall x: [[[x \text{ dress}] \vee [x \text{ hood}]] \wedge [LRRH \text{ wears } x]] [x \text{ pretty}]); \\ [H \text{ cap}] \vee [H \text{ hood}], [LRRH \text{ wears } H] \\ \text{?[H pretty]} \\ \hline \text{?[[H dress] \vee \neg [H cap]]} \end{array}$$

In essence, what we are doing here is to match the goal $?[H \text{ pretty}]$ to the consequent $[x \text{ pretty}]$ of the universal conditional, chaining back to the intermediate subgoal

$$?[[[H \text{ dress}] \vee [H \text{ hood}]] \wedge [LRRH \text{ wears } H]].$$

But in this backward chaining, we are also allowed to use any number of minor premises to reduce the new goal. In the present case we can use the first minor premise to replace $[H \text{ hood}]$ with $\neg[H \text{ cap}]$ in the intermediate subgoal, and the second minor premise to delete $[LRRH \text{ wears } H]$ from it (technically, replacing it with \top). This gives the final subgoal shown above.

A point of difference from RI is that in matching a (part of a) goal to a part of a premise we use different notions of “matchable variables” in goals and premises. In premises, the matchable variables – the ones we are allowed to substitute for – are defined as before; but in a goal, the matchable variables are those that are bound by *positively* occurring *existential* quantifiers or *negatively* occurring *universal* quantifiers. This plays no role above since the goal contains no variables; but it is easy to see that an existential goal like

$$?(\exists y [y \text{ pretty}])$$

should be satisfiable by a premise like `[H pretty]`, and this involves unifying the existential variable `y` with the constant `H`.

For completeness we give the formal statements of the two goal chaining rules for the nonprobabilistic case, with no minor premises. We will then go through another detailed example.

$$\frac{MAJ^+(\Phi), ?GOAL^+(\Phi')}{?\neg(MAJ_\sigma^+(\neg(GOAL_\sigma^+(\top))))} \quad \frac{MAJ^+(\Phi), ?GOAL^+(\Phi')}{?GOAL_\sigma^+(\neg(MAJ_\sigma^+(\perp)))}$$

where σ unifies Φ with Φ' . The first rule is *sound* if Φ' contains no unmatchable free variables which are bound in $GOAL^+(\Phi')$ as a whole (e.g., a variable bound by a top-level universal quantifier). The second rule is sound if Φ contains no unmatchable free variables which are bound in $MAJ^+(\Phi)$ as a whole (e.g., a variable bound by a top-level existential quantifier).

We have so far suppressed episodic variables in explaining **RI** and **GC**. In view of the importance of such variables in our semantic representation of NL sentences, we now give a detailed illustration of goal chaining (with use of supplementary premises) based on episodic formulas. Consider the following general “explanatory” axiom:

$$\begin{aligned} &(\forall x: [x ((\text{attr predatory}) \text{ animal})] (\forall y: [y \text{ creature}] \\ &\quad (\forall e_1: [[y \text{ near } x] ** e_1] \\ &\quad\quad (\forall e_2: [e_2 \text{ during } e_1] \\ &\quad\quad\quad [[x \text{ attack } y] ** e_2] \\ &\quad\quad\quad\quad \rightarrow (\exists e_3: [e_3 \text{ same-time } e_2] \\ &\quad\quad\quad\quad\quad [[x \text{ hungry}] ** e_3] \vee [[x \text{ enraged}] ** e_3])))))) \end{aligned}$$

A predatory animal attacks a nearby creature only when it is hungry or enraged.

Note that `x`, `y`, `e1` and `e2` are matchable variables. Suppose we want to know *if the wolf was ever enraged*. Then this goal can be posed as

$$?(\exists e_4: [e_4 \text{ before } \textit{Now}] [[\text{Wolf enraged}] ** e_4]),$$

where we observe that `e4` is a matchable variable. Since the goal has no unmatchable variables, we use the first **GC** rule. Note that the matrix of the goal matches the second disjunct in the consequent of the general axiom, with substitution $\{\text{Wolf}/x, e_3/e_4\}$. So applying the first **GC** rule, we obtain the following new goal (after simplifying and distributing negation):

$$\begin{aligned} &?[[\text{Wolf } ((\text{attr predatory}) \text{ animal})] \wedge \\ &\quad (\exists y [[y \text{ creature}] \wedge \\ &\quad\quad (\exists e_1 [[y \text{ near } \text{Wolf}] ** e_1] \wedge \\ &\quad\quad\quad (\exists e_2 [[e_2 \text{ during } e_1] \wedge \\ &\quad\quad\quad\quad [[[\text{Wolf attack } y] ** e_2] \wedge \\ &\quad\quad\quad\quad\quad (\forall e_3: [e_3 \text{ same-time } e_2] \\ &\quad\quad\quad\quad\quad\quad [[e_3 \text{ before } \textit{Now}] \wedge \neg [[\text{Wolf hungry}] ** e_3]])]]))]]))]. \end{aligned}$$

Suppose now that our knowledge base contains the axiom

$$(\forall x: [x \text{ wolf}] [x ((\text{attr predatory}) \text{ animal})])$$

as well as the particular fact `[Wolf wolf]`. Then the initial conjunct `[Wolf ((attr predatory) animal)]` of our goal formula will be reduced *via* the first axiom to `[Wolf wolf]` and this in turn will immediately be eliminated *via* the second axiom. Thus, we are left with subgoal

$$\begin{aligned} &?(\exists y [[y \text{ creature}] \wedge \\ &\quad (\exists e_1 [[y \text{ near } \text{Wolf}] ** e_1] \wedge \\ &\quad\quad (\exists e_2 [[e_2 \text{ during } e_1] \wedge \\ &\quad\quad\quad [[[\text{Wolf attack } y] ** e_2] \wedge \\ &\quad\quad\quad\quad (\forall e_3: [e_3 \text{ same-time } e_2] \\ &\quad\quad\quad\quad\quad [[e_3 \text{ before } \textit{Now}] \wedge \neg [[\text{Wolf hungry}] ** e_3]])]]))]]). \end{aligned}$$

The new goal asks, “Did the wolf attack a nearby creature sometime in the past, but was not hungry?”¹⁷ Suppose now the knowledge base contains facts: $[[\text{Wolf attack Fox}] ** E]$, $[[\text{Fox near Wolf}] * E]$, and $[E \text{ before Now}]$. Then the question could be further simplified to

? $(\forall e_3: [e_3 \text{ same-time } E] \neg [[\text{Wolf hungry}] ** e_3])$.

If this cannot be answered, then we would go back to the previous goal and attempt to prove it using other facts.

The probabilistic version of RI produces conclusions that are annotated with a lower bound on the certainty (degree of belief) of those conclusions. The bounds are computed as a product of the bounds associated with the minor premises used and, if the major premise is a probabilistic conditional, with the numeric strength of that conditional. This is not quite as crude as it sounds, since some provision is made to avoid repeated use of the same evidence to strengthen (or weaken) belief in the same conclusion. (This is done by keeping track of support sets in the inference process.) Also for linear inference chaining using a nonrepetitive sequence of simple probabilistic conditionals, the computed probability bounds conform the probabilistic semantics in (Bacchus et al., 1996). When multiple proofs or disproofs are found for the same proposition, with various lower bounds on the probabilities that the conclusion is true/ false, these probability bounds are combined through a multiple-evidence scheme essentially like the “noisy-OR” technique in Bayes nets (Pearl, 1988). (This involves some rather crass independence assumptions.)

Forward inference chaining using RI is terminated when the *expected interestingness* of the conclusions being drawn falls below a threshold value. Thus, for instance, we would tend to pursue the consequences of LRRH being attacked, but would be unlikely to reason that LRRH is a person and therefore has a head, and also must have a mother who also has a head, and so on. The expected interestingness of a proposition is the product of its interestingness and its lower bound on certainty. Predicates, individuals (constants), and propositions (sentences) all have interestingness ratings. Those of predicates are currently pre-set; for instance action predicates are generally rated as more interesting than atelic ones, and of course some actions, like marrying or attacking someone, are rated higher than others, such as walking or resting. Among atelic predicates, being a person is more interesting than being a rock, and being terrified is more interesting than being comfortable. The ratings of individuals and propositions evolve as information accumulates. The idea is that an individual is interesting to the extent that we know a lot of interesting facts about it; and in turn, a proposition is interesting to the extent that it involves interesting individuals and predicates. This may sound circular, but in fact can be implemented consistently. We also allow for inheritance of interestingness from premises to consequences, and from effects to their causes (i.e., causes of interesting effects are apt to be interesting themselves). Salience in context might be expected to be important as well, but some preliminary experiments suggested this may not be particularly important for inference termination.

RI and GC do most of the work needed to generate immediate consequences of new inputs and to answer questions. However, for question-answering there is also another class of goal-directed methods consisting of standard natural deduction rules such as proving a conjunction by proving the conjuncts, proving a conditional by assuming the antecedent and proving the consequent, and proving a disjunction by proving one of the disjuncts while assuming the negation of the others. These rules are used for breaking down a given goal, forming a goal tree whose leaves are then tackled by using GC. For rules that use assumption-making, the assumptions may be used to trigger forward inferencing *via* RI; assumptions (and their consequences) are retracted once a proof attempt has terminated.

With the kinds of EL inference rules described so far, EPILOG is able to make some quite complex inferences and to answer questions based on logically represented simple narratives or telegraphic messages (Namioka et al., 1992). The control structure for question answering (for questions presented as logical goals) may be sketched as follows. For a given question, simultaneous attempts are made to prove the corresponding goal and its negation. (For Wh-questions, variable bindings are tracked in these attempts.) An agenda containing potential knowledge-access actions and goal-chaining actions for subgoals at the leaves of the current goal trees is used to prioritize the steps of the proof and disproof attempts. Knowledge-access actions are guided by a systematic, automatically maintained classification of all formulas in terms of keys consisting of $\langle \text{predicate, argument} \rangle$ or $\langle \text{predicate, argument type} \rangle$ pairs (with an indication of the argument’s

¹⁷The fact $(\forall e (\forall e_1 (\forall e_2 [[e_2 \text{ same-time } e_1] \rightarrow [[e_2 \text{ before } e] \leftrightarrow [e_1 \text{ before } e]]])))$ would also be needed eventually.

role), and by “climbing” *type hierarchies* in which these arguments or argument types participate. (We will show some sample classifications under Implementation Status and Test Scenarios.) Multiple factors are taken into account in the prioritization of agenda items; in the case of goal-chaining actions these include: the location (including depth) of the subgoal in its goal tree; whether the proposed goal-chaining action matches an antecedent or consequent clause of the selected major premise; and the interestingness and complexity of the subgoal.

In addition, an important feature of the inference process is the use of multiple *specialists* to provide fast inferences about taxonomies, times, parts, sets, etc. These can greatly accelerate proof attempts by “evaluating” and simplifying certain kinds of terms and formulas in derived clauses or goals, and by directly detecting inconsistencies between certain kinds of subformulas (e.g., incompatible types in a type hierarchy, or cycles in a set of temporal ordering relations), where we might otherwise need lengthy disproofs. These remarks bring us close to implementation issues, about which we will have a little more to say in the section on Implementation Status and Test Scenarios.

Simulative Inference

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, we need to develop ways of enabling a story understanding system to make inferences about mental processes by simulation. In other words, the system should be able to temporarily treat the beliefs of another agent as if they were its own, then “observe” what further beliefs it would derive from those assumed, and then ascribe those additional beliefs to the other agent.

This appealing idea has a considerable history in AI, with some studies aimed at developing logically rigorous models of sound simulative inference (e.g., (Creary, 1979; Haas, 1986; Konolige, 1986)), and others leaning more toward practical goals (e.g., (Moore, 1977; Ballim and Wilks, 1991; Chalupsky and Shapiro, 1996)). Kaplan and Schubert (Kaplan and Schubert, 1997; Kaplan, 1998) offer a thorough formal analysis of simulative inference in a computational setting. The model is based on viewing belief retrieval and augmentation in terms of an ASK-TELL mechanism that operates on an agent’s belief store (which can take any form, not necessarily a set of formulas). ASK is an algorithmic query mechanism that returns “yes” or “no” for any query formula, indicating respectively that the formula is believed or not believed. (A formula that is not believed need not be disbelieved, i.e., it may be that neither the formula nor its negation is believed.) TELL is an algorithmic belief augmentation mechanism that attempts to add a formula to the agent’s beliefs. (This may fail, e.g., if contradictions are encountered.) The main results in the cited references concern the conditions on ASK and TELL under which simulative inference is sound, and there are also restricted completeness results.

We will not discuss simulative inference in detail here, since its implementation for EL remains largely a research issue. We merely mention that we envisage an implementation in two parts, one relying on a goal-driven ASK mechanism and the other on an input-driven TELL mechanism. ASK and TELL would be much like the goal-driven and input-driven inference mechanisms we have already described, except that they would make only very simple inferences, to assure fast termination. (Keep in mind that ASK is intended as a model of belief retrieval, not problem solving.) Goal-driven simulative inference would be triggered by goals of form

$$[[\alpha \text{ believes } \beta] * \eta],$$

(i.e., in a certain situation η , agent α believes proposition β), and would consist of an attempt to “evaluate” the goal to truth or falsity, by running ASK on query β within a *belief space* for agent α in situation η . (More accurately, the query would be a formula Φ , assuming that proposition $\beta = (\text{That } \Phi)$). It may also be feasible to have the simulation return subgoals (e.g., beliefs to be proven by ordinary inference) in cases where the result is neither truth nor falsity. In a system that reasons about beliefs, a belief space for another agent is a way of configuring or clustering the known beliefs of that agent so as to facilitate

reasoning about them. Belief space mechanisms have been incorporated into many knowledge representation systems, including our EPILOG system. For the purpose of simulative inference, the crucial computational requirement is that running ASK or TELL in the belief space of another agent should yield precisely the same results as if the beliefs of the other agent were the system’s own, i.e., as if they had been stripped of the belief-wrapper `[α believes ...]` and integrated into the system’s own belief space.¹⁸

Input-driven simulative inference would be triggered by storage of formulas of form

`[[α learns β] ** η],`

i.e., there is some event η that consists of some agent α learning (coming to believe, from some external source) proposition β . In this case simulative inference would consist of an attempt to add β (or more exactly Φ , as above) to the belief space for α , using the TELL mechanism. In general this will trigger a cascade of further inferences, and the output of the simulation would consist of “significant inferences” observed in this simulative use of TELL. For instance, if upon asserting (TELLing) Φ in α ’s belief space, Ψ is inferred and this is rated as a sufficiently “interesting” inference, then something like the formulas

`[[α infers (That Ψ)] ** η_i],`

`[η cause-of η_i],`

would be included among the outputs of the simulation. This expresses the prediction that α infers that Ψ holds (and thus will *believe* that it holds) *as a result* of learning β (i.e., that Φ holds).

We think this sort of prediction is particularly important in story understanding. When humans learn new information, they often become instantly aware of certain significant consequences of that information, and potentially act on those consequences. For instance, a person encountering a dangerous animal in the woods (thus “learning” about its proximity) would immediately think of the possibility of being attacked, and hence might take some appropriate action (freezing, fleeing, etc.). So to anticipate and understand the behavior of story characters, we need to actively anticipate their thoughts, and this would be greatly facilitated by input-driven simulative inference.

We should emphasize that simulative inference, though potentially extremely useful, cannot be a stand-alone method of reasoning about beliefs. Observe, for instance, that we cannot apply simulative inference to premises

`[[A believes P1] \vee [A believes P2]], \neg [A believes P2]`

to obtain

`[A believes P1],`

even though this is a trivial deduction. So the way to view simulative inference is as special “attachment” techniques (in the terminology of (Konolige, 1986)) that are integrated into the regular goal-driven and input-driven inference mechanisms, reducing certain formulas to truth or falsity (or to subgoals), and adding many belief inferences that would be very difficult to obtain by ordinary logical inference.

Many theoretical and practical difficulties will need to be dealt with in further work on simulative inference in EL. One is that the theory of propositions in our situational logic does not quite fit with the computational model of belief in (Kaplan and Schubert, 1997), or other sentence-based theories. For instance, in EL a compound (but quantifier-free) sentence is semantically indistinguishable from logically equivalent sentences built up out of the same constituents. (E.g., `[$\neg\Phi \vee \Psi$]` is semantically indistinguishable from `[[$\Phi \wedge \neg\Psi$] \rightarrow [$\neg\Phi \vee \Psi$]].) So belief in one is the same as belief in the other. While it is easy to particularize the computational model to conform with such a constraint, we may not want this, and in that case our situation theory would require significant changes. Another problem is the time-dependence of beliefs. On the one hand, we want an understanding system to be able to “time-stamp” beliefs (or in our case, “situation-stamp” them), since beliefs can and do change. On the other hand, we want to assume by default that a belief is still held if there is no reason to think it has been abandoned. In essence, this is the frame problem for beliefs. It is also unclear under what conditions simulative inference in a logic as expressive as EL will be sound; or even for the first-order subset of EL how we can ensure that ASK and TELL (goal-driven and input-driven inference) will satisfy the conditions for soundness identified in (Kaplan and Schubert, 1997). From a practical perspective, the main problem is how to make ASK and TELL belief-space independent in the sense required by simulative inference.`

¹⁸One issue that arises is how to treat self-references by the other agent.

(e_0 : LRRH started off for Granny's)

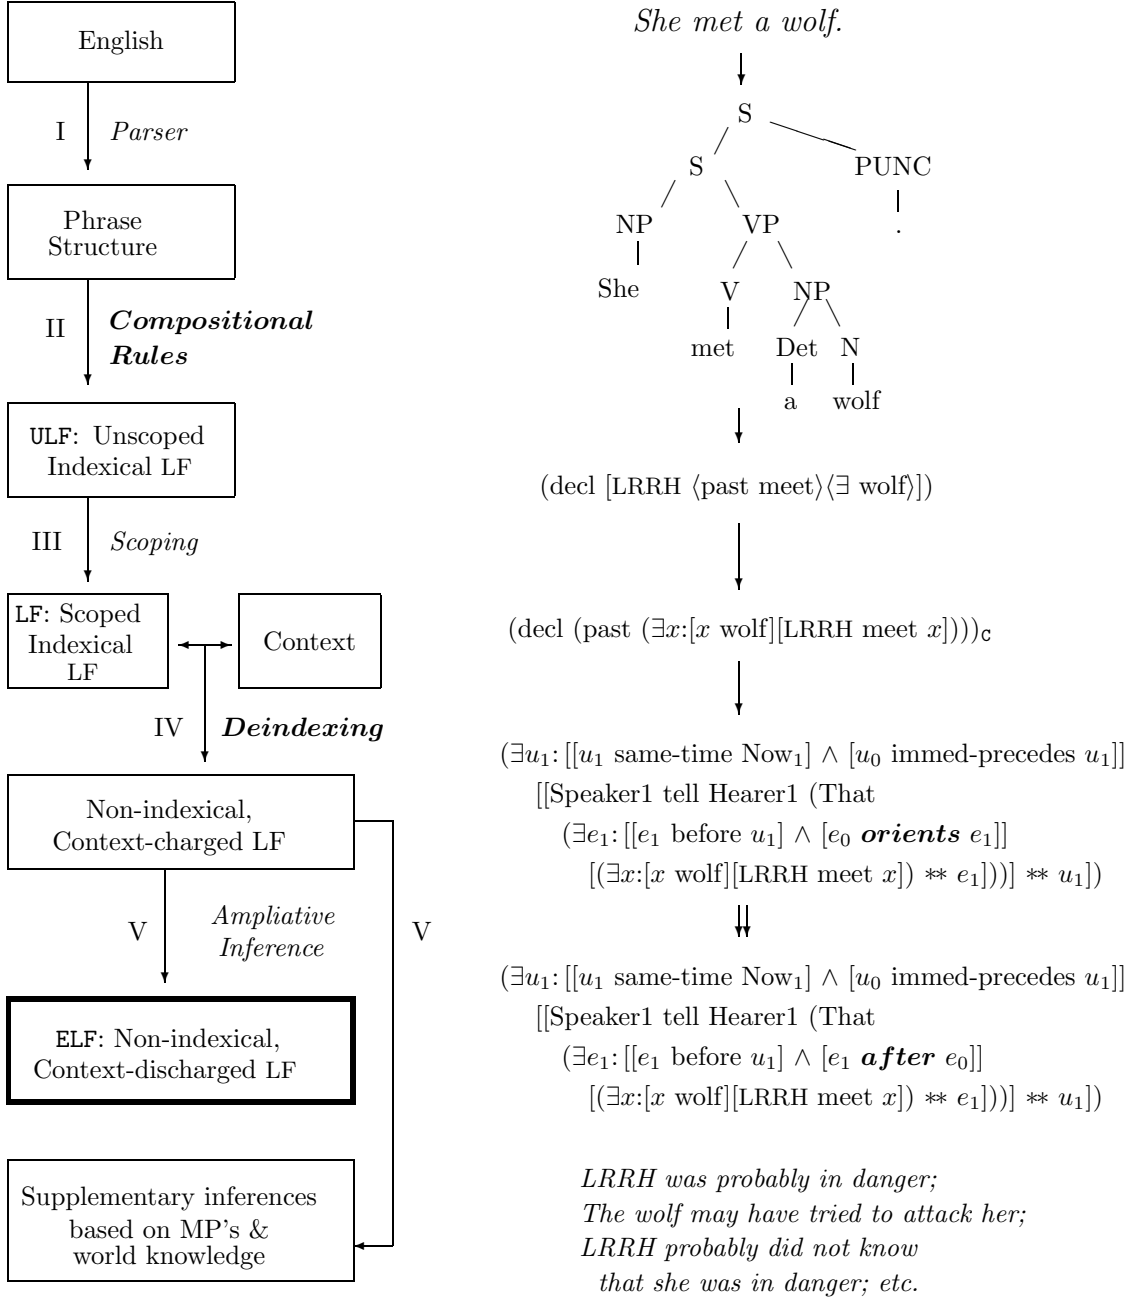


Fig. 2. The Conceptual Stages of NL Understanding

LF (16). The subscripted **C** indicates the explicit context structure with respect to which the scoped LF is to be interpreted. Among other things, this consists of a “tense tree,” whose purpose is to facilitate context-dependent tense-aspect interpretation, a “clock” which generates a succession of *Now*-points for speech times, and hearer and speaker parameters.

Stage IV: Deindexing

The scoped, indexical translation is deindexed with respect to this context **C** in stage IV, so as to obtain a nonindexical logical form usable for inference. The computation of the nonindexical ELF from the LF is driven by a simple, recursive deindexing mechanism that makes use of the tense tree in context structure **C**. The deindexing rules handle tense, aspect, and many temporal PP-adverbials and their interaction; their effect is to bring 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. Note that u_0 in Fig. 2 (stages IV and V logical forms) is the utterance episode of the previous sentence, i.e., (13), and e_0 is the episode introduced by it, i.e., that of Little Red Riding Hood’s starting off for Grandmother’s cottage. Now_1 is the speech time of sentence (14). Application of appropriate deindexing rules transforms LF (16) into ELF (17) as shown below.

$$\begin{aligned}
 (17) \quad & (\exists u_1: [u_1 \text{ same-time } Now_1] \wedge [u_0 \text{ immediately-precedes } u_1]) \\
 & \quad [[\text{Speaker1 tell Hearer1 (That} \\
 & \quad \quad (\exists e_1: [e_1 \text{ before } u_1] \wedge [e_0 \text{ orients } e_1]) \\
 & \quad \quad \quad [[e_1 \text{ in-loc Forest}] \wedge \\
 & \quad \quad \quad \quad (\exists x: [x \text{ wolf}] \wedge \\
 & \quad \quad \quad \quad \quad (\exists e_2: [e_2 \text{ at-about } e_1] \\
 & \quad \quad \quad \quad \quad \quad [(\exists e_3: [e_3 \text{ impinges-on } e_2] \\
 & \quad \quad \quad \quad \quad \quad \quad [[e_3 \text{ lasts-for (K ((num 3) (plur day)))] \wedge \\
 & \quad \quad \quad \quad \quad \quad \quad \quad (\neg [x \text{ eat}])] ** e_3)] \\
 & \quad \quad \quad \quad \quad \quad \quad \quad \quad ** e_2)]] \\
 & \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad [LRRH \text{ meet } x])] \\
 & \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad ** e_1)]] \\
 & \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad ** u_1)].
 \end{aligned}$$

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.”

This deindexing mechanism is compositional in the sense that operators **pres**, **past**, **futr**, **perf**, etc., contribute separately and uniformly to the meanings of their operand formulas, driving the generation and traversal of tense trees in deindexing. We cannot include a detailed introduction to tense trees here. Instead, we will have to confine ourselves to a sketchy intuitive exposition, an example of the notation we use for deindexing rules, and pointers to our prior writings on the subject ((Hwang and Schubert, 1992; Hwang, 1992; Hwang and Schubert, 1993b)).

Tense trees “grow” downward, though one tree can “embed” another *via* certain horizontal links. They can be thought of as being generated as a byproduct of a depth-first traversal of the (indexical) logical form (LF) of a sentence (viewed as a nested list structure), in the course of deindexing that LF. In fact they reflect the tense-operator structure of the LF: each branch signifies the embedding of some clause of the LF by an operator; and the *direction* of a branch indicates the “temporal orientation” of the corresponding operator. In particular, the branches generated by **past**, **perf**, and **futr** operators point down to the left, straight down, and down to the right respectively. (Imagine time as progressing from left to right.) The tense “tree” for a single sentence is typically just a zig-zag *path*; it is the layering of such paths on top of each other for successive clauses or sentences that may lead to multiple branches at a single node, particularly as a result of tense changes.

What’s the use of a tense tree? In the course of its generation and subsequent re-traversals *episode tokens* are generated and placed at the tree nodes, and temporal relationships among tokens are automatically asserted, capturing the deindexed meaning of the tense operators as well as interclausal and intersentential connections. The actual work is done by a set of deindexing rules, one for each operator, named **Pres**, **Past**,

Futr, **Perf**, etc., accordingly. Each rule can be viewed declaratively as consisting of an *equivalence* and an *equation*. The following deindexing rule for **past** is the only one we will look at:

$$\text{Past: } (\text{past } \Phi)_T \leftrightarrow (\exists e_T: [[e_T \text{ bef}_T \text{ Emb}_T] \wedge [\text{Last}_{\nearrow T} \text{ orients } e_T]] [\Phi_{\circ \nearrow T} ** e_T])$$

$$\text{Tree transformation: } (\text{past } \Phi) \cdot T = \uparrow (\Phi \cdot (\circ \nearrow T))$$

The equivalence specifies what $(\text{past } \Phi)$ *means*, in terms of the episode tokens in a given tense tree T . (Note the introduction of the characterization operator ‘**’!) $[e_T \text{ bef}_T \text{ Emb}_T]$ says that the new episode e_T characterized by Φ is *at* or *before* (“ bef_T ”) Emb_T , the episode at the nearest node embedding T (dominating the root) – usually this is the speech (or narration) event. $[\text{Last}_{\nearrow T} \text{ orients } e_T]$ says that the last event that was stored at the past-node (before e_T was placed there) provides a “point of orientation” for e_T (cf., (Leech, 1987; Webber, 1988)). We comment further on the ‘**orients**’ relation below.

The equation in the **Past** rule states a structural constraint that a tense tree T must satisfy if it is generated by the LF $(\text{past } \Phi)$. But we can equally well view it as a recursive tree-modification rule, reading the right-hand side “inside-out, right-to-left”, like a Lisp expression: the arrow operators indicate tree traversal in the direction of the arrow, and the open dot dictates the placement of a new episode token e_T at the node reached. More exactly, the equation prescribes a left downward traversal from the current node (generating a new branch if necessary), followed by placement of new episode token e_T at the node reached, followed by letting Φ have its effect on the tree (this is guaranteed to bring us back to the same node), followed by return to the start node.

Only four operators actually induce downward branching in tense trees, viz., **past**, **perf**, **futr** and **fpres** (future present, as in “LRRH won’t recognize the wolf when she *arrives* at Grandmother’s”). **pres** adds an episode token but does not cause branching. Horizontal embedding branches are generated for the surface speech act (the act of telling the hearer something, asking something, etc.) and for subordinate clauses (e.g., ones headed in the LF by the proposition-forming operator **That**). For most other operators (e.g., **prog**, and predicates), deindexing simply means moving the dependence on the tense tree T inward to their operands – there is neither a syntactic transformation nor an effect on the tense tree. However, as we have seen, adverbials of certain sorts (e.g., temporal and locative ones) are an important exception; they are syntactically transformed into conjoined predications about episodes (much as in traditional approaches, such as (Reichenbach, 1947; Davidson, 1967; Dowty, 1982)). For details on the interpretation and deindexing of temporal and other adverbials, see (Hwang, 1992; Hwang and Schubert, 1993c; Hwang and Schubert, 1994).

We conclude our look at tense trees with the following figure, showing how the tense tree resulting from processing (13) is extended when (16) (the LF for (14)) is deindexed.




Fig. 3. Tense Tree Transformation

The left diagram shows the tense tree induced by (13); the right one shows the final tree after (16) has been completely deindexed, yielding (17). The dashed link corresponds to the two embedding speech acts (the “telling that” implicit in the two narrative sentences, (13) and (14)), and the dotted link corresponds to the syntactic embedding of the relative clause “who had not eaten for three days”.

Stage V, Part 1: Narrative Inferences

Let us return to the *orienting* relation $[e_0 \text{ orients } e_1]$ in (17). This is generated by the **Past** rule, and asserts that e_0 , the event of LRRH starting out, is the “point of orientation” for the subsequently described event e_1 of LRRH meeting a wolf. (Note that e_0 , e_1 were placed next to each other in the tense tree above; this is what enables the generation of the orienting relation.) Orienting relations contribute to narrative

coherence, and their automatic derivation is one of the most important benefits of the tense tree mechanism. However, ‘*orients*’ does not quite have the status of a logical predicate; rather, it is what we call a *context-charged* relation. Such relations “suggest” various possibilities (e.g., various possible temporal or causal relations), and the idea is that their meaning is to be “discharged” through *narrative inferences*. These are nondeductive (probabilistic or default) inferences comprising part of stage V of the understanding process, the ampliative inference stage. Narrative inferences hypothesize alternative meanings for context-charged or otherwise ambiguous expressions in the LF using various features of the current LF and previous ELFs to assign *a priori* likelihoods to these alternatives. Other inference processes, termed “implicit question-answering” in the next subsection, are assumed to perform the final adjudication among various alternatives so as to arrive at an overall sentence interpretation “coherent with” the already interpreted prior discourse. In that sense, the meaning of context-charged relations still depend on discourse context, though not on the explicit context structure \mathcal{C} .

In our example, $[e_0 \text{ orients } e_1]$ 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 characterizations of e_0 (LRRH’s starting off) and e_1 (her meeting a wolf) and the circumstances described, this hypothesis would be deemed the most probable, but in other cases the most probable particularization of the *orients* relation may be a subepisode relation, a causal or explanatory relation, or any of the discourse relations that have been discussed in the literature.

Note that besides the *orients* relation, (17) contains another context-charged relation, namely $[e_3 \text{ impinges-on } e_2]$. This is generated by the deindexing rule for the *perf* operator, and is intended to be particularized into either *until* or *before* through narrative inferences. It relates the episode e_3 of the wolf’s not eating for three days (reported in past perfect tense) to the past reference point e_2 (which coincides temporally with the meeting event e_1). The fact that the characterization of e_3 is negative and hence atelic is taken to provide strong evidence for the *until* interpretation of *impinges-on*; i.e., the not-eating episode lasts until the reference time, and hence until the meeting. Thus, taking for granted the final adjudicative process (or assuming that the feature-based evidence is in this case already decisive), (17) would be particularized to

$$(18) (\exists u_1: [[u_1 \text{ same-time Now}_1] \wedge [u_0 \text{ immediately-precedes } u_1]] \\ \text{[[Speaker1 tell Hearer1 (That} \\ (\exists e_1: [[e_1 \text{ before } u_1] \wedge [e_1 \text{ after } e_0]] \\ [[e_1 \text{ in-loc Forest}] \wedge \\ (\exists x: [[x \text{ wolf}] \wedge \\ (\exists e_2: [e_2 \text{ at-about } e_1] \\ [[\exists e_3: [e_3 \text{ until } e_2] \\ [[e_3 \text{ lasts-for (K ((num 3) (plur day)))] \wedge \\ (\neg [x \text{ eat}]]] ** e_3)] \\ ** e_2)]]) \\ [LRRH \text{ meet } x]]] \\ ** e_1]])) \\ ** u_1]).$$

Note that $[e_0 \text{ orients } e_1]$ has been particularized into $[e_1 \text{ after } e_0]$, and $[e_3 \text{ impinges-on } e_2]$ into $[e_3 \text{ until } e_2]$.

Although we assumed at the beginning that the referent of the pronoun “she” in (14) had been resolved to LRRH in (15), in actuality it would be resolved in stage (V), simultaneously with the discharging of context-charged relations. In fact, we conjecture that reference resolution could naturally be accomplished in a manner paralleling the derivation of temporal relations. First we would add a kind of anaphora deindexing to the temporal deindexing in stage IV, consisting of augmenting anaphoric expressions (terms corresponding to pronouns and definite descriptions) with lists of readily accessible antecedents. Accessibility would be determined from the structure of the LF of the current sentence and the ELFs of prior sentences. Then in stage V, narrative inferences aimed specifically at reference resolution would assign prior probabilities to the accessible antecedents based on features of the current LF and prior ELFs, and “implicit question answering” would perform the final adjudication.

Stage V, Part 2: Adjudication through Implicit Q-A

The narrative inferences posited above should be viewed as part of a broader stream of input-driven inferences, triggered by the information in the current LF and prior ELFs in conjunction with meaning postulates and world knowledge. These more general inferences are indicated by the second “branch” of stage V indicated in Fig. 2, leading to “supplementary inferences”; some sample inferences we would expect to obtain are indicated at the bottom right of the figure.

However, the problem is that in general the inference stream is not univocal. Rather, there will be various *alternative* ways to extrapolate inferentially from the current LF and prior ELFs, depending both on alternative ways of discharging context-charged relations, and alternative kinds of world knowledge that can be brought to bear. For example, the pronoun “she” in (14) could theoretically refer to Grandmother rather than LRRH, and such an interpretation would lead to very different inferences. Also, instead of conjecturing that the hungry wolf was a hazard to LRRH, we might alternatively conjecture that he would take an interest in the cake in LRRH’s basket. In fact, such an inference would be essential for making sense of a continuation such as “The wolf slyly inquired whether he could help LRRH carry the basket”.

While we need not sort out all such alternatives immediately to achieve understanding (and may never need to sort out some of them), we do clearly need to adjudicate among alternative forward inferences in order to achieve a *coherent* interpretation of a narrative. We think that our earlier description of the EL inference mechanisms provides a reasonably clear and specific picture of how forward inferencing might work; but we have so far offered few clues about how to do the sorting-out that achieves global coherence. We suggest that this cannot be primarily a matter of deciding what global picture is most “plausible” in terms of our *world* knowledge, since coherence is in a sense a “contractual” matter between narrator and reader (or speaker and hearer), not just a matter of how the world is. In other words, the narrator is under a conventional obligation to *make* the story cohere, and the reader relies on this convention in choosing among alternative inferences.

Of course, some *a priori* biases will be introduced by the more or less likely particularizations of context-charged relations and the more or less likely conclusions based on world knowledge in the form of probabilistic conditionals. As well, the mechanisms we have mentioned for combining probabilistically weighted evidence for or against a conclusion reached *via* multiple inference pathways will cause some belief shifts. However, in general such processes will merely lead to adjustment of degrees of belief in various alternatives, not to a “crystallization” of a particular global interpretation. The following is a speculative discussion of a process we call “implicit question answering” (IQA), which we regard as the key to arriving at a coherent interpretation.

The idea is that 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* with less than complete certainty 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 confirm or disconfirm it (to some degree). The assumptions and inferences that provided the confirming or disconfirming answer are then considered correct. The key point for us (and the reason for our “question-answering” terminology) is that *the narrator is free to answer the tacitly raised questions 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.

A positive example of IQA

We begin with an illustration of the positive case:

- (19) a. John dropped the glass on the floor.
b. It broke.

Inference chaining based on (19a), along with axioms about “dropping,” “glasses,” and about fragile objects striking hard surfaces would quickly lead to a rather probable (but not certain) prediction that the glass broke. In our terminology, therefore, the question of *whether the glass broke* is raised. Now in (19b), the pronoun may *a priori* refer to either the glass or the floor. If it refers to the floor, then (19b) neither confirms nor disconfirms any predictions triggered by (19a). But if it refers to the glass, then (19b) directly supports

the predicted breakage of the glass. Thus according to the IQA principle, the latter interpretation is chosen. Note also that a tentative inference from the orienting relation computed for (19a) and (19b), namely, that the breaking was right after, and caused by, the dropping will thereby be confirmed as well (assuming that such an immediately following causal consequence was predicted from (19a)).

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 expectations 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 abductive approaches of Charniak and Goldman (Charniak, 1988; Charniak and Goldman, 1988) and Hobbs *et al.* (Hobbs et al., 1993) 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 *disconfirmation* 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.¹⁹

A negative example of IQA

A suitable illustration is (19a) plus the denial of (19b):

- (20) a. John dropped the glass on the floor.
b. It didn’t break.

In this case, it is the *denial* of a prediction from (20a) that the glass broke, i.e., that the glass *didn’t* break, that is supported by (20b), with the pronoun resolved to “the glass.” Again, resolving the pronoun to “the floor” neither confirms nor disconfirms any questions raised by (20a), and hence that possible way of resolving the pronoun is not confirmed. By contrast, approaches like those of Charniak and Goldman or 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 general knowledge certainly supports the conclusion that *the floor* didn’t break, the pronoun would be resolved to refer to the floor.²⁰

A final example of IQA

The examples above are simpler than anything likely to be encountered in real narratives. The (a)-sentence forcefully suggests a particular prediction (breakage of the glass), and the (b)-sentence directly confirms or disconfirms that prediction. More commonly, narrative sentences tend to evoke a variety of possible explanations for the reported situation or event, and a variety of possible predictions. So implicit question-answering in general involves searching for corroborative or antagonistic connections between tentative explanations and predictions evoked by a new sentence and those evoked by prior sentences. The following example is more realistic.

- (21) a. John heard steps behind him.
b. He began to run.

A spontaneous explanatory inference from (21a) is likely to be that there was someone behind John, quite close to him (and he knew this). In turn, this leads to the (possibly very tentative) conjecture that John may believe himself in danger, and may try to get away from the person behind him. Other, more benign possibilities are that the person behind John wishes to catch up and communicate with him, or simply intends to walk past him. If the latter are the most likely possibilities in John’s mind, we would not expect any special actions from him, except perhaps to attend to the approaching walker. (Of course, prior context

¹⁹The proposals in (Lascarides and Asher, 1991; Lascarides et al., 1992) do invoke discourse conventions to sort out interpretations of narratives, instead of making this a matter of world knowledge (alone), but as far as we can tell would not deal properly with negative examples.

²⁰It might be countered that the resolution of the pronoun in both (19) and (20) is the result of “centering,” where the verb object in (a) is the preferred center. However, this is disconfirmed by “John accidentally dropped the cutting-board on the glass. Fortunately, it didn’t break.”

may disable one or another inference chain, or suggest less conventional ones. Also, the inference of danger from (21a) 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 (21b) also suggests multiple alternative explanations, considered in isolation: John (the only possible referent for *he* in this case) may simply be in a hurry, or he may be trying to get away from someone or something near him, or he may be exercising. (These seem like the most probable explanations.) Once again, only one of these possible inferences, namely the second, bears on any of the questions raised by (21a). In particular, this inference confirms that John is trying to get away from someone near him, and hence that inference, and everything leading up to it, and the relevant interpretation of (21b), are deemed correct.

We leave the discussion of implicit question-answering here. Developing the details of such a theory, and implementing it in EL, remains a major challenge for the future.

To conclude the discussion of the stages of understanding, we emphasize again that though the stages have been described as if they ran sequentially, and the start-up implementations in which EL has played a role were in fact sequential (viz., the TRAINS system (Allen and Schubert, 1993; Traum et al., 1996) and the message processing application (Namioka et al., 1992)), the intention is to interleave them 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.

Implementation Status and Test Scenarios

While there is as yet no complete story understanding system that includes all of the modules and processes described in the previous sections, the EL knowledge representation along with the input-driven and goal-driven inference mechanisms have been fully implemented in the EPILOG system (Miller et al., 1991). Also, various aspects of our approach have been incorporated into two complete prototype NLP systems: the TRAINS 91-93 system (Allen and Schubert, 1993; Traum et al., 1996) and a message processing application at Boeing company (Namioka et al., 1992). In addition, we have carried out detailed analyses of several story excerpts, and verified that EPILOG is able to perform many of the inferences necessary for understanding.

We begin with an overview of EPILOG, continue with brief descriptions of the two prototype systems, and then discuss our work on story fragments at some length, concluding with a note on the frame problem in narratives.

The EPILOG System

The EPILOG system is the practical knowledge representation and inference system for EL, and represents the culmination of many years of implementation work on a succession of NL-oriented representations. It allows for the full EL syntax, and besides the input-driven and goal-driven inference mechanisms, deploys an array of specialist subsystems to help with both of these inference modes (see (Miller and Schubert, 1988; Gerevini and Schubert, 1995) and further references therein). As shown in Fig. 4, there is a type specialist for efficiently handling facts such as that wolves are animals, using a preorder numbering scheme on tangled type hierarchies; a part-of specialist, also based on numbered hierarchies; a time specialist using a “timegraph” structure for fast transitive inference; a color specialist based on a cylindrical color space; specialists for efficient string and set manipulation, arithmetic, etc. The specialists are used for immediate simplification of terms and formulas when these are stored, and for “generalized resolving”, e.g., the detection of inconsistencies between simple formulas such as $[x \text{ wolf}]$ and $\neg[x \text{ animal}]$, or $[H \text{ red}]$ and $[H \text{ green}]$. Some of these specialists, such as the time and set specialists, are dynamic in the sense that they accumulate information in their own specialized data structures as new inputs are presented. The specialists are uniformly tied into the general inference mechanisms *via* a specialist interface, making it easy to add new specialists. This interface also enables indirect communication between specialists, so that they can help each other without knowing of each other’s existence. Note also the response generator indicated in the figure, which is able to give rough-and-ready verbalizations in English of EL formulas.

An important design goal for EPILOG was to assure efficient use of *relevant* knowledge in an inference task even if the knowledge base is very large. This is accomplished by two kinds of knowledge indexing, one

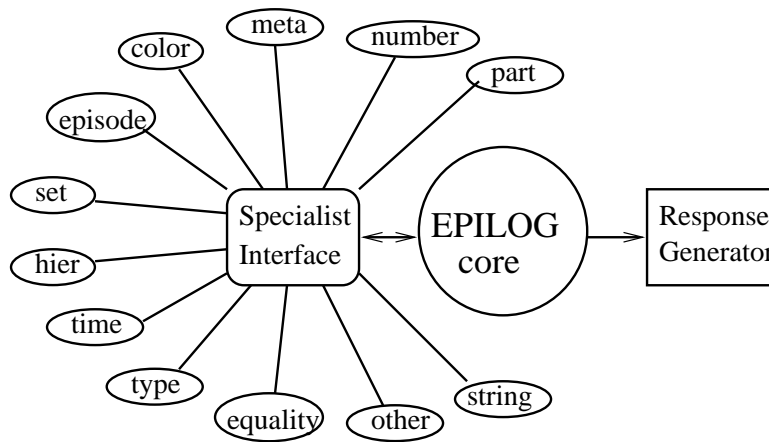


Fig. 4. EPILOG Architecture

based on *predicate-argument* pairs and the other on *topic-argument* pairs. Predicate-argument indices are the primary means of accessing relevant formulas in both input-driven and goal-driven inference. They are based on classifying formulas in terms of the more “interesting” atomic subformulas occurring within them, using as keys the predicates in those subformulas paired with one argument or argument type. For instance, a formula describing what happens when a predator attacks another creature might contain a subformula $[x \text{ attack } y]$, where x is typed as a ‘predator’ and y as a ‘creature’, through quantifier restrictions. In that case, the formula as a whole would be hash-indexed under $(\text{predator attack subj})$ and $(\text{creature attack obj})$, indicating that the formula involves predicate ‘attack’ with a subject of type ‘predator’ and an object of type ‘creature’. EPILOG classifies input formulas automatically in this way, and then uses this classification to find formulas to match against subformulas in Rule Instantiation and Goal Chaining. The restriction of the classification to the more “interesting” atomic subformulas is important in that it prevents triggering an excessively broad inference process; for instance, we do not want to classify the formula about predators attacking other creatures using the subformula $[y \text{ creature}]$, as this would cause EPILOG to “think of” attacks by predators whenever *any* creature is mentioned.

Predicate-argument indexing for belief formulas of form $[\alpha \text{ believe } (\text{That } \Phi)]$ uses extended key lists, obtained by adding key words corresponding to α and *believe* to those extracted from Φ ; similarly for other modal predicates, such as ‘hope’, ‘say’, etc. (We also regard fictional tales as lying within the scope of a modal predicate; e.g., $[\text{LRRH-story has-it } (\text{That } \dots)]$.) This facilitates uniform inference about the beliefs (etc.) of other agents, in a manner paralleling inference based on its own beliefs. (However, it is not clear whether the parallelism is sufficiently exact to provide a basis for simulative inference.)

Topic-argument indices consist of an individual or type paired with a topic, e.g., $(\text{LRRH } \text{tp.appearance})$, or $(\text{wolf } \text{tp.feeding})$. They provide a way of generating descriptive output, as an end in itself or when a question cannot be answered exactly. For instance, they allow EPILOG to access and output everything it knows about LRRH’s appearance or emotional attitudes, or about a wolf’s feeding habits. If a question like “Does LRRH like the wolf?” cannot be answered exactly, the related information “LRRH likes animals”, “LRRH loves Grandmother”, etc., (if previously supplied) could in principle be returned. Topic-argument indexing is based on user-supplied *topic indicators* for non-type (i.e., non-noun-like) predicates. There is a predetermined (but extensible) hierarchy of several dozen topics under which facts about any given thing or type of thing can be classified, such as appearance-related topics (color, shape, texture, etc.), behavior-related topics (such as feeding, locomotion, etc.), function/use (for artifacts), and so on. Topic indicators specify what topic a predicate indicates with respect to each of its argument positions. For instance, ‘eat’ indicates *tp.feeding* with respect to the subject and *tp.existence* with respect to the object (since eating something terminates or at least threatens its existence).

In the course of its development, EPILOG has been tested on numerous simple deduction and probabilistic inference problems, some of which will be mentioned in the subsections that follow. Though it was designed for shallow inference on large knowledge bases rather than for “deep” reasoning, EPILOG can handle some quite challenging theorem proving problems, including the “Steamroller” (see (Stickel, 1986)), when the hierarchy information in this problem is placed in a type hierarchy.

Two Prototype Systems

The first prototype system we will describe made use of EL as a *semantic* representation, but not directly for inference. The second system emphasized the use of the EPILOG inference engine, adhering only very loosely to our conception of how to derive a semantic representation from English.

TRAINS 91-93 (Allen and Schubert, 1993; Traum et al., 1996) helped a user solve simple problems in transporting goods by rail (including simple manufacturing steps like having a factory produce orange juice from oranges) in a small simulated railroad network. The 1991 system used a GPSG-like grammar and parser and semantic rules associated one-to-one with phrase structure rules to derive unscoped logical forms from user inputs (and also from the system’s own responses); these were then scoped to obtain indexical LFs, and deindexed with the aid of tense trees to obtain final nonindexical ELFs. Since there was a pre-existing frame-like knowledge representation for domain problem solving, ELFs were not employed directly for inference in this system, but rather were used as sources of information for subsequent modules aimed at speech act analysis, dialogue management, and planning. TRAINS 93 was similar in its architecture, but used an integrated scoping, referent determination and tense deindexing module, and improved modules further “downstream”. For a discussion of the various knowledge representations used, the pros and cons of a nonuniform approach, and further references see (Traum et al., 1996). From our EL perspective, a particularly gratifying aspect of the “TRAINS experience” was that EL as a semantic representation of natural language, and our method of deriving ELFs, proved readily applicable in an interactive problem solving domain very far removed from the story understanding scenarios that motivated us originally. We take this as evidence for the generality and domain-independence of our approach to semantic representation.

The Boeing message processing application was developed for the Boeing Commercial Airplane Reliability and Maintainability (R&M) division (Namioka et al., 1992). The goal was to partially automate the process of extracting information from the thousands of telexes received monthly by the division, reporting problems and repairs on the Boeing fleet. The start-up system used a GPSG parser, originally developed by P. Harrison and subsequently adapted for use on the telegraphic messages (Jenkins et al., 1990). The phrase structure rules used were well-adapted to computing approximate ELFs, though there was not a full-fledged version of a semantic interpreter. The features of EL that made it particularly suitable for this application were its expressiveness, its direct connection to language, its natural handling of causation, and its probabilistic rules, allowing expert system-like knowledge encoding. For example, EL made it easy to represent such notions as “reporting that ...” (i.e., modal operators), “severe wear” (i.e., modifiers), “an impact by some object causing formation of a crack” (causation), and rules (generalizations) such as

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

Moreover, the input-driven and goal-driven inference mechanisms of EPILOG were well-suited to extraction of the desired information from messages.

The ultimate limitations of both the TRAINS-93 system and the R&M prototype system were the fact that large amounts of linguistic and world knowledge are needed to deal with even a small set of dialogues or messages. Particularly gaps in linguistic knowledge are apt to cause failure, leading to very brittle systems. Recent NLP systems have striven to overcome brittleness through partial parsing and partial interpretation. In addition, the world knowledge “bottleneck” is often finessed by specializing systems so that they seek out particular kinds of information from texts or dialogues in very restricted domains. To the extent that the kinds of information sought are highly predictable and schematized, many syntactic and semantic details can be ignored. Such approaches can be practically rewarding, but it seems to us that genuine, domain-independent language understanding ultimately depends on dealing successfully with those syntactic and semantic details. We expect that the problem of acquiring many of the details will ultimately yield to automated learning techniques – but a framework as rich as EL will still be required to accommodate them.

Story Fragments

Since our primary interest is in the use of EL in story understanding, we have examined several story fragments in detail, to determine how readily the inferences that seem to be made spontaneously by people in the comprehension process could be generated by EPILOG from EL representations of the fragments.²¹

²¹For some further details see <http://www.cs.rochester.edu/research/epilog/>

Most of these tests were carried out without the benefit of a NL front end that could handle all of the grammatical, semantic, scoping and deindexing issues (not to mention ambiguities) involved in the chosen examples, and so the ELFs used as inputs to EPILOG were hand-constructed. In the initial trials, these ELFs were not exactly the forms we would expect to obtain mechanically, for various reasons: our conception of the ELF was still evolving; in some cases the *frame problem* posed difficulties (e.g., an exploding bomb ceases to be a bomb, but most static properties persist; knowing which ones do can be crucial); some involved simulative inference, which was and remains unimplemented; and in some cases more accurate representations would have required extensive use of schematic *meaning postulates* (MPs) involving predicate variables, sentence variables, etc., and the mechanisms for applying such MPs were still under development. We will discuss three examples, two of them very sketchily since they were based on “inaccurate” ELFs, and the third, whose representation is faithful to our theory of logical form, in a little more detail.

A principle we firmly adhered to in working out these examples is to specify the requisite knowledge at the most *general* level possible. This is extremely important, since it is all too easy to invent highly specific rules that will appear to deal successfully with any given, small example, or even with certain narrow *classes* of examples. For instance, it would be a “cheat” to have a rule (conditional formula) for the Little Red Riding Hood story that says that if a child encounters a wolf, the wolf will make plans to eat the child. A more defensible rule would be one that says that if a predator encounters a nonpredatory creature no larger than itself, and is either hungry or enraged, it may try to attack (and perhaps subdue and eat) that creature. Many semi-practical systems that extract semantic content from certain types of stories (such as reports on business mergers or terrorist incidents) depend heavily on using rules that work only for those types of stories. Our interest, however, is in domain-independent language understanding.

Terrorist story excerpt

As a first example, we briefly discuss an excerpt from a *Wall Street Journal* news item:

- (23) An explosives-laden car blew up near the office of an [Afghan] guerrilla group in a crowded Shiite Moslem neighborhood in Beirut.

Simplifications we allowed ourselves here were the treatment of “explosives-laden” and “Shiite Moslem” as single lexical items; conjunctive interpretation of “explosives-laden car” (i.e., as a quantifier restriction $[[x \text{ explosives-laden}] \wedge [x \text{ car}]]$ rather than $[x ((\text{attr explosives-laden}) \text{ car})]$) and similarly for “crowded Shiite Moslem neighborhood”; and the temporal linking of the episode of the car being explosives-laden with the episode of the office existing as such, and the latter episode with the episode of the neighborhood being crowded.

Some inferences that are indicative of understanding (23) are the following:

- *A car was destroyed, and so (probably) was an office, and some furniture;*
- *Probably someone was killed, and hence was dead.*

EPILOG made these inferences, among others, based on one MP and 12 general facts, including

- *Anything in a location where a bomb explodes is probably destroyed;*
- *Anyone who is destroyed is killed;*
- *Anyone who is killed becomes dead;*
- *If something is an office, there is a good chance that some furniture, a computer, and documents are in it;*

etc. Besides drawing the obvious conclusions by input-driven inference, EPILOG also was able to answer questions framed in terms of quantifiers such as “Was there something near the location where the explosion took place?”, “Was no-one killed?”, etc., posed as ELF query formulas. With additional rules, EPILOG’s understanding of the excerpt could have been deepened; for instance, EPILOG could have inferred the target organization with a rule to the effect that detonation of an explosive device at or near the premises of some individual, group or organization indicates that some hostile agent or agency intended this as an attack against the individual, group or organization.

The wolf and the woodcutters

One fragment of the Little Red Riding Hood story that we studied at length is the following.

- (24) 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, genuine understanding of the third sentence entails being able to explain why the wolf decided against eating *Little Red Riding Hood*, and how the presence of woodcutters nearby affected the wolf's decision. In particular, one has to know that when some agent *dares* not do something, he must think it possible that his attempt to do it would have adverse consequences for himself; then one has to simulate his reasoning process to guess what unpleasant consequences he anticipates.

Since we have not implemented simulative inference, we did not aim at full comprehension of the passage but rather at generating a set of plausible inferences under the hypothesis that the wolf attempts to eat LRRH, with the woodcutters nearby. In addition to this hypothesis, EPILOG was given several dozen pieces of general knowledge. Among the inferences it generated in input-driven mode were ones leading to the conclusion that the wolf may be severely punished. The following summarizes the relevant inference chain (consisting of about 30 inference steps) in words (Schubert and Hwang, 1989). We state major intermediate conclusions, with parenthetical indications of the steps leading to them.

- *The wolf's attempt to eat LRRH is an instance of the action type, "trying to eat LRRH" (based on a meaning postulate about the relation between action instances and action types).*
- *Trying to eat LRRH involves attacking her (because trying to eat a living creature involves attacking it in order to subdue it).*
- *There will be an instance of the wolf attacking LRRH during his attempt to eat her (from an axiom about what it means for one action to involve another).*
- *The wolf's attack on LRRH is an extremely wicked action (because for any creature to attack a child is extremely wicked).*
- *The attack is conspicuous for a nearby woodcutter (because there are some woodcutters nearby, and woodcutters are human, and neither the wolf nor LRRH are tiny relative to the woodcutter, and for a sizable creature to attack a sizable thing is conspicuous to a human being).*
- *The woodcutter may notice the attack (because if a person is near something, that thing is likely to be within plain sight of the person, and a conspicuous action within plain sight of a person is apt to be noticed by the person).*
- *The wolf may be severely punished (because if a human being notices something extremely wicked going on, the wrongdoer may eventually be severely punished).*

A simulative inference process would now ascribe this reasoning to the wolf, and conclude that he anticipates possible punishment if he were to act on his desire to eat LRRH, thus explaining why he dared not eat her. Of course, we do not claim that the above inference chain is somehow uniquely correct. Changes to the knowledge base could have produced other variations, but we suggest that *something* like the above is needed.

We do think it possible to understand the passage at a more superficial level with a less elaborate inference chain. We could just reason that the wolf must perceive some hazard from the proximity of the woodcutters, if he were to attempt to eat LRRH, and leave it at that. But it is our sense that most human readers of the story "fill in" a more detailed scenario, and so it is important for an automated system to be able to do likewise. Imagine asking a reader, "Why do you suppose the wolf feared the nearby woodcutters?" The answer would surely be something like, "Well, the woodcutters might go after the wolf and perhaps kill him", which presupposes much of the reasoning above.

The wolf enters Grandmother’s cottage

The story excerpt for which we have performed the most nearly complete analysis and knowledge base construction is the following

- (25) The wolf drew out the peg and the door flew open. Then he sprang upon the poor old lady and ate her up in less than no time, for he had been more than three days without food.

We should also provide the immediately preceding sentence, since we use some facts from it (without doing a full analysis): *The worthy grandmother was in bed, not being very well, and cried out to him, “Pull out the peg and the latch will fall.”* We will first outline a set of inferences needed for understanding, then present the logical forms of the sentences in the fragment and the representations of a few bits of background knowledge needed to produce those inferences.

Again the above excerpt may be explained in various ways by humans. Similarly, the system could in principle explain the story in several ways, depending on the kind of background knowledge supplied. The following is a clause-by-clause outline of the inferences enabled by the knowledge base we built. EPILOG was able to make most of the inferences mentioned, except for those based on Implicit Question Answering (IQA), which (as we explained) are beyond the system’s current capabilities. Since IQA is needed to resolve references, we helped the system along by manually supplying the equations that resolve the referring expressions, at the point of need. In this way we were able to “walk” the system through the passage.

1.1 “The wolf drew out the peg, ...”

- *The latch will fall (since Grandmother said so).*
- *The door will then be ready to be opened, and the wolf could push it open (causing it to become open) and enter the cottage.*
- *The wolf’s reason for drawing out the peg was to bring about the above events. (We assume the system already knows that the wolf wanted to enter the cottage as the story mentioned earlier that the wolf knocked on the door.)*

1.2 “. . . and the door flew open.”

- *This confirms that the door became open, as anticipated. By the IQA principle, this also confirms that the wolf pushed it open, though the narrator’s failure to mention this tends to undercut the inference (i.e., it may have opened spontaneously).²²*
- *It is very likely that the wolf entered the cottage when the door opened (since the wolf had wanted the door to open so that he could enter the cottage).*
- *When he entered the cottage, he was near Grandmother and alone with her in the cottage, and she was weaker than he (being old and ill), and he knew this.*
- *Since the wolf was probably wicked, he would probably use the opportunity to do something bad to Grandmother, in particular, attack her (knowing she was weaker than him).*

2.1 “Then he sprang upon the poor old lady ...”

- *His “springing upon” her suggests — if we view him as a person — that he will assault and perhaps kill her, since this is one way of initiating an assault. But since this is also an instance of a predator springing upon a weaker creature, he may be preparing to subdue, kill, and eat her. In either case, the expectation that the wolf will do something bad to Grandmother is confirmed, and the assessment of the wolf as very wicked is strengthened.*

2.2 “. . . and ate her up in less than no time ...”

- *His eating her confirms one of the previous predictions. By the IQA principle, this also confirms he first subdued and killed her, though the narrator’s failure to mention this tends to undercut the inference.*

²²We do not at present deal with this “undercutting” phenomenon.

- *Furthermore, his eating Grandmother very quickly suggests that he was either very hungry or in a great hurry.*

2.3 “. . . for he had been more than three days without food.”

- *No-one can eat if he has no food, and a creature will be very hungry if he hasn’t eaten for a day or more. That confirms one of the alternative explanations of the wolf’s haste.*
- *We now know that the wolf is extremely wicked, and that LRRH is likely to arrive at the cottage shortly, and that the wolf knows this, and that he may therefore want to do something bad to LRRH upon her arrival, and hence may stay at the cottage to await her.*

After having been “walked through” the excerpt, EPILOG is also able to answer the following kinds of questions with very little effort. (Questions are shown in EL, followed by English glosses of the EL questions and answers).

?(Wh e: [e episode] (The e’: [e cause-of e’] [[Wolf pull-out Peg] ** e’]))
Why did the wolf pull out the peg?
(Because he wanted to enter Grandmother’s cottage.)

?(∃e: [[Wolf hungry] ** e]
 (The e’: [e right-after e’] [[Wolf eat-up Granny] ** e’]))
Was the wolf hungry right after the episode of his eating up Grandmother?
(Probably not.)

Note that the answer to the second question is qualified with an epistemic probability (e.g., .75), as it is based on an uncertain inference. That is, the system predicts that the wolf was probably not hungry after eating Grandmother, since we have an axiom stating that if a creature eats a large portion of food, it is likely to be full immediately afterwards, with minimum degree of certainty .75.

We now show how the excerpt is translated and give a small sampling of the background knowledge required for understanding the story as above. We conclude with some remarks on the frame problem. Readers not interested in technical details may safely skip the rest of this subsection.

Story representation

We show the initial, unscoped logical forms (ULFs), scoped but still indexical logical forms (LFs) and completely deindexed episodic logical forms (ELFs), in that order. We omit speech act operators.

1. *The wolf drew out the peg and the door flew open.*

⇒ ⟨∧ [⟨The wolf⟩⟨past draw-out⟩⟨the peg⟩
 [⟨The door⟩ (⟨past fly⟩ open)]]
 ⇒ [(past (The x₁: [x₁ wolf] (The x₂: [x₂ peg] [x₁ draw-out x₂])) ∨
 (past (The x₃: [x₃ door] [x₃ (fly open)])))]
 ⇒ [(∃ e₂₁: [[e₂₁ before Now₁₁] ∧ [E₁₀ orients e₂₁]] [[Wolf pull-out Peg] ** e₂₁]) ∧
 (∃ e₂₂: [[e₂₂ before Now₁₂] ∧ [e₂₁ orients e₂₂]] [[Door (fly open)] ** e₂₂]])

2. *Then he sprang upon the poor old lady and ate her up in less than no time, for he had been more than three days without food.*

⇒ [⟨∧ ((adv-e at-time-Then)
 [He ⟨past spring-upon⟩⟨The ((attr poor2) ((attr old) lady))⟩]
 [He ((adv-a (in-manner (very quick))) (⟨past eat-up⟩ Her))⟩]
 because ⟨past (perf ((adv-e (lasts-for (K ((num 3) (plur day))))))
 [He without (K food))⟩]⟩]
 ⇒ [[[past (The x₄: [x₄ ((attr poor2) ((attr old) lady))]]
 ((adv-e at-time-Then) [He spring-upon x₄]])] ∨
 (past [He ((adv-a (in-manner (very quick))) (eat-up Her))]]]

A meaning postulate about monotone predicate modifiers

MP 6. For α a monotone predicate modifier (such as ‘very’) and π a monadic predicate,

$$(\forall x [[x (\alpha \pi)] \rightarrow [x \pi]])$$

Two pieces of relevant world knowledge

WK 1. If a predatory animal springs upon a nonpredatory creature, that probably is part of an attack intended to subdue the creature, allowing it to be killed and eaten.

$$\begin{aligned} & (\exists x: [x ((\text{attr predatory}) \text{ animal})] (\exists y: [[y \text{ creature}] \wedge (\neg [x \text{ predatory}])] \\ & \quad (\exists e_1: [[x \text{ spring-upon } y] ** e_1]))) \\ \rightarrow_{.85, e_1} & (\exists e_2: [[e_2 \text{ involves } e_1] \wedge [[x \text{ attack } y] ** e_2]] \\ & \quad [[x | e_2] \text{ done-with-intention} \\ & \quad (\text{Ka } ((\text{adv-a } (\text{for-purpose } (\text{Ka } \lambda z (\text{seq } [z \text{ kill } y] [z \text{ eat } y]))) \\ & \quad (\text{subdue } y)))))) \end{aligned}$$

% ‘seq’ is an operator that maps n formulas, $n \geq 1$, into a sequence of formulas.

WK 2. Creatures are very hungry when they have not eaten for more than a day.

$$\begin{aligned} & (\exists n: [[n \text{ number}] \wedge [n \geq 1]] \\ & \quad (\exists e_1: [e_1 \text{ lasts-for } (\text{K } ((\text{num } n) (\text{plur } \text{day}))))] \\ & \quad \quad [(\exists x: [x \text{ creature}] (\neg [x \text{ eat}])) ** e_1])) \\ \rightarrow_{.9, n, e_1} & (\exists e_2: [[e_2 \text{ at-end-of } e_1] \wedge [e_1 \text{ cause-of } e_2]] \\ & \quad [[x (\text{very hungry})] ** e_2]) \end{aligned}$$

A note on the frame problem

One problem that arises quite frequently in narratives is that we need to assume the persistence of certain properties or relationships, in the absence of explicit or implicit information to the contrary. As mentioned before, this is a version of the well-known *frame problem* in AI.

For instance, at a point in the story prior to the excerpt about the wolf entering Grandmother’s cottage, we are told that the wolf knocked on the door. From this we would make the inference that he probably wanted to enter. When the narrator says “the door flew open,” we are quite sure that the wolf still wanted to enter, and so infer that he did so. (Note that the story does not say this!)

Dealing with the frame problem has generally been treated in AI as a matter of designing the right logic, rather than as a matter of having the right knowledge. In fact, this has been a major impetus behind the development of nonmonotonic logics. We think, on the contrary, that persistence is not a matter of logic, but something to be inferred from the right sorts of knowledge about the world, about one’s own knowledge, and about narration. One of us has made specific proposals about the relevant world knowledge in a non-narrative setting (Schubert, 1990; Schubert, 1994), but this is not the place to pursue the matter. We want to suggest, however, that in a narrative setting persistence inferences hinge on certain narrative *conventions*, to the effect that the narrator is expected to let the reader know about changes in relevant properties and relationships, either explicitly or by providing a basis for inferring the changes. For instance, coming back to the wolf’s evident intention to enter Grandmother’s cottage, if the wolf had dropped that intention after knocking (e.g., because no-one answered), the story would have given us a basis for inferring this. Since it did not, we infer that the wolf’s intention persists.

We have formulated this idea as a narrative axiom. For a formal statement of the axiom, we refer the reader to (Hwang and Schubert, 1993a); here we just state the idea in words:

A state evoked by a narrative and characterized by an atelic sentence extends temporally to any given later episode unless I (the reader of the story) have been led to disbelieve this.

An “evoked episode”, in the context of a narrative, is one for which a token has been generated as a result of input-driven inference chaining. In particular, all episodes occurring in orienting relations are considered evoked episodes. Like **orients**, the property of being “evoked” is context-charged, i.e., it is a property that

a thing has in virtue of its role in the narrative, rather than its role in the “world.” The question of whether the reader has been led to believe or disbelieve a particular claim is settled simply by introspection: if at a certain point in the story the reader can verify with minimal inference effort (and with reasonably high probability) ($\text{That } \Phi$), where Φ concerns entities or events in the story, then that’s what the reader has been led to believe.²³ Disbelief is just belief in the negation.

We have not implemented belief introspection as a means of evaluating formulas that assert beliefs of the system, and so we have not used the above axiom. It plays no direct role in the excerpts we have discussed because of their brevity. However, note that we made a number of assumptions in the reasoning process for the final excerpt that would require persistence inferences from *earlier* parts of the story in a more complete system. They include the supposition that the wolf is wicked, that he wants to enter Grandmother’s cottage, that Grandmother is in the cottage and is ill and alone, that drawing out the peg on the door will make the latch fall, that LRRH is on her way to the cottage, and that the wolf knows all this.

Conclusions and Future Work

The main concern of our research has been to develop an adequate semantic representation (SR) and knowledge representation (KR) for general NLU systems, especially ones aimed at understanding narratives.

As a logic, EL is in a state analogous to that of various nonmonotonic, conditional and probabilistic logics whose semantics remains an active area of research. In other words, the syntax is better worked out than the semantics. We have, however, specified the semantic types of all EL expressions, adding various constraints on the structure of situations and other classes of domain entities and on the interpretations of the basic operators, predicates and functions of EL (Hwang, 1992; Hwang and Schubert, 1993b). We also have (unpublished) proofs of the validity of various axiom schemas, the soundness of RI and GC, and the persistence of (certain kinds of) information through the situation ordering. However, not only are revisions needed in all of this because of our revisions to the semantics of ‘**’ and ‘*’, but we still need to demonstrate that there exist (nontrivial) models of the sort we have posited.

There are two extreme views about the deployment of knowledge representations whose semantics is not fully understood. One extreme is to reject such deployment altogether, on the grounds that there will be no guarantee that all reachable conclusions are justifiable and all justifiable conclusions are reachable, and this may be disastrous in certain critical applications. The other extreme is to reject formal denotational semantics altogether, on the grounds that the design of representations and inference mechanisms (especially in language understanding) is a cognitive modelling task in which the correspondence between symbols and the world has no role to play.

We think both extreme views are harmful and account in part for the rift that remains within “knowledge representation and reasoning” as practiced outside and within NLP. We firmly believe that the most effective strategy for arriving at a representation adequate for general NLU and commonsense reasoning is one of progressive refinement and theoretical deepening, starting with a representation that allows us to express linguistic meaning and commonsense knowledge easily and directly, fits with a conceptually modular view of the language understanding process (in the spirit of “divide and conquer”), and readily supports all the sorts of inferences that people are apt to make. The refinement and deepening should be driven by concurrent theoretical investigations and experimental implementations.

The EL representation meets many of our desiderata. It does so by combining ideas from Montague grammar, situation semantics and DRT, and adding a number of new ideas concerning the semantics of situations, actions, propositions and facts, times, quantification and tense and aspect. The deindexing of indexical logical forms with the aid of tense trees is a crucial stage in our approach, systematically deriving episodic logical formulas that can then be used for inference. EL has been implemented and tested on small but realistic text samples, including ones from the particularly challenging genre of fairy tales, and also has been incorporated into complete prototype NLP systems. The results so far are encouraging, suggesting that it is indeed possible to grapple simultaneously with a wide spectrum of problems in natural language

²³In contrast with the perspective taken in *autoepistemic logic* (Moore, 1985), we regard it as crucial that an agent should need only a very small computational effort to check whether it believes something or not (as in (Kaplan and Schubert, 1997; Kaplan, 1998)).

understanding. More specifically, the following are the contributions of EL to knowledge representation for natural language processing.

- EL is an *expressive* SR/KR that allows the content of most English sentences and most world knowledge to be represented in an intuitively comprehensible, computationally usable 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 (extensional) generic conditionals, as well as restricted quantifiers, modal operators, predicate and sentence nominalization operators, and predicate and sentence modifiers. These features have been brought together for the first time in a logic for narrative understanding. Also, the mapping from English to EL is transparent and modular, handling many combinations of tense, aspect and adverbials (this was not discussed in detail here for space reasons, but see (Hwang, 1992; Hwang and Schubert, 1994)).
- In theory and in practice, EL allows linguistic and domain knowledge to be strictly separated from parsing and inference control structure, so that the former can be expanded and revised independently of the latter.
- The main rules of inference in EL, RI and GC, provide input-driven and goal-driven inference modes, and are able to combine multiple premises in one fell swoop. Since the rules allow the use of probabilistic conditionals, they support expert system-like combination of evidence in the derivation of probable explanations or predictions. Furthermore, these rules have been successfully implemented in EPILOG, and integrated uniformly with an array of specialized inference techniques.

Although we think that EL and the associated interpretive and inference mechanisms provide a good “first draft” of a comprehensive foundation for language understanding, many formidable challenges remain. The following are among the most important issues for future work.

- The representation and semantics of several linguistic phenomena remains unclear. This includes quantifier modifiers (“almost all”), comparatives, questions and Wh-nominals (“why she left”), clausal adverbials (“even though she left”), various uses of “except” and “only”, etc. A particularly important area for further investigation is the semantics of *generic passages* (Carlson and Spejowski, 1997) – extended passages that describe typical or habitual patterns of events or relationships. We believe that much of our world knowledge consists, in effect, of such generic passages. (*Scripts* and *frames* can be viewed as representations of such passages (Schank and Abelson, 1977; Minsky, 1975; Schubert, to appear).
- As discussed above, we need to extend and deepen the semantics of EL itself, fully justifying various basic axioms and the RI and GC rules, and demonstrating the existence of models.
- The probabilistic constructs and inferences need to be put on a firmer theoretical foundations. Ideally, we would like to develop an analogue of Bayes net reasoning within a logical setting, resting on analogous conditional independence assumptions (i.e., something like independence of a conclusion of all KB hypotheses (other than ones derived from the conclusion), given the truth or falsity of just those hypotheses that support a one-step probabilistic inference to the conclusion or its negation).
- Inference control needs to be further improved. In particular forward inference termination through “interestingness” and probability criteria is not always intuitively satisfactory.
- We need to gain a better understanding of both *simulative* and *introspective* inference in EL, and to implement these important modes of reasoning.
- Closely related to this is the need for further work on the *frame problem* in the context of narratives; as we indicated, we think this depends on introspective (autoepistemic) reasoning.
- We need to develop in detail a computational theory of “implicit question answering” as outlined earlier, as a means of arriving at a coherent interpretation of extended texts.
- Some of the EPILOG specialists need further development, and some new ones are needed, most importantly a specialist for “imagining” objects interacting and in motion (including complex deformable objects like people, plants, clothing, newspapers, etc.). For instance, when we’re told “John tossed the umbrella in the trunk of his car”, how do we infer that the umbrella was probably folded up?

- A major long-range goal in the development of EPILOG is to equip it with a goal-directed component, so that it will be capable of exhibiting independent initiative, e.g., by asking questions. Of course, EPILOG (like many other inference engines) already does goal-directed inference, but these inferences are aimed at answering questions, not guiding action. We are thinking in terms of a set of overarching “drives”, such as cooperativeness and curiosity, guiding a behavioral component consisting primarily of a conversation and reasoning manager.
- We are eager to link up EPILOG with a broad-coverage, preference-seeking parser for English. In the shorter term, we are planning to experiment with available parsers, and in the longer term we expect to employ an ambitious parser of our own, one particularly well-suited to our interpretive goals, that has been under sporadic development for many years (based on ideas in (Schubert, 1984; Schubert, 1986)).
- Last and perhaps most importantly, we need ways of breaking through the “knowledge bottleneck”, both in linguistic and world knowledge.

One approach to this problem is *knowledge bootstrapping* through linguistic input. This presupposes a certain minimal linguistic competence and basic world knowledge to begin with. The basic grammatical knowledge can be hand-coded (automated grammar acquisition so far does not lead to grammars capable of supporting semantic analysis). However, this leaves the task of disambiguation, which is knowledge-intensive and a major obstacle to getting off the ground. Some work we are currently undertaking is aimed at the accumulation of “head patterns” from linguistic corpora, for use in disambiguation. In essence, head patterns are patterns of co-occurrence of predicates and other operators (such as modifiers) with particular arguments or types of arguments. (As such they are related to the knowledge indexing scheme we use.)

To help accumulate some basic world knowledge, we intend to avail ourselves of (and extend) work that has been done in extracting semantic information from lexicons, especially type hierarchy information and part-of information. We also think that text corpora, particularly fictional ones (because of the everyday *minutiae* they touch on), are a potentially rich source of knowledge about the properties various kinds of entities are likely to have, the actions and relationships they participate in, and ultimately the causal and habitual patterns of events they tend to be involved in.

With a reasonably comprehensive grammar, pattern-based disambiguation, and some rudimentary world knowledge in place, and presupposing the solution of some of the more important remaining representational problems (particularly the representation of comparatives and generic passages), we would be ready to provide further knowledge by telling it to the system in English. This would be an effective way to continue the bootstrapping process, provided that the system had some capacity to detect its own knowledge gaps, and could muster the initiative to ask.

Acknowledgements

The authors are grateful to Phil Harrison for his insights and suggestions during much of the development of EL and to James Allen for many useful comments in the context of the TRAINS project. We owe a large debt of gratitude to Stephanie Schaeffer, who implemented EPILOG, happily accepting and actualizing ever more ambitious specifications and giving us continuous and crucial feedback on how our ideas were working out (or not working out) in practice. This research was supported in part by NSF Research Grants IRI-9013160 and IRI-9503312, ONR/DARPA Research Contract No. N00014-92-J-1512, the Boeing Co. in Seattle under Purchase Contracts W-278258 and W288104, and a Killam Memorial Scholarship (CHH).

References

- Ali, S. S. and Shapiro, S. C. (1993). Natural language processing using a propositional semantic network with structured variables. *Minds and Machines*, 3(4):421–451.
- Allen, J. F. and Schubert, L. K. (1993). Language and discourse in the TRAINS project. In Ortony, A., Slack, J., and Stock, O., editors, *Communication from an Artificial Intelligence Perspective*, pages 91–120. Theoretical Springer-Verlag, Heidelberg.

- Alshawi, H. and van Eijck, J. (1989). Logical forms in the Core Language Engine. In *Proc., 27th Annual Meeting of the Assoc. for Computational Linguistics (ACL-89)*, pages 25–32, Vancouver, B.C., Canada.
- Bacchus, F., Grove, A. J., Halpern, J. Y., and Koller, D. (1996). From statistical knowledge bases to degrees of belief. *Artificial Intelligence*, 87:75–143.
- Ballim, A. and Wilks, Y. (1991). *Artificial Believers: The Ascription of Belief*. Lawrence Erlbaum, Hillsdale, NJ.
- Barwise, J. (1989). *The Situation in Logic*. CSLI, Stanford, CA.
- Barwise, J. and Perry, J. (1983). *Situations and Attitudes*. MIT Press, Bradford Books, Cambridge, MA.
- Carlson, G. N. (1982). Generic terms and generic sentences. *J. of Philosophical Logic*, 11:145–181.
- Carlson, G. N. and Pelletier, F. J. (1995). *The Generic Book*. University of Chicago Press, Chicago.
- Carlson, G. N. and Spejewski, B. (1997). Generic passages. *Natural Language Semantics*, 5:1–65.
- Chalupsky, H. and Shapiro, S. C. (1996). Reasoning about incomplete agents. In *Proc., 5th Int'l. Conf. on User Modelling*.
- Charniak, E. (1988). Motivation analysis, abductive unification, and nonmonotonic equality. *Artificial Intelligence*, 34:275–295.
- Charniak, E. and Goldman, R. (1988). A logic for semantic interpretation. In *Proc., 26th Annual Meeting of the ACL*, pages 87–94, Buffalo, NY.
- Chierchia, G. (1985). Formal semantics and the grammar of predication. *Linguistic Inquiry*, 16(3):417–443.
- Chierchia, G. (1993). Questions with quantifiers. *Natural Language Semantics*, 1:181–234.
- Cohen, A. (1997). Default reasoning and generics. *Computational Intelligence*, 13(4):506–533.
- Creary, L. G. (1979). Propositional attitudes: Fregean representation and simulative reasoning. In *Proc., 6th Int'l. Joint Conf. on Artificial Intelligence*, pages 176–181, Tokyo.
- Cullingford, R. (1981). Sam. In Schank, R. C. and Riesbeck, C. K., editors, *Inside Computer Understanding*, pages 75–119. Lawrence Erlbaum, Hillsdale, NJ.
- Dahlgren, K., McDowell, J., and Stabler, Jr., E. P. (1989). Knowledge representation for commonsense reasoning with text. *Computational Linguistics*, 15(3):149–170.
- Davidson, D. (1967). The logical form of action sentences. In Donald Davidson and Gilbert Harman, editors, *The Logic of Grammar*, pp.235–245. Dickenson Publ., Encino, CA., 1975. (Reprinted from *The Logic of Decision and Action*, Nicholas Rescher, ed., U. of Pittsburg Pr., 1967.).
- Devlin, K. (1991). *Logic and Information*. Cambridge University Press, Cambridge, UK.
- Dowty, D. R. (1982). Tenses, time adverbs and compositional semantic theory. *Linguistics and Philosophy*, 5:23–55.
- Fodor, J. D. and Sag, I. A. (1982). Referential and quantificational indefinites. *Linguistics and Philosophy*, 5:355–398.
- Geach, P. T. (1962). *Reference and Generality*. Cornell University Press, Ithaca, NY. (amended ed. 1968).
- Gerevini, A. and Schubert, L. K. (1995). The temporal reasoning tools timegraph i-ii. *Int. J. of Artificial Intelligence Tools*, 4:281–299.
- Groenendijk, J. and Stokhof, M. (1991). Dynamic predicate logic. *Linguistics and Philosophy*, 14:39–100.

- Haas, A. R. (1986). A syntactic theory of belief and action. *Artificial Intelligence*, 28:245–292.
- Hirst, G. (1988). Semantic interpretations and ambiguity. *Artificial Intelligence*, 34:131–177.
- Hirst, G. (1991). Existence assumptions in knowledge representation. *Artificial Intelligence*, 49:199–242.
- Hobbs, J. R. (1985). Ontological promiscuity. In *Proc., 23rd Annual Meeting of the ACL*, pages 61–69, Chicago, IL.
- 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, J. R., Stickel, M. E., Appelt, D. E., and Martin, P. (1993). Interpretation as abduction. *Artificial Intelligence*, 63:69–142.
- Hwang, C. H. (1992). *A Logical Approach to Narrative Understanding*. PhD thesis, University of Alberta.
- Hwang, C. H. and Schubert, L. K. (1992). Tense trees as the 'fine structure' of discourse. In *Proc., 30th Annual Meeting of the ACL*, pages 232–240, Newark, DE.
- Hwang, C. H. and Schubert, L. K. (1993a). Episodic Logic: A comprehensive, natural representation for language understanding. *Minds and Machines*, 3(4):381–419.
- Hwang, C. H. and Schubert, L. K. (1993b). Episodic Logic: A situational logic for natural language processing. In Peter Aczel, David Israel, Y. K. and Peters, S., editors, *Situation Theory and its Applications*, volume 3, pages 303–338, Stanford, CA. CSLI.
- Hwang, C. H. and Schubert, L. K. (1993c). Interpreting temporal adverbials. In *Proc., ARPA Workshop on Human Language Technology*, pages 138–143, Princeton, NJ.
- Hwang, C. H. and Schubert, L. K. (1994). Interpreting tense, aspect and time adverbials: A compositional, unified approach. In Gabbay, D. M. and Ohlbach, H. J., editors, *Proc., 1st Int'l. Conf. on Temporal Logic*, pages 238–264, Bonn, Germany. Springer-Verlag.
- Iwańska, L. (1993). Logical reasoning in natural language: It is all about knowledge. *Minds and Machines*, 3(4):475–510.
- Iwańska, L. (1997). Reasoning with Intensional Negative Adjectivals: Semantics, Pragmatics and Context. *Computational Intelligence*, 13(3):348–390.
- Jenkins, T., Gaillard, A., Holmback, H., Namioka, A., Darvish, J., Harrison, P., and Lorbeski, M. (1990). Automated message understanding: A real-world prototype. In *Proc., IEA/AIE-90*, Charleston, SC.
- Kadmon, N. (1987). *On Unique and Non-Unique Reference and Asymmetric Quantification*. PhD thesis, University of Massachusetts, Amherst.
- Kamp, H. (1981). A theory of truth and semantic representation. In Groenendijk, J., Janssen, T., and Stokhof, M., editors, *Formal Methods in the Study of Language*, pages 277–320. Mathematical Centre-tracts, U. Amsterdam, Amsterdam, The Netherlands.
- Kaplan, A. N. (1998). Simulative inference about nonmonotonic reasoners. In *Proc., 7th Conf. on Theoretical Aspects of Rationality and Knowledge (TARK VII)*, Evanston, IL.
- Kaplan, A. N. and Schubert, L. K. (1997). Simulative inference in a computational model of belief. In Bunt, H., Kievit, L., Muskens, R., and Verlinden, M., editors, *IWCS II: 2nd Int'l. Workshop on Computational Semantics*. Dept. of Computational Linguistics, Tilburg University, The Netherlands. Extended and revised version in Bunt, H. and Muskens, R., *Computational Semantics* (Studies in Linguistics and Philosophy Series), Kluwer, to appear.
- Kolodner, J. L. (1981). Organization and retrieval in a conceptual memory for events or CON54, where are you? In *Proc., 7th Int'l. Joint Conf. on Artificial Intelligence*, pages 227–233, Vancouver, BC.

- Konolige, K. (1986). *A Deduction Model of Belief*. Morgan Kaufmann, Los Altos, CA.
- Lascarides, A. and Asher, N. (1991). Discourse relations and defeasible knowledge. In *Proc., 29th Ann. Meet. of the Assoc. for Computational Linguistics (ACL'91)*, pages 55–62, Berkeley, CA.
- Lascarides, A., Asher, N., and Oberlander, J. (1992). Inferring discourse relations in context. In *Proc., 30th Annual Meeting of the Assoc. for Computational Linguistics (ACL'92)*, pages 1–8, Newark, DE.
- Leech, G. (1987). *Meaning and the English Verb (2nd ed.)*. Longman, London, UK.
- 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, *Meaning, Use, and Interpretation of Language*, pages 302–323. Walter de Gruyter, Berlin.
- McAllester, D. A. and Givan, R. (1992). Natural language syntax and first-order inference. *Artificial Intelligence*, 56(1):1–20.
- Miller, 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.
- Miller, S. A. and Schubert, L. K. (1988). Using specialists to accelerate general reasoning. In *Proc., 7th Nat'l. Conf. on Artificial Intelligence (AAAI-88)*, pages 161–165, St. Paul, MN.
- Moens, M. and Steedman, M. (1988). Temporal ontology and temporal reference. *Computational Linguistics*, 14(2):15–28.
- Montague, R. (1973). The Proper Treatment of Quantification in Ordinary English. In Hintikka, J., Moravcik, J., and Suppes, P., editors, *Approaches to Natural Language*. D. Reidel Publ. Co.
- Moore, R. C. (1977). Reasoning about knowledge and action. In *Proc., 5th Int'l. Joint Conf. on Artificial Intelligence (IJCAI-77)*, pages 223–227.
- Moore, R. C. (1985). Semantic considerations on nonmonotonic logic. *Artificial Intelligence*, 25(1):75–94.
- Namioka, A., Hwang, C. H., and Schaeffer, S. (1992). Using the inference tool EPILOG for a message processing application. *Int'l. J. of Expert Systems*, 5(1):55–82.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufman, San Mateo, CA.
- Perrault, C. (1961). *Perrault's Complete Fairy Tales*. Dodd, Mead & Co., New York, NY. Translated from French by A. E. Johnson and others.
- Reichenbach, H. (1947). *Elements of Symbolic Logic*. Macmillan, New York, NY.
- Schank, R. C., Goldman, N. M., III, C. J. R., and Riesbeck, C. K. (1975). *Conceptual Information Processing*. North Holland, New York, NY.
- Schank, R. C. and Leake, D. B. (1989). Creativity and learning in a case-based explainer. *Artificial Intelligence*, 40:353–385.
- Schubert, L. K. (1984). On parsing preference. In *Proc., 10th Int. Conf. on Computational Linguistics (COLING-84)*, pages 247–50, Stanford Univ., Stanford, CA.
- Schubert, L. K. (1986). Are there preference trade-offs in attachment decisions? In *Proc., 5th Nat. Conf. on Artificial Intelligence (AAAI-86)*, pages 601–605, Philadelphia, PA.
- Schubert, L. K. (1990). Monotonic solution of the frame problem in the situation calculus: An efficient method for worlds with fully specified actions. In Kyburg, H. E., Loui, R., and Carlson, G. N., editors, *Knowledge Representation and Defeasible Reasoning*, pages 23–67. Kluwer Academic Publ., Dordrecht.

- Schubert, L. K. (1994). Explanation closure, action closure, and the sandewall test suite for reasoning about change. *J. of Logic and Computation*, 4(5):679–799.
- Schubert, L. K. and Hwang, C. H. (1989). An episodic knowledge representation for narrative texts. In *Proc., 1st Int'l. Conf. on Principles of Knowledge Representation and Reasoning (KR'89)*, pages 444–458, Toronto, Canada.
- Schubert, L. K. and Pelletier, F. J. (1982). From English to logic: context free computation of ‘conventional’ logical translations. *Am. J. of Computational Linguistics*, 8:26–44. Also in *Readings in Natural Language Processing*, Barbara Grosz, Karen Sparck Jones and Bonnie Webber, eds., 293–311, Morgan Kaufman, Los Altos, CA, 1986.
- Schubert, L. K. and Pelletier, F. J. (1989). Generically speaking, or, using discourse representation theory to interpret generics. In Chierchia, G., Partee, B. H., and Turner, R., editors, *Property Theory, Type Theory, and Semantics, V.2: Semantic Issues*, pages 193–268. Kluwer Academic Publ., Boston.
- Stickel, M. (1986). Schubert’s steamroller problem: formulations and solutions. *J. of Automated Reasoning*, 2:89–101.
- Traugott, J. (1986). Nested resolution. In *Proc., 8th Int. Conf. on Automated Deduction (CADE-8)*, pages 394–402, Oxford, UK. Springer-Verlag.
- Traum, D., Schubert, L., Poesio, M., Martin, N., Light, M., Hwang, C. H., Heeman, P., Ferguson, G., and Allen, J. (1996). Knowledge representation in the TRAINS-93 conversation system. *Int. J. of Expert Sys., special issue on Knowledge Representation and Inference for Natural Language Processing*, 9(1):173–223.
- Webber, B. L. (1988). Tense as discourse anaphor. *Computational Linguistics*, 14(2):61–73.
- 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.