Framing the Donkey: Towards a Unification of Semantic Representations with Knowledge Representations

Lenhart Schubert
Dept. of Computer Science, University of Rochester
Rochester, NY 14627-0226, USA. Tel.: 716-275-8845
Email: schubert@cs.rochester.edu

Abstract

This paper proposes a new way of handling difficult cases of anaphora in language at the level of logical form. This includes donkey anaphora and references into script-like or frame-like knowledge. These cases of anaphora are important, since they occur in virtually every sentence expressing commonsense, general facts about the world – the kinds of facts a language understanding or commonsense reasoning system needs to know. The proposed approach, called "dynamic skolemization", can be implemented within a more-or-less standard logical framework, avoiding the need for logics in which the context of semantic evaluation changes dynamically, such as Discourse Representation Theory or Dynamic Predicate Logic. The approach leads to script-like or frame-like representations, and thus provides a linguistic motivation for them (at the level of complex, structured knowledge rather than at the level of predicate-argument structure).

Introduction

The phenomenon of "donkey anaphora" has been much discussed in the linguistic semantics literature (e.g., see (Carlson & Pelletier 1995)), but less so in AI, despite – as I will argue – its importance for knowledge representation and acquisition. The phenomenon is exemplified by sentences such as the following:

1. If *a farmer* buys *a donkey*, *he* pays for *it* in cash.
2. If John bought *a donkey* at the fair, he paid for *it* in cash.

Note that (1) has a generic flavor, in that it quantifies over cases where a farmer buys a donkey. (2) admits both a generic and a specific reading. On the generic reading, it quantifies over past instances of John buying a donkey (perhaps with different instantiations of "the fair", or perhaps with a single instantiation of a very long fair). On the specific reading, (2) seems to refer to just a single visit by John to the fair, and a possible donkey-purchase during that visit.

Both sentences and both readings of (2) illustrate the problem of donkey anaphora, but the specific reading of (2) does so most simply and starkly. If we interpret indefinites in terms of existential quantification (as has been customary in philosophical logic and much of AI), then the logical form of the specific reading of (2) appears to be

$$[(\exists x)B(x) \wedge B(J,x)] \;\rightarrow\; P(it,M),$$

ignoring implicit events. But how are we to disambiguate the anaphoric pronoun? We cannot set it to x, since this would lead to a free occurrence of x outside the scope of its quantifier, where this free occurrence is semantically disconnected from the prior bound occurrence. We also cannot appeal to a "wide-scope reading" of the indefinite, giving the existential quantifier wider scope than the conditional and thus bringing the pronoun within its scope. This would give a reading to the effect that something x is such that if it is a donkey and John bought it at the fair, then he paid for it in cash. But such a statement is true virtually independently of John's donkey purchases and manner of payment, as long as there is *some*thing that is not a donkey, or *some*thing that John didn't buy, thus rendering the antecedent of the conditional false and the conditional true.

Finally, we cannot appeal to some principle of quantifier scope widening analogous to conversion to prenex form, since this would convert the negatively embedded existential quantifier to a wide-scope universal. If it happens that John bought several donkeys at the fair, truth of (2) would then require John to have paid for *all* of them in cash. But this is wrong (though it is sometimes an implicature), as can be seen by augmenting (2) as follows:

3. John took enough money to the fair to buy a donkey. So, if he bought a donkey at the fair, he paid for it in cash.

Clearly we do not want to infer that if John was so impressed with the livestock offerings that he went on a donkey-buying spree, he paid for *all* his donkeys in cash. Yet (2) remains true in this context. The point is even clearer for the following variant of (2):

4. If John bought a donkey at yesterday's fair, he rode it home.

This seems perfectly true even in a case where John bought two donkeys, but rode only one home while leading the other. Another argument against the scope-widening solution is that it plainly leads to incoherence for variants of (1) involving explicit generic quantification, such as

5. If a farmer buys a donkey, he {*usually/ sometimes/ rarely*} pays for it in cash.

Why, then, should the KR community care about this problem? I claim that much of the knowledge needed for commonsense reasoning is of this form. For instance, the following are some general facts that everyone presumably knows:

6. If a hungry predator encounters a smaller creature, it may well try to attack and eat it.
7. If a conspicuous action occurs within plain sight of a person (in a normal state of alertness), the person will notice it.
8. If an adult attacks a child, that is an extremely wicked action.
9. If someone does something extremely wicked, and others find out about it, they will want to punish him severely.
10. If one enters an enclosure, one will then be inside it (and outside it immediately before).
11. If you eat at a restaurant, you enter (it), get a table to sit at, select a meal from a menu, place your order with a waiter, etc.

The first five items are taken loosely from the encodings in (Schubert & Hwang 1989; Hwang & Schubert 1993a) of the knowledge needed to understand some small fragments of a fairy tale (Little Red Riding Hood). We found that dozens of axioms of this sort, encoded in EL (Episodic Logic, an NL-like representation) were typically required to enable the EPILOG system to make sense of a sentence in a story (also represented in EL), i.e., to infer the "obvious" causal, temporal, part-of, and other coherence relations. Of course, the claim that story understanding is a knowledge-intensive activity is a truism in AI nowadays (thanks to the work of Roger Schank and his coworkers, and many others in the 70's and 80's), but the present point concerns the particular form of the required facts when stated in ordinary language. Evidently, all five sentences involve at least one instance of donkey anaphora. (In the third sentence, the deictic *that* is in effect a reference to an action whose existence is implicit in the antecedent.)

The last sentence is intended to echo the kind of knowledge often encoded in script-like representations. Setting aside the role of *you*, the consequent clauses clearly involve tacit or explicit references to the restaurant in question, and these references are again of the "donkey" variety.

The need to convey huge numbers of mundane generalities to machines remains a frustrating bottleneck in the effort to endow machines with ordinary understanding and common sense. There are several possible ways of attacking the problem, the most direct being the systematic hand-coding of the requisite facts (Lenat 1995). However, this approach is very labor-intensive, and the task would be greatly lightened if the bulk of the necessary knowledge could be conveyed directly through language.

The problem of donkey anaphora presents a particularly formidable roadblock in this endeavor. There are some proposed solutions that deal plausibly with examples like (1-4), as indicated in the next section. But in the first place, these solutions depend on adopting a logic with a quite radically altered, "dynamic" semantics, and corresponding nonstandard inference mechanisms. The possibility of translating into a more conventional framework improves prospects for bridging the gap between work in linguistic semantics and knowledge representation. More importantly, I will show in the next section that the dynamic semantics approach encounters difficulties that would prevent its broader use for "telling machines about the world".

In section 3, I will then develop the proposal I call *dynamic skolemization*, showing how it deals both with standard donkey anaphora and the more challenging cases described in section 2. Dynamic skolemization can be augmented so as to introduce new concepts (predicates) that "summarize" the content of the formulas comprising the scope of the skolemized variables. When this is done, the representations becomes quite concise, and strongly reminiscent of scripts and frames in AI.

In section 4 I discuss some further issues in the application of dynamic skolemiztion, and possible limitations of the approach. In the final section I reiterate my conclusions.

Earlier Solutions: Dynamic Semantics

The difficulties posed by donkey anaphora have spurred new developments in the logical analysis and representation of ordinary language. Sentence (1), repeated here as (12), will serve to illustrate two well-known approaches very briefly:

12(a) If a farmer buys a donkey, he pays for it in cash

(b)
$$\boxed{\begin{array}{|l|}\hline \quad x\ y \\ \hline F(x) \\ D(y) \\ B(x,y) \\ \hline \end{array} \;\Rightarrow\; \begin{array}{|l|}\hline \quad z\ w \\ \hline P(z,w) \\ z = x \\ w = y \\ \hline \end{array}}$$

(c) $(x, y\ :\ F(x), D(y), B(x,y)) \quad \Rightarrow \quad (z, w\ :\ P(z,w), z = x, w = y)$

13. $[(\exists x, y)F(x) \wedge D(y) \wedge B(x,y)] \ \rightarrow\ P(x,y)$

Perhaps the best-known approach to anaphora is that of (Kamp 1981) and (Heim 1982), known as DRT. The representation of (2a), repeated here as (12a), would be a Discourse Representation Structure (DRS) as shown in (12b). This DRS contains a conditional whose antecedent and consequent are again DRS's. The indefinites *a farmer* and *a donkey* are represented as variables called "discourse referents", comprising the syntactically separate "universe" of the antecedent DRS. Corresponding "conditions" $F(x)$ and $D(y)$ appropriately constrain them in the antecedent DRS. The treatment of the definites in the consequent DRS is analogous, and the anaphoric references are resolved by the equations in the consequent DRS. An alternative, nonpictorial syntax could be as illustrated in (12c).

What is crucial here is the semantics, in particular the "dynamic" way in which discourse referents are added to the domain of a partial assignment function in the course of semantically evaluating a DRS. Roughly speaking, a conditional DRS like (12b) is taken to be true relative to a model and a partial assignment if every way of making the antecedent true by adding its discourse referents (here, x and y) to the domain of the assignment can be extended to a way of making the consequent true by adding its discourse referents (here, z and w) to the domain of the assignment.

Clearly, then, the way DRT treats the anaphoric pronouns in the consequent of a conditional is both syntactically and semantically dependent on the presence of coreferring discourse referents in the antecedent. Syntactically, the translations of the pronouns need to be equated to "accessible" referents (such as those in the antecedent, or in the universe of an embedding DRS), and semantically the equations succeed in making the desired connection because of the way partial assignment functions get extended in the antecedent, and the way this extension is "carried forward" to the consequent.

A disadvantage of DRT as a meaning representation is that it requires a mapping from surface linguistic form to logical form that is hard to reconcile with a compositional semantics. DPL (Groenendijk & Stokhof 1991) provides a compositional alternative. DPL would assign (12) the logical form shown in (13). Note that in DPL we can actually substitute the quantified variables x and y for the anaphoric pronouns, much as we were tempted to do – but couldn't make formal sense of – in the case of *it* in (2). The way DPL makes sense of (13) is by treating formulas much like computer programs, where variables are assigned values in earlier subformulas, and these values persist to later subformulas. In particular, an existential formula $(\exists x)\Phi$ is thought of as denoting a *nondeterministic assignment* to x, followed by a *test* Φ. Truth corresponds to "executability" (for some initial values of the variables and some nondeterministic choices during execution). Technically, this is formalized in terms of state-change semantics, where a state is a (total) variable assignment; i.e., the semantic value of $(\exists x)\Phi$ relative to a model consists of the pairs of states such that the formula can (nondeterministically) transform one into the other. For a conditional $\Phi \rightarrow \Psi$, successful execution (truth) requires that for every way of executing Φ, the resultant state allows execution of Ψ.

Again, the dynamic change in the context of evaluation induced by the prior material is apparent, as is the need for "accessible" prior referents. Another significant point to note is that both approaches lead to a universal interpretation of indefinites in the antecedent of a conditional. In DRT, this is because of the iteration over *all* ways of making the antecedent true (in the present example, all ways of picking a farmer and a donkey such that the farmer buys the donkey). Similarly, in DPL this is because of the iteration over successful executions of the antecedent. So this semantics is somewhat appropriate for (12) but hardly for the specific reading of (2). However, one could define a "specific conditional" within DRT or DPL that does justice to intuitions about (2). EL, mentioned above, would assign much the same logical form as DPL to (12) (neglecting events), but would employ a distinctive "generic conditional" to obtain a quasi-universal reading. For the material conditional, EL's particular brand of DRT-like dynamic semantics would give the desired specific reading for (2).

The reliance of these theories on the presence of accessible referents for interpreting anaphora can be a serious handicap. This is apparent for passages where "dependent entities" are introduced in one sentence, and then referenced in another, as in

14. There were several instances at the fair where a farmer bought a donkey. In each instance, *the farmer* paid for *the donkey* in cash.

The dynamic semantics approaches provide no easy way to logically represent the referential definites in the second sentence so that they refer to the appropriate farmer and donkey in each of the instances under discussion. This is apparent if we write down a DPL-like or EL-like logical form for these sentences (enriching the basic syntax slightly so as to allow restricted quantifiers and some set-theoretic constructs):

15. (a) $(\exists S : \textit{set-of-instances}(S))$
$(\forall e : e \in S)$
$(\exists x : F(x))(\exists y : D(y))B(x, y, e).$
(b) $(\forall e : e \in S)(\exists e' : e' \sqsubseteq e)$
$(\textit{The}\ z : F(z))(\textit{The}\ w : D(w))P(z, w, e').$

This interprets the "instances" in (13) as donkey-purchases by farmers, and accordingly introduces a Davidson-like event variable into B ("buys") and likewise into P ("pays in cash for"). While I do not regard this approach to event representation adequate (see (Hwang & Schubert 1993b)) it will do for present purposes.

Note, first of all, that the referential connection between the bound occurrence of S at the beginning and the free occurrence later on will be made successfully by the dynamic semantics approaches mentioned. However, the same is not true of the definites

(*The* $z : F(z)$) and (*The* $w : D(w)$). Substituting the farmer-variable x from the first sentence for z, and the donkey-variable y for w, simply does not lead to a coherent result in any of the theories, since $(\exists x : F(x))$ and $(\exists y : D(y))$ lie within the scope of a universal quantifier. Intuitively, the first sentence posits various values of x and y depending on the choice of $e \in S$, and in the second sentence z and w should refer to these values as a function of e, but no such function is available.

The problem is that the desired referents are too deeply embedded in the prior discourse, so that they cannot do the required context-setting for interpreting the anaphoric descriptions. In fact, the only way we appear to have of solving the reference problem is to reformulate the second sentence so that it *repeats* the content of the first sentence antecedently, making the required referents directly accessible:

16(a) In each instance, if *a farmer* bought *a donkey, the farmer* paid for *the donkey* in cash.

(b) $(\forall e : e \in S)\ \underline{[(\exists x : F(x))(\exists y : D(y))B(x, y, e)]} \rightarrow (\exists e' : e' \sqsubseteq e)P(x, y, e')$

However, besides being cumbersome and *ad hoc*, this importation of material ought to be redundant. After all, the discourse prior to (16a) already characterized the instances in question as instances in which a farmer buys a donkey, so why should we have re-specify this property?

This is not the worst case. There are familiar kinds of reference where the required antecedent material is not merely submerged in the prior discourse, but hidden in the tacit knowledge of the speaker and hearer. The following is an example:

17(a) John *dined at Mario's*. He left a big tip for *the waiter*.

(b) $(\exists e)D(J, M, e)$. $(\exists x : T(x))(\exists e' : e' \sqsubseteq e)$ $(\mathit{The}\ y : W(y))L(J, x, y, e')$

Any of the dynamic semantics approaches would now require retrieving and instantiating a "dining out" scenario, at least in part, and incorporating this antecedently into the representation of the second sentence. Only then would we be able to resolve the definite description (*The* $y : W(y)$) to the appropriate referent, i.e., the waiter who served John. This approach would be laborious and hard to systematize, keeping in mind that stored scenarios could be quite complex, and that multiple scenarios could be relevant to a single sentence (e.g., "At the business dinner at Mario's, the waiter spilled coffee on the CEO just when he was about to put up some pie-charts"). These are just the sorts of problems that hampered full exploitation of "scripts" in story-understanding research in the 70's and 80's.

My goal here is to provide a straightforward and concise method of interpreting sentences involving anaphora as in (16) and (17), as well as the more innocuous cases of donkey anaphora. The idea is to make available *Skolem functions* (constants, in the simplest cases) for establishing referential connections. For instance, in the case of (17) we want to make available a "waiter function" that picks out the waiter of any dining-out event, as a function of that event. In the following I develop such an approach, beginning with simple cases of (non-donkey) anaphora.

Dynamic Skolemization

Skolem constants

Let's begin with a very simple example of cross-sentential anaphora, the sort that invites the use of Skolem constants – and has been dealt with in that way in countless natural-language systems:

18(a) John bought a donkey. He paid for it in cash.

(b) $(\exists x : D(x))(\exists e)B(J, x, e)$. $(\exists e')P(J, it, e')$.

(c) $D(A) \wedge B(J, A, E)$. $P(J, A, E')$.

(18b) shows the preliminary logical form of (18a) (again with Davidsonian event arguments). In (18c), the variable x for the newly introduced donkey has been skolemized to A and the variable e for the newly introduced buying event has been skolemized to E. The anaphoric pronoun has also been replaced by A, unproblematically making the cross-sentential connection. Note that (as in some previous examples) there is also an implicit relation between the paying event and the buying event, i.e, the paying is understood to be part of the buying. Various theories have been proposed for deriving this sort of connection systematically (e.g., (Webber 1988; Hwang & Schubert 1992; Lascarides *et al.* 1992)). By whatever means the connection is derived, the skolemization of event variables facilitates its representation, e.g., as $E \sqsubseteq E'$.

We should note in passing that skolemization is not a logically sound transformation. Rather, it amounts to a logical strengthening, since predicating constraints on a constant will filter out some possible models. But the stregthening is trivial in the sense that it is guaranteed to preserve satisfiability, and that no new sentences are entailed, other than some that involve the Skolem constant. (See the "conservative extension" theorem below. This is why skolemization is legitimate in refutation proofs, such as resolution proofs.)

If one now considers a conditional sentence like (2) (repeated here as (19a)), one feels intuitively that the anaphoric connection is made just as in (18a), the non-conditional case.

19(a) If John bought a donkey at the fair, he paid for it in cash.

(b) $[(\exists x)D(x) \wedge B(J, x)] \rightarrow P(J, it)$

(c) $[D(A) \wedge B(J, A)] \rightarrow P(J, A)$.

One conceptualizes the situation of John buying a donkey, and in doing so, establishes some conceptual token for the donkey. This token is then available for reference, whether the context-setting sentence supplies a fact or a proviso.

Of course, the discourse referents of DRT and the existential variables of DPL provide just such tokens, but I have argued that they do so only very locally and transiently. Skolem constants and functions could provide a simple, "heavy-duty" alternative, but seem to fail us here. If we assign to (19a) the preliminary logical form (19b) (neglecting events for simplicity) and then attempt to skolemize and resolve the referent as earlier in (18c), the result (19c) is patently defective. This is apparent from the fact that (19c) is the skolemized form of $(\exists x)[D(x) \wedge B(J,x) \rightarrow P(J,x)]$, i.e., "There is a thing x such that if it is a donkey and John bought it, then he paid for it in cash"; but this is trivially true if we can find any x that is not a donkey or was not bought by John, rendering the antecedent false.

However, it is interesting to note that (19c) fails as a representation of (19a) by being too weak, rather than too strong. Given (19a) and given that John bought a donkey, it should follow that there is a donkey that John bought, and paid for in cash. But if we replace (19a) as a premise in this argument by (19c), the conclusion no longer follows. This is because the existence of a donkey that John bought fails to entail $D(A) \wedge B(J, A)$, the antecedent of (19c).

This suggests a simple remedy: we supplement the skolemization in (19c) with the following stipulation:

20. $(\exists x : D(x))B(J,x) \ \rightarrow \ D(A) \wedge B(J,A)$.

I will use the term *dynamic skolemization* for the two steps of (i) skolemizing an existential variable, and (ii) stipulating a supplementary condition like (20) above, relating the unskolemized existential formula to its skolemized form. (This is stated more formally, and generalized, below.) I will also refer to the supplementary condition itself as a "Skolem conditional".

Given (20), falsity of $D(A) \wedge B(J,A)$ implies that John didn't buy a donkey. So $D(A) \wedge B(J,A)$ can no longer be trivially false, and (19c) trivially true, as a result of A not being a donkey or John not having bought it. Some additional points are worth noting. First, the Skolem conditional can be turned into a biconditional, by existential generalization. This makes it natural to think of (20) as a partial *definition* of the newly introduced constant A. As such it provides a *conservative extension* of any theory framed in a vocabulary exclusive of A, as shown later.

Second, when we resolve the pronoun in (19a,b) to the constant A, introduced through dynamic skolemization, the resultant reading is *non-universal.* To obtain a past-habitual reading of (19a), a generic or habitual quantifier (arguably, over events) would have to be introduced into the logical form (19b). Examples of this sort will be seen later.

A third, related point concerns the logical consequences of readings based on dynamic skolemization. Suppose we are told after (19a) that John *did* buy a donkey (within the relevant time frame), namely Archie (A'):

$D(A') \wedge B(J, A')$.

It does *not* follow that John paid for Archie in cash. This is because it does not follow that $A' = A$. If we want these consequences, we could add the following (conditional) uniqueness implicature to our interpretation of (19a):[1]

$[(\exists x : D(x))B(J,x)] \ \rightarrow \ (\exists! x : D(x))B(J,x)$.

An alternative would be to regard the report of John's purchase of Archie as implicating the uniqueness of this donkey-purchase. In fact, both implicatures may well be present.

Abstracting from the above example, we can define dynamic skolemization as follows, for an existential sentence that occurs outside the scope of all other quantifiers:

21. Given an occurrence of a (closed) sentence of form $(\exists x : \Phi)\Psi$ in the provisional logical form of an English sentence being interpreted,
    (a) assert the Skolem conditional
    $(\exists x : \Phi)\Psi \ \rightarrow \ \Phi_{C/x} \wedge \Psi_{C/x}$,
    where C is a new constant;
    (b) replace the original occurrence of $(\exists x : \Phi)\Psi$ by the consequent of the above implication.

The notation $\Phi_{C/x}$ stands for the result of substituting C for all free occurrences of x in Φ.[2] The "provisional logical form" of a sentence refers to the logical form we would obtain *before* determining referents of pronouns and definites, but *after* analyzing phrase structure, applying rules of logical-form composition corresponding to phrase structure rules (grounding these in a lexicon that supplies logical translations of lexical items), and scoping quantifiers. Of course, phrase structure, lexical semantics, and quantifier scope are all sources of ambiguity, so an actual analysis sytem would in general have to entertain multiple provisional logical forms for a given (partial) sentence. The different reference possibilities corresponding to these different logical forms may lead to more or less coherent interpretations, and in this way may well influence the ultimate resolution of the various types of ambiguity.

Let us note some properties of dynamic skolemization, beginning with the following simple

Theorem. *The Skolem conditional (21a) provides a*

[1] I take such an implicature to be present in many contexts, and to be the result of the hearer's assumption that if the speaker thought there might have been multiple instances of the indefinite, and such instances are relevant to the purposes of the exchange, then he would have made this possibility explicit. This is analogous to Gricean "quantity scale implicatures" in sentences such as "*Some / Most* students passed the exam", or "If *some / most* students passed the exam, John will be pleased", where the speaker seems to assume that not all students passed the exam.

[2] I use the particular variable x and constant C here for readability, but these are of course arbitrary.

conservative extension *of any theory framed in a vocabulary exclusive of the new Skolem constant C.*

What this means is that no new sentences follow from the theory together with the definition, other than ones involving the defined term C. The theorem is easily proved by considering any formula Φ not involving C that is entailed by a given theory $\mathcal{T}$ (not involving C) together with the Skolem conditional. For any model $\mathcal{M}$ of theory $\mathcal{T}$, let $\mathcal{M}' = \mathcal{M}$ if the antecedent of the Skolem conditional is false in $\mathcal{M}$; otherwise let $\mathcal{M}'$ differ from $\mathcal{M}$ only in assigning a value to C that makes the consequent of the Skolem conditional true. Then $\mathcal{M}'$ is a model of $\mathcal{T}$ together with the Skolem conditional, and so by assumption satisfies Φ. But $\mathcal{M}$ and $\mathcal{M}'$ agree on Φ, hence $\mathcal{M}$ satisfies Φ. Thus $\mathcal{T}$ entails Φ. I leave the generalization of the theorem (and proof) for the case of Skolem *functions* (below) to the reader.

Another property that (21) ought to satisfy is that it should give the expected result in simple declarative contexts – keeping in mind that it was motivated by existential sentences within conditional antecedents. It is easy to see that it does. For instance, consider (once again) (18a). First, note that the skolemization specified by (21b) yields the sentence $D(A) \wedge B(J, A)$ (*modulo* the choice of arbitrary new constant) as previously specified in (18c). But since this sentence occurs at the top level of the text being interpreted, it is *asserted.* As such it entails the Skolem conditional (previously given as (20)), since it coincides with the consequent. Hence the Skolem conditional is redundant. In other words, in the case of simple declaratives (21) just describes ordinary skolemization of top-level existentials.

Skolem functions, concept definitions

The generalization of dynamic skolemization (21) to existential formulas containing free variables (i.e., embedded within the scope of other quantifiers) is straightforward:

22. Given an occurrence of a formula $(\exists y : \Phi)\Psi$, containing free variables $x_1, ..., x_m$, in the provisional logical form of an English sentence being interpreted,
    (a) assert the Skolem conditional
    $$(\forall x_1)...(\forall x_m)[(\exists y : \Phi)\Psi \rightarrow \Phi_{f(x_1,...,x_m)/y} \wedge \Psi_{f(x_1,...,x_m)/y}],$$
    where f is a new m-place function symbol;
    (b) replace the original occurrence of $(\exists y : \Phi)\Psi$ by the consequent of the above implication.

A simple illustration is provided by the following type of quantified donkey sentence, often used as central example in the literature; (we neglect events):

23(a) Every farmer who bought a donkey paid for it in cash.

(b) $(\forall x : F(x) \wedge (\exists y : D(y))B(x, y))\, P(x, it)$

(c) Assert:
$(\forall x)[(\exists y : D(y))B(x, y) \rightarrow D(f(x)) \wedge B(x, f(x))]$
Subst.:
$(\forall x : F(x) \wedge D(f(x)) \wedge B(x, f(x)))\, P(x, f(x))$

Note the resolution of the anaphoric pronoun in (23c). The reading this encodes is *non-universal* for donkeys, i.e., any one farmer purchasing multiple donkeys need only have paid for one in cash. I believe this is correct, based on examples analogous to (3) and (4), despite the frequently present universal implicature.[3] For example, an analog of (4) is

24. Every farmer who bought a donkey rode it home,

which seems true even if the farmers involved rode only one donkey each.

It is clear that repeated application of (22) to a formula with multiple existential quantifiers, $(\exists y_1 : \Phi_1)...(\exists y_n : \Phi_n)\Psi$ will be equivalent to simultaneous skolemization of the y_i to $f_1, ..., f_n$. I will henceforth freely apply it in this way. The following is an interesting variant of (1), illustrating the case $n = 2$; (25b) is the Skolem conditional, and (25c) the logical form of (25a) after substitution:

25(a) Whenever a farmer bought a donkey, he paid for it in cash.

(b) $(\forall e)[(\exists x : F(x))(\exists y : D(y))B(x, y, e) \rightarrow F(f(e)) \wedge D(g(e)) \wedge B(f(e), g(e), e)]$

(c) $(\forall e)[F(f(e)) \wedge D(g(e)) \wedge B(f(e), g(e), e) \rightarrow (\exists e')\, P(f(e), g(e), e')]$

The *whenever*-clause calls for universal event quantification, and this leads to the introduction of a farmer-function f and donkey-function g on events. In this case, we obtain a reading that effectively quantifies universally over farmers and donkeys, assuming that in distinct action predications such as $B(x, y, e)$, $B(x', y', e')$, e and e' cannot be the same buying-event unless x and y are the same as x' and y' respectively. This seems intuitively correct.

There remains a certain inelegance in these examples of dynamic skolemization from a practical computational perspective, in that we are repeating variants of the original existential formula, $(\exists y_1 : \Phi_1)...(\exists y_n : \Phi_n)\Psi$, three times: twice in the Skolem conditional (in the antecedent and consequent) and once in the logical form of the given text. I will therefore introduce a space-saving notation, called *concept definition* that avoids redundancy.

Before defining the notation formally, the idea is just this: when we encounter a wff $(\exists y_1)...(\exists y_n)\Phi$ that is to be skolemized, we first define a new predicate that is simultaneously equivalent to (i) this wff, for all values of the free variables (those not bound by the $(\exists y_i)$), and (ii) the skolemized version of the wff. We then

[3] I believe that a reasonable account of the universal implicature is possible, in terms of interpretive strategies in the face of violated uniqueness implicatures.

substitute the newly defined predicate into the source text. Formally, the concept definition schema is this:

> 26. For π a new m-place predicate constant, $f_1, ..., f_n$ new m-place function constants, and Φ a formula possibly containing free occurrences of variables $x_1, ..., x_m$ and $y_1, ..., y_n$ (and no others),
> $(\mathbf{Def}\ \pi\ \ (x_1, ..., x_m)\ (f_1, ..., f_n)\ \Phi_{\underline{f}/\underline{y}})$
> abbreviates (with underlining indicating vectors)
> $(\forall x_1)...(\forall x_m)[\pi(x_1, ..., x_m) \leftrightarrow (\exists y_1)...(\exists y_n)\Phi \leftrightarrow \Phi_{\underline{f}(x_1,...,x_m)/\underline{y}}]$

$x_1, ..., x_m$ are the *variables* of the Def-schema, $f_1, ..., f_n$ are the *roles* (role functions, Skolem functions), and $\Phi_{\underline{f}/\underline{y}}$ is the *body* of the definition. Note that in this body we are using function symbols (without arguments) in place of variables, so that the result will not be a well-formed formula. This avoids having to repeat the variables $x_1, ..., x_m$ on which each role depends.

Applying the Def-schema to (25) as described above, we get a more concise and transparent representation of the dynamic skolemization (now formulated as a concept definition, (27a)) and the resultant logical form, (27b); *FBD* is the newly defined concept (mnemonic for "farmer buys donkey"):

27(a) $(\mathbf{Def}\ FBD\ (e)\ (f,g)\ F(f) \wedge D(g) \wedge B(f,g,e))$
(b) $(\forall e)\ FBD\ (e) \rightarrow (\exists e')P(f(e), g(e), e')$

We are now well-equipped to return to the farmers at the fair, and Mario's restaurant. Looking back at (15a), note that we will first skolemize S, the set of instances mentioned, to some new constant (but I will retain S here). The next existential sentence encountered is $(\exists x : F(x))(\exists y : D(y))B(x,y,e)$. Hence we get the definition

$(\mathbf{Def}\ FBD(e)\ (f,g)\ F(f) \wedge D(g) \wedge B(f,g,e))$.

Substituting the defined concept in the provisional logical form (15a), we obtain (28a); we can now easily resolve the references in (15b) to $f(e)$ and $g(e)$, as shown in (28b):

28(a) $(\forall e : e \in S)FBD(e)$.
(b) $(\forall e : e \in S)(\exists e' : e' \sqsubseteq e)P(f(e), g(e), e')$.

To deal with the example of Mario's restaurant, (17), we naturally need to presuppose some prior knowledge of "what happens when" a person dines out. While people are unlikely to first acquire this information in the form of a tidy verbal package, it would help with knowledge bootstrapping if this were possible for computers. So suppose we tell our NLP system something like the following:

> Generally, when a person dines at a restaurant, s/he enters the restaurant, gets seated at a table, selects a meal from a menu, tells the order to a waiter, etc.

I will assume that *generally* quantifies over dining-at-restaurant episodes here, binding an episodic variable:

$$\begin{array}{l}(Gen\ e : (\exists x : P(x))(\exists y : R(y))D(x,y,e)) \\ \quad (\exists e_1 : e_1 \sqsubset e \wedge starts(e_1, e))\ Enter(\hat{x}, \hat{y}, e_1) \wedge \\ \quad (\exists e_2 : e_2 \sqsubset e \wedge \hat{e} \prec e_2)\ Get\text{-}seated(\hat{x}, e_2) \wedge \\ \quad (\exists e_3 : e_3 \sqsubset e \wedge \hat{e}_2 \prec e_3)\ Select(\hat{x}, e_3) \wedge \\ \quad (\exists w : W(w))(\exists e_4 : e_4 \sqsubset e \wedge \hat{e}_3 \prec e_4) \\ \qquad\quad Tell\text{-}order(\hat{x}, w, e_4) \wedge \\ \qquad\quad \ldots \qquad\qquad)\end{array}$$

Note that the first line of this expression comprises the *restrictor* of the *Gen* quantifier. The symbols with carets are an *ad hoc* abbreviation for anaphoric terms, used for readability. For instance, $Enter(\hat{x}, \hat{y}, e_1)$ abbreviates

$(The\ z : R(z))Enter(s/he, z, e_1)$,

i.e., s/he enters the restaurant; the particular symbols used, such as $\hat{x}$ and $\hat{y}$, are intended to suggest to the reader the expected referents of these terms.

The first existential sentence encountered is the restrictor of the *Gen* quantifier. Thus we apply dynamic skolemization and concept definition to it:

$(\mathbf{Def}\ \ PDR\ \ (e)\ (f,g)\ P(f) \wedge R(g) \wedge D(f,g,e))$.

Having substituted the defined predicate in the restrictor, and knowing that x is now $f(e)$ and y is $g(e)$, we can also resolve $\hat{x}$ and $\hat{y}$ to $f(e)$ and $g(e)$ respectively in the matrix of the *Gen* construct. Next we apply dynamic skolemization and concept definition to the matrix. Here I allow myself a slight liberty, skolemizing all of the conjoined existential sentences at once (processing them in succession would yield multiple definitions rather than a single one, but would otherwise be equivalent):

$$\begin{array}{l}(\mathbf{Def}\ \ PDRsteps\ \ (e)\ (f_1, f_2, f_3, h, f_4, ...) \\ \quad f_1 \sqsubset e \wedge starts(f_1, e) \wedge Enter(f, g, f_1) \wedge \\ \quad \ldots \\ \quad W(h) \wedge f_4 \sqsubset e \wedge f_3 \prec f_4 \wedge Tell\text{-}order(f, h, f_4) \wedge \\ \quad \ldots \qquad\qquad)\end{array}$$

Substituting in the matrix of *Gen*, we then obtain

$(Gen\ e : PDR(e))PDRsteps(e)$.

Having accommodated the background knowledge in this way, we are now ready to interpret the definite description *the waiter* in (17), repeated here as (29a). The result is (29b), and after further skolemization, (29c):

29(a) John dined at Mario's. He left a big tip for *the waiter*.
(b) $D(J, M, E)$.
$(\exists x : T(x))(\exists e : e \sqsubset E)L(J, x, h(E), e)$.
(c) $D(J, M, E)$. $T(C) \wedge E' \sqsubset E \wedge L(J, C, h(E), E')$.

There is still a hidden ambiguity, however. As far as the stored background knowledge is concerned, the agent of E is $f(E)$ and the restaurant is $g(E)$. To link these to *John* and *Mario's* respectively, we again need to invoke a "uniqueness of roles" assumption concerning dining-out events:

$(\forall e)(\forall x, y, x', y')D(x, y, e) \wedge D(x', y', e) \rightarrow$

$x = x' \wedge y = y'$.

Again, this is an intuitively reasonable assumption.

Scripts, Frames and Knowledge Bootstrapping

The reader has probably noticed a close resemblance between the above encodings of the farmers-at-the-fair and dining-at-Mario's examples on the one hand, and scripts and frames as understood in AI (e.g., (Schank & Abelson 1977; Minsky 1975)) on the other. For instance, Minsky suggests that a person entering an office retrieves and instantiates an office frame featuring *slots* for the expected parts and furnishings (floor, ceiling, walls, door, desk, chair, etc.). These slots quickly become bound to perceived parts and furnishings, and the knowledge associated with them helps to guide further perception and action. Now this is just the sort of structured representation we would obtain by skolemizing a description of the form, "An office generally has a floor, a ceiling, four walls, a door, a desk, etc., with such-and-such properties and related in such-and-such ways...". Assuming a provisional logical form like

$(Gen\ x{:}\,Office(x))\,[(\exists y_1 : Floor(y_1)\ \ldots\,]$,

we would define of a concept such as $Office\text{-}interior(x)$ corresponding to the body of the *Gen*-formula, with roles (Skolem functions) that pick out the parts and furnishings of the office. It would then be straightforward to interpret the definites in general and particular sentences like

30. An office belonging to a professor often has piles of books and papers on *the desk*.

31. Mary went into her office and sat down at *the desk*.

Note that the *Gen*-formula for an office does not involve any existential quantifiers in the restrictor, and so does not give rise to a concept definition for the restrictor (in contrast with the examples of the farmers at the fair and Mario's restaurant). Thus the formula supplies information about the pre-existing concept of an office, rather than about newly introduced episodic notion such as that of dining at a restaurant. This is typically how frames have been employed in knowledge-based systems, i.e., they are used to supply information about familiar, usually non-episodic concepts such as that of an office, a wage earner, a particular product, etc., expressible with a lexical or compound noun.

In the case of general episodic knowledge, the tie-in with scripts is by now obvious. The dining-out example was of course intended as a parallel to Schank and Abelson's (1977) restaurant script. While their script is not formally interpreted, the name of the script can be viewed as a defined concept similar to *PDR* above; further, it involves roles for the participating entities that are clearly analogous to some of the skolemized roles in the definitions of *PDR* and *PDRsteps*, and it involves steps that are analogous to the skolemized subepisodes in *PDRsteps*. An interesting difference is that it is not possible to access the roles in a script "from the outside", as seems to be required for representing general or particular sentences such as

31. When a person dines at a *fancy* restaurant, s/he pays *the waiter* directly, not at a cash register.

32. When Mary dined at Mario's, *the waiter* turned out to be an an ex-classmate of hers.

For (31), a traditional script-based approach would have to create a variant script for fancy dining, as a context in which to place the new fact (assuming that the pre-existing restaurant script is noncommittal about how the bill is paid). This could be done in some *ad hoc* way by modifying a copy of the regular script, or directly inserting an alternative track into that script. In the approach based on dynamic skolemization, we would automatically create a separate concept of dining at a fancy restaurant, and, recognizing that dining at a fancy restaurant entails dining at a restaurant, make use of the role functions in the *PDR* and *PDRsteps* definitions to pick out the waiter (and the implicitly referenced bill). For (32), the traditional approach would expand out an instance of the restaurant script, thus gaining access to the waiter in that script instance. In the approach based on dynamic skolemization, it is sufficient to recognize Mary's dining at Mario's as an instance of a person dining at a restaurant (hence instantiating *PDR*); this immediately makes available the role functions in *PDRsteps* (as was illustrated for 29a).

It has been known for a long time that frame slots involve Skolem-like functions, logically speaking (Hayes 1979). Furthermore, role-functions have been proposed in the planning literature for picking out steps in goal-oriented courses of action (e.g., Kautz 1991). However, the representational devices employed in scripts, frames and plans were conceived on purely computational and introspective grounds, with no formal links to the syntax/semantics interface. So it seems auspicious that considerations arising from problems in linguistic anaphora lead directly to script/frame-like representations, and in doing so also provide solutions to some traditionally difficult reference problems.

Schank and his collaborators were very interested in the problem of *learning* scripts by generalization from particular narratives containing repetitive patterns of events (Schank 1982). While this type of generalization learning is undoubtedly of crucial importance for AI, the present work clears some obstacles from a more direct path to the acquisition of such knowledge, namely explicit description of the general patterns in ordinary language. Lenat (1995), as mentioned earlier, has already mounted a large-scale effort to encode general knowledge directly in a frame-like representation. But this has depended so far on familiarizing the contributors to this effort with the internal representation, and entrusting to their intuitions the correct formulation of commonsense facts in this representation. Such an effort would surely benefit from an effective linguistic input channel. Language is the shared medium for the exchange of commonsense human knowledge, and as such would provide a much faster, much more natural means for knowledge bootstrapping; as well, if the mapping from the internally stored information to language were straightforward and systematic, browsing and verifying the stored knowledge, and any inferences drawn from it, would become much easier.

Further Issues

A variety of issues remain concerning the scope and details of dynamic skolemization. Most require further research, but I want to at least touch briefly on some of them.

First, a question of interest in linguistic semantics is how the proposed interpretation of donkey sentences fares with respect to the "proportion problem" (e.g., (Kadmon 1987)). For example, consider the following sentence:

33. In most instances, if a farmer owns a donkey, he is poor.

Suppose that the facts relevant to (33) are that 99 out of 100 farmers own just one donkey and are poor, while one out of 100 farmers owns 100 donkeys and is not poor. Then (33) seems true on its preferred reading, despite the fact that less than half of all farmer-owned donkeys are owned by poor farmers. The question here is how we individuate instances of farmers owning donkeys. If we treat instances as Davidsonian events (i.e., as tacit arguments of verbs, in this case *owns*), then each instance of a particular farmer owning a particular donkey is distinct, and (33) will turn out false.

One option here is to equate instances with farmers, yielding the desired reading that most farmers who own a donkey are poor. A disadvantage of this approach is that it calls for a rather complex mapping from surface form to logical form, allowing quantificational binding of an indefinite sentential subject (*a farmer*) by a frequency adverbial (*in most instances*). Be that as it may, dynamic skolemization will not complicate matters. If in the provisional logical form the farmer-variable is bound by *most* rather than by $\exists$, then the only Skolem function introduced will be one that picks out a donkey corresponding to a given donkey owner. As it happens, this function is not needed in the matrix clause of (33), since this contains no reference to the donkey. If it contained such a reference, the Skolem function would unproblematically resolve it.

To me a more attractive option is to treat instances as events (more generally, situations/ eventualities/ episodes), but not as Davidsonian ones. We regard events as being determined not just by atomic predications, but also by more complex formulas such as $(\exists y : D(y))O(x, y)$, as in Situation Semantics (Barwise & Perry 1983) or Episodic Logic (Hwang & Schubert 1993b). Then the ownership events over which (33) quantifies need no longer be so "fine-grained" that each donkey determines a distinct event. This can account for the preferred reading. However, the details would require a digression into SS or EL. Suffice it to say that dynamic skolemization would still work in such a setting, with some complications arising from the relative scopes of existential quantifiers and the operators that link formulas to events ($\models$ in SS, ** and * in EL).

A second issue concerns reference to disjunctions of indefinites, as in

34. If a farmer buys *a donkey or a mule*, he pays for *it* in cash.

Here unmodified application of dynamic skolemization will yield two Skolem functions, one for each disjunct. Yet we want to cover both alternatives with the referential pronoun (*it*). (We should note that the same problem arises in the dynamic semantics approaches.) One possible solution is to invoke the logical equivalence

$(\exists x)\Phi \vee (\exists y)\Psi \;\leftrightarrow\; (\exists x)(\Phi \vee \Psi[x/y])$

before skolemizing. Or, placing the burden on the mapping from surface form to logical form, we could allow for an interpretation of a disjunction (NP_1 or NP_2) of indefinite noun phrases as one of

$(\lambda x[\mathrm{NP}'_1(x) \vee \mathrm{NP}'_2(x)])$,

$\langle \exists x[[x = \mathrm{NP}'_1] \vee [x = \mathrm{NP}'_2]]\rangle$,

depending on whether the indefinite NPs are interpreted as predicates or as quantified terms. The angle brackets indicate an unscoped existential quantifier, and NP'_i denotes the logical translation of NP_i. (In the second version, this again contains an unscoped $\exists$-quantifier.) Unscoped quantifiers are eventually assigned sentential scope within the complete logical translation of a clause. I would argue that such a shared-variable interpretation is even possible (though dispreferred) for *con*junctions. Witness, for instance, the sentence

35. Our northern neighbor, and the best friend this country has, is currently quarreling with us over fishing rights.

Here we are referring to some x (e.g., Canada, from a US perspective) such that x is our northern neighbor *and* is the best friend this country has.

A third issue concerns reference to indefinites in negated contexts, as in the following example:

36(a) John didn't buy a donkey and pay for it in cash.

(b) $\neg[(\exists x : D(x))B(J, x) \wedge P(J, it)]$.

Let us take (36b) as the provisional logical form of (36a), where the existential quantifier has narrow scope relative to the conjunction. If we now apply the dynamic skolemization schema (21), we obtain the skolemized version $\neg(D(A) \wedge B(J, A) \wedge P(J, it))$. With *it* resolved to A, we have $\neg(D(A) \wedge B(J, A) \wedge P(J, A))$.

Does this (along with the Skolem conditional) capture the intuitive meaning of (36)? Not quite. Rather, it expresses the (weaker) proposition that if John bought just one donkey, then he didn't pay for it in cash. For, if John bought a donkey, then by the Skolem conditional, $D(A) \wedge B(J, A)$ holds, and hence from $\neg(D(A) \wedge B(J, A) \wedge P(J, A))$, $\neg P(J, A)$ holds, i.e., John did not pay in cash for the one donkey he bought. This is a consequence we want, but we also want the stronger consequence that John did not pay in cash for *any* donkey he bought, no matter how many he bought.

We can obtain a correct skolemized version of (36a) by treating the existential quantifier in (36b) as having wider scope than the conjunction. (I deliberately left the scopes in (36b) ambiguous.) For then we can resolve the anaphoric pronoun, setting it to x, *before* skolemizing. If we now skolemize, we still obtain $\neg(D(A) \wedge B(J, A) \wedge P(J, A))$. But the Skolem conditional now states that if John bought a donkey and paid for it in cash, then A is such a donkey. Since the skolemized sentence denies that A is a donkey that John bought and paid for in cash, it follows that there is no such donkey.

However, if the point of skolemization is to facilitate the interpretation of anaphoric expressions, then we need not have skolemized at all here. The situation is just as if the sentence had been

37. John didn't buy every donkey (in sight) and pay for it in cash,

where there are no indefinites to skolemize. Here, too, correct interpretation of the pronoun appears to call for wide-scoping of the quantifier (in this case $\forall$ rather than $\exists$), so that the pronoun can be identified with the quantified variable. My conjecture is that dynamic skolemization is inappropriate in negated contexts, and is not needed for representing reference relations confined to such contexts.

But what about apparent cases of external reference into negated contexts, as in

38. John didn't buy a DONKEY. It was a mule. ?

The result of dynamically skolemizing the indefinite in the negated context would again be unsatisfactory. In particular, it would fail to capture the entailment that John bought a mule. This is to be expected, since the interpretation of (38) depends upon a presupposition, mediated by the indicated stress on A DONKEY, that John did buy SOMEthing. It appears, then, that the correct referent for the pronoun comes from (the skolemized form of) that presupposition, rather than from the negative sentence.

Finally, there is the issue of anaphoric reference within modal contexts, as illustrated by (38):

38. Mary believes that if John bought a donkey at the fair, he paid for it in cash.

A little thought shows that it would be incorrect to apply dynamic skolemization without change here. We would attribute to Mary a conditional belief about a particular individual A; but she might hold that belief simply because she believes that the antecedent is false (A is not a donkey that John bought). Yet at the same time she may fail to believe that if John bought a donkey, he paid for it in cash. I think a plausible approach may be to confine the entire process of dynamic skolemization to the modal context concerned. In particular, in processing the provisional logical form derived from (38), we stipulate that Mary believes that if John bought a donkey, then he bought a donkey A (where A is a new constant). We then recast (38) to say that Mary believes that if John bought A where A is a donkey, he paid for A in cash. Thus we are attributing to Mary the same dynamic skolemization process we are employing ourselves (according to the proposed theory).

Concluding Remarks

I have shown how one can dynamically skolemize indefinites encountered in natural language input, so that the resultant constants and functions are available for subsequent reference. This provides a new way of dealing with donkey anaphora and other cases of anaphora standardly handled through some version of dynamic semantics. The advantages are that the resulting representations are context-independent and admit a standard semantics, and that referential connections can be made in cases that are problematic for dynamic semantics.

The kinds of anaphora at which dynamic skolemization is aimed pervade ordinary language, particularly the sort of language used to express commonsense generalizations about the world. So the possibility of dealing straightforwardly with heretofore refractory cases should prove very helpful in the effort to impart commonsense knowledge to machines through language.

Moreover, I showed that the representations of general facts obtained through dynamic skolemization, aided by concept definitions, are strikingly similar to the frame-like and script-like representations that are the stock-in-trade in knowledge representation and reasoning. This convergence between certain kinds of linguistically and nonlinguistically motivated representations provides further reason for optimism about the possibility of a unified approach to representation and inference.

I have also given sketchy indications that certain traditional conundrums such as the proportion problem and reference to disjunctions of indefinites should be resolvable within the proposed framework, or at least do not become more difficult within that framework. Whether dynamic skolemization is a "universal tool" for representing reference, or whether certain negated or modal contexts depend in an essential way on a dynamic semantics, is a matter for further research.

Acknowledgements

This work was supported in part by NSF grant IRI-9503312 and by ARPA/SSTO grant F30602-95-1-0025.

References

J.F. Allen and L.K. Schubert. The TRAINS project. Tech. Rep. 382, Dept. of Computer Science, Univ. of Rochester, Rochester, NY., 1991. Also slightly revised as *Language and discourse in the TRAINS project*, in A. Ortony, J. Slack, and O. Stock (eds.), Communication from an Artificial Intelligence Perspective: Theoretical Springer-Verlag, Heidelberg, pp. 91-120.

J. Barwise and J. Perry. *Situations and Attitudes.* Bradford Books, MIT Press, Cambridge, MA.

G.N. Carlson and F.J. Pelletier. *The Generic Book.* Univ. of Chicago Press, Chicago, 1995.

J. Groenendijk and M. Stokhof. Dynamic predicate logic. *Linguistics and Philosophy 14*, pp. 39-100.

P.J. Hayes. The logic of frames. In D. Metzing (ed.), *Frame Conceptions and Text Understanding*, de Gruyter, New York, 1979.

I. Heim. *The Semantics of Definite and Indefinite Noun Phrases.* Ph.D. Dissertation, U. Mass., 1982.

C.H. Hwang and L.K. Schubert. Episodic Logic: a comprehensive, natural representation for language understanding. *Minds and Machines 3*(4): Special Issue on KR for NLP, pp. 381-419, 1993a.

C.H. Hwang and L.K. Schubert. Episodic Logic: a situational logic for natural language processing. In P. Aczel, D. Israel, Y. Katagiri and S. Peters (eds.), *Situation Theory and its Applications, v.3 (STA-3)*, CSLI, Stanford, CA, pp. 303-338, 1993b.

C.H. Hwang and L.K. Schubert. Tense trees as the 'fine structure' of discourse. *Proc. of the 30th Ann. Meet. of the Assoc. for Computational Linguistics (ACL-92)*, U. Delaware, Newark, DE, June 28 - July 2, pp. 232-240, 1992.

N. Kadmon. *On Unique and Non-Unique Reference and Asymmetric Quantification.* Ph.D. Dissertation, U. Mass., Amherst, MA.

H. Kamp. A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, and M. Stokhof (eds.), *Formal Methods in the Study of Language*, Mathematical Centre-tracts, U. Amsterdam, Amsterdam, Netherlands.

H.A. Kautz. A formal theory of plan recognition and its implementation. In J.F. Allen, H.A. Kautz, R.N. Pelavin, and J.D. Tenenberg, *Reasoning about Plans*, Morgan Kaufmann, San Mateo, CA, pp. 69-125.

A. Lascarides, N. Asher, and J. Oberlander. Inferring discourse relations in context. In *Proc. of the 30th Ann. Meet. of the Assoc. for Computational Linguistics (ACL-92)*, U. Delaware, Newark, DE, June 28 - July 2, pp. 1-8, 1992.

D. Lenat. CYC: A large-scale investment in knowledge infrastructure. *Comm. of the ACM 38*(11), pp. 33-38, 1995.

M. Minsky. A framework for representing knowledge. In P.H. Winston (ed.), *The Psychology of Computer Vision*, McGraw-Hill Book Co., New York, pp. 211-277, 1975.

R.C. Schank. *Dynamic Memory: A Theory of Reminding and Learning in Computers and People.* Cambridge Univ. Press, New York.

R.C. Schank and R.P. Abelson. *Scripts, Plans, Goals and Understanding.* Lawrence Erlbaum Assoc., Hillsdale, NJ, 1977.

L.K. Schubert and C.H. Hwang. An episodic knowledge representation for narrative texts. *Proc. of the 1st Int. Conf. on Principles of Knowledge Representation and Reasoning (KR-89)*, May 15-18, Toronto, Canada, pp. 444-458, 1989.

B.L. Webber. Tense as discourse anaphor. *Computational Linguistics 14*(2), pp. 61-73, 1988.