

Discourse, anaphora and parsing*

Mark Johnson
Center for the Study of Language and Information
and Dept. of Linguistics
Stanford University

Ewan Klein
Centre for Cognitive Science
[formerly School of Epistemics]
Edinburgh University

Abstract

Discourse Representation Theory, as formulated by Hans Kamp and others, provides a model of inter- and intra-sentential anaphoric dependencies in natural language. In this paper, we present a reformulation of the model which, unlike Kamp's, is specified declaratively. Moreover, it uses the same rule formalism for building both syntactic and semantic structures. The model has been implemented in an extension of PROLOG, and runs on a VAX 11/750 computer.

1. Introduction

The study of anaphora has been a central issue in both theoretical and computational linguistics. Studies of anaphora in theoretical linguistics usually concentrate on describing the constraints on sentence-internal anaphora (e.g. (Reinhart 1983)). However, recent work by Hans Kamp (Kamp 1981) suggests that it is possible to describe some important aspects of inter-sentential anaphora while still respecting the constraints of intra-sentential anaphora. In this paper we construct a model of anaphoric dependencies which is based on Kamp's theory of Discourse Representation (DRT), but expressed in the same declarative formalism that we use for describing syntactic structure. Unlike the standard DRT approach to constructing discourse representations, our model avoids any mention of left-to-right *processing*. Note that we are not denying that there are left-to-right dependencies in anaphora, nor are we denying that these dependencies ultimately arise from the fact that earlier parts of a discourse are processed before later parts of that discourse. Rather, we claim that such dependencies should not be stated implicitly in the specification of the processing strategy, but are better expressed as part of the formal description of the model.

The idea of separating a computer program into two distinct parts: a logical specification of the problem to be solved, and a proof procedure that "interprets" this specification to actually solve the problem has been a prominent idea in recent work on logic programming, especially in the work of Kowalski. We connect directly into this tradition, in that our specification of DRS theory is provided in the form of an extended Horn-clause logic formalism.

Our system thus consists of two parts: a logical specification of DRS theory, written in a language that we have dubbed PrAtt (for Prolog with Attributes), and a simple theorem prover (interpreter) which is capable of deducing the DRSs that correspond to various input sentences using the logical specification of DRS theory.


In terms of Kowalski's (Kowalski 1979b) famous maxim "Algorithm = Logic + Control", the logical specification of the DRS theory corresponds to the "Logic", while the inference technique used by the inference engine corresponds to the "Control". Currently our inference engine uses a simple top-down proof technique (inherited from Prolog, in which

the inference engine is written), so the system as a whole (= logical specification of DRS theory + top-down theorem prover) functions essentially as a top-down predictive parser. However, this top-down behaviour is a property of the theorem prover only, and one could replace the theorem prover component with a more sophisticated proof technique such as Earley deduction (Pereira and Warren 1983) resulting in a system that used a generalization of Earley's parsing algorithm. Such a change would be a change to the theorem-prover only, since both systems would use the same logical specification of DRT.

2. A naive model of anaphoric dependency

In this section we give a brief overview of the basic ideas involved in the model. We do this by presenting a "naive" model provides the core of the analysis developed in section 5, but which ignores the complexity of syntactic structure and quantifier binding. The naive model enables us to explain our stance, independent of these complicating factors, on matters such as the changing nature of the discourse context over time, the mechanisms used to describe reference, semantic gender agreement, etc.

The following diagram (1) illustrates a naive declarative model of anaphoric dependency, where all that is required to license an anaphoric pronoun is the presence of a possible antecedent to its left.



In the naive model we conceive of a discourse context as simply consisting of a set of individual names, or *reference markers*. These represent the entities which are available to be talked about in the discourse, and play a similar role in our framework as the discourse entities of (Webber 1979). In particular, they provide the set of possible antecedents for anaphoric noun phrases. We make the simplifying assumption that the only way for a reference marker to find its way into the context is by courtesy of an indefinite description. We assume that reference markers are *typed*, and we adopt the convention of using 'f' as a marker for female gender entities, 'm' for male gender entities, and 'x' for neuter or indeterminate gender entities.¹

The vertical bars in (1) represent moments of time in the analysis of the discourse; each moment is associated with a discourse context. In this way, we can characterize a developing discourse context as a series of discrete states, each of which is localized at specific point in time and unchanging during the course of the parse.² In the diagram above, these

¹ These three types of reference marker correspond to the genders available for pronominal agreement in English. However, it seems plausible that a more complicated account of agreement would be required for those languages (e.g. French and German) in which gender marking is semantically arbitrary.

* The research reported in this paper was conducted at the The Center for the Study of Language and Information, and was made possible in part by a gift from the System Development foundation, we gratefully acknowledge financial support from the National Science Foundation (grant BNS-8309780); and Klein also acknowledges financial support from the U.K. Science and Engineering Research Council (Advanced Fellowship). Earlier versions of this paper were presented at the Summer Meeting of the Association for Symbolic Logic, July 15-20 1985, Stanford University, and the Autumn Meeting of the Linguistics Association for Great Britain, September 18-20 1985, University of Liverpool.

We would like to express our appreciation for comments and suggestions from Jo Calder, Glyn Morrill, Carl Pollard, Fernando Pereira, John Perry, Ivan Sag, Stuart Shieber, and Henk Zeevat.

contexts are shown above the bar that corresponds to the point of time at which they hold. Thus, at the beginning of the discourse the context was empty (i.e. the null set), while after the phrase *a woman* was uttered the context contained the single reference marker *f*. Consequently, $\{f\}$ serves as the context for *kissed*.

We view the meaning of a linguistic expression α as a relation between the context that precedes α and the context that follows α . That is, the meaning of α has the general form shown in (2):

$$(2) \quad \text{Preceding-Context} \mid \alpha \mid \text{Following-Context}$$

Consequently, in the naive model the discourse context is determined by a series of equations relating the context which immediately precedes a lexical item to the context which immediately follows that item. For individual words, this relation is part of the lexical specification. To illustrate, the semantic contribution of *woman* is given here by (3a), or more generally, as (3b).

$$(3a) \quad \{\} \mid \text{woman} \mid \{f\}$$

$$(3b) \quad C \mid \text{woman} \mid C \cup \{f\}$$

The anaphoric pronoun *her* behaves in a very different fashion to indefinite noun phrases. Rather than adding a reference marker to the following context, it looks in the preceding context for a reference marker of the right sort (i.e. one that agrees with it in number and gender). If there is no such antecedent marker, the pronoun cannot be interpreted as anaphoric. The meaning of anaphoric *her* is the relation in (4).

$$(4) \quad C \mid \text{her} \mid C \text{ iff } f \in C$$

where *f* is the reference marker associated with *her*

A sequence of discourse contexts is well-formed for a string if all of the relations associated with the lexical items in the string hold; i.e. the discourse contexts are a solution to the relational equations. Sometimes these equations will have a single solution; in that case, the discourse is unambiguous. However, usually the equations have multiple solutions, which means, in effect, that the discourse has many interpretations. This arises, in the present discussion, when a pronoun has several possible antecedents.³ On the other hand, it is also possible that the equations have no solution at all. This case arises when a pronoun is used in a discourse context that contains no appropriate reference marker at all.

At a more abstract level, we can view this model as one in which the context is a stream of reference markers, which is threaded from one lexical item to the next. The equations associated with individual lexical items act as (possibly non-deterministic) operators on their input stream to produce an output stream, which serves as the input to the following lexical item.

One of the main virtues of this simple picture is that it invites comparison with other ideas. Our proposed notion of meaning is clearly reminiscent of the claim in (Barwise and Perry 1983) that meaning is a relation between different types of situation, though it also has its roots in earlier work on indexical semantics, such as (Stalnaker 1972). Second, it is also

² It seems that this technique of factoring a single non-monotonic representation into a series of monotonic ones is applicable in many areas other than the one discussed here. At an abstract level it is similar to the technique discussed by (Kowalski 1979a). It is also similar to the use of *difference lists* in logic programming, since the "content" of a particular element is the difference between its "output" and its "input".

³ In such a case, our program merely *enumerates* all possible interpretations, which results in the familiar combinatorial explosion of solutions. A better technique, which we cannot explore here, would be to factor out the ambiguity and localize it in the representation.

reminiscent of the technique used in logic programming known as difference lists (Pereira 1985) or threading.

3. Discourse Representation Theory

The naive model presented in the last section ignored all syntactic and lexical interactions with the "left-to-right" nature of anaphoric dependency. The fatal flaw of this account is that it fails to explain the anaphoric properties of universally quantified NPs. The data which shows this is well known, and some illustrative cases are given in (5) to (7).

- (5) a. A woman_i went home. She_i was tired.
b. Every woman_i went home. She_i was tired.

- (6) a. Every man_i thought he_i was ill.
b. Lee gave every woman_i her_i prize.

- (7) a. Every man saw a woman_i. She_i was going home.
b. Every woman who kissed a man_i loved him_i.

(5) shows that a universal NP does not normally act as an antecedent for pronouns in a following sentence.⁴ According to the variable-binding paradigm of anaphora, this follows because a universal can only enter into an anaphoric relation with pronouns that are in its scope. For our current purposes, it is not important whether scope is determined in terms of a tree-geometrical notion like c-command (Reinhart 1983), or in terms of function-argument structure, as proposed by (Ladusaw 1980) and (Bach and Partee 1980); in either case, it is clear that the scope of the universal in (5) is that portion of the first sentence that we have italicised. Examples (6) illustrate cases where a universal does enter into an anaphoric relation with a pronoun in its scope (again indicated by italicisation). (7) is intended to indicate the interaction between indefinites and universals. In (7a), the indefinite has narrower scope than the universal, and it is thereby incapable of acting as an antecedent for a pronoun such as the following *she* which is outside the scope of the universal. By contrast, when both the indefinite and the pronoun fall within the scope of a universal, as in (7b), an anaphoric link is permissible. Note that (7b) is a so-called 'donkey' sentence.

The study of these syntactic and lexical effects has been a central theme of modern theoretical linguistics, but most work within this paradigm has concentrated almost exclusively on intra-sentential anaphora. However, recently (Kamp 1981), (Heim 1982) and (Haik 1984) have developed theories capable of providing a unified account of the main properties of intra- and inter-sentential anaphora. We will base our account on Kamp's Discourse Representation Theory, and in this section, we briefly outline those aspects of Kamp's model which are of most relevance to us.

DRT is intended to explicitly capture the distinctions in anaphoric potential exhibited by (5a) and (5b), while simultaneously providing a basis for truth-conditional semantic interpretation. Thus (5a) would be associated with a DRS of the form (8).

$$(8) \quad \boxed{\begin{array}{c} f \\ \text{woman}(f) \\ \text{went-home}(f) \\ \text{tired}(f) \end{array}}$$

⁴ Sentences like (i) are exceptions to this generalization:

- (i) Every man will like this car. He'll certainly want to drive it.

Rather than abandoning the generalization altogether, it seems more fruitful to adopt the hypothesis that such discourses involve 'modal subordination' (Roberts 1986) of the second sentence to the first. However, we do not understand the precise mechanics of this process.


A discourse representation has two parts: a 'universe' consisting of set of discourse markers (in this case a singleton set) and set of conditions. The sentence *A woman went home* licenses the introduction of the reference marker *f* into the universe of the DRS, and this marker is also entered as the argument of the predicate *went-home*. When *She was tired* is analyzed, the pronoun can be interpreted as anaphoric on a preceding NPs if the marker licensed by that NP is 'accessible'; i.e. if the marker belongs to the universe of the immediately enclosing DR or a superordinate one. Since *f* is accessible, the pronoun *her* can be identified with it to yield the condition *tired(f)*.

Before turning to sentences involving universal NPs, it will be useful to consider in a little more detail the procedure for constructing a DR like (8) proposed by (Kamp 1981). Kamp's rules pivot on the noun phrases in a sentence, and depend particularly on any determiners in the noun phrases. It is useful to think of every determiner as having a semantic *restrictor* and a semantic *scope*. The determiner will bind an argument position in each of these. Thus, in a simple intransitive sentence like the first sentence of (5), the restrictor of *a* is *woman()*, while its scope is *went home()*, where the empty parentheses indicate an open argument position. Given an existing (possibly empty) DRS *K*, a sentence of the form *[[a Res] Scope]* is "processed" in the following manner:

- (i) add a new reference marker *x* to the universe of *K*;
- (ii) fill the argument slot in *Res* by *x*, and add the resulting clause to the conditions of *K*; and
- (iii) fill the argument slot in *Scope* by *x*, and recursively call any applicable construction rules to process the resulting string.

Let us turn now to sentences involving universals. The DR associated with (5b) is illustrated in (9).

(9)



The universal quantifier *every* triggers the introduction of two subordinate DRSs, linked by the relation $=>$; this corresponds roughly to implication in first order logic. When we come to analyze the second sentence of the discourse, *She was tired* the reference marker licensed by *every woman* is trapped in the subordinate DRS; it is not accessible at the top level of the discourse. Consequently, the only option is to treat the pronoun *she* as non-anaphoric, which we have indicated here by associating it with a distinct reference marker.

When we consider sentence-internal anaphora, the antecedent-introducing potential of *every* and *a* converge. For example, in both of the following sentences, *he* can be anaphoric to the subject NP:

(6a) Every man_i thought he_i was ill.

(10) A man_i thought he_i was ill.

Although it may not be obvious from the examples given so far, DR theory correctly predicts that the reference markers associated with an indefinite or universal NP in subject position will be anaphorically accessible to pronouns that it c-commands.⁵ To see why, we need to consider in a little more

⁵ It might be argued that DR theory fails to provide an adequate semantic distinction between a 'c-command binding' relation and a 'discourse anaphora' relation, as proposed for example by (Reinhart 1983) in order to account for the strict/sloppy ambiguity in VP ellipsis. Whether this criticism is justified or not depends in large part on the appropriate analysis of such ellipsis phenomena in the DR framework. For some discussion, see (van Eijck 1985), (Klein 1985), (Roberts 1984).


detail the way in which DR's are constructed on Kamp's approach.

Construction rules apply to sentences on a top-down, left-to-right basis. Given a sentence like (6a) or (10), the first constituent to be processed is the subject NP. We either stay in the current DR, if the determiner is *a*, or 'push down' to an embedded DR if the determiner is *every*. (This embedded DR is, therefore, the antecedent box of the conditional like that displayed in (9).) A discourse marker *x_i* is introduced into the universe of whatever is now the current DR, and *x_i* also becomes the argument of the subject nominal (e.g. *man(m_i)*) and the first argument of the predicate VP (e.g. *m_i thought he was ill*). When the VP is processed, there are again two cases, depending on whether the subject determiner was *a* or *every*. In the first case, we enter the new conditions licensed by the VP into the current DR. In the second case, we close off the current (antecedent) DR, and open a new embedded DR which forms the consequent box of the conditional. Kamp claims that the reference markers *accessible* as antecedents to a given pronoun occurrence consist of those reference markers which are present in the universe of either the current DR or of any DR's which are superordinate to the current DR. Of two DR's *K₁* and *K₂*, *K₁* is superordinate to *K₂* if:

- (i) *K₂* is embedded in *K₁*, or
- (ii) if *K₁* is the antecedent of a conditional of which *K₂* is the consequent, or
- (iii) if there is some *K₃* such that *K₁* is superordinate to *K₃* and *K₃* is superordinate to *K₂*.

This is illustrated in (11) diagram below, where the lightly shaded boxes are all superordinate to the darkly shaded box.

(11)



Consider now what follows when we come to process the NP *he* in either (6a) or (10). It can be anaphorically linked to any reference marker which is accessible to it, and this will of course include the marker *x_i* introduced by the subject NP.

Let us now attempt to summarise the salient features of DRT. Note, first, that every noun phrase is associated with a 'space'⁶ in a Discourse Representation. Referential terms – which we take to include definite and indefinite descriptions, proper names, and definite pronouns – are entered into an existing space. By contrast, universally quantified NPs induce a new subspace.

Second, the space associated with an NP represents both the quantificational scope of the NP and its anaphoric domain.

Third, the boundaries of these spaces are not coterminous with clause or sentence boundaries. A clause containing universal NPs will induce a number of subspaces; conversely, the space associated with a referential NP can encompass indefinitely many sentences of a given discourse.

Fourth, the space of an indefinite NP which occurs within the scope of a universal NP is the same as the space of the universal.

4. The flow of anaphoric information


In the last section we showed how DRT is able to simultaneously describe both the semantics of quantification and the anaphoric 'range' of referential noun phrases in terms of a single discourse representation. The standard version of DRT depends crucially on processing notions in order to explain the failure of anaphora in examples like (12).

(12) He_i liked a boy_i.


Since the reference marker for *a boy* is not introduced into the DR until after the pronoun *he* is introduced, it is unavailable as a possible antecedent. That is, the failure of anaphora is explained by assuming that the pronoun's antecedent is assigned at the time at which it is introduced into the DRS, and that the reference marker for the noun phrase is introduced after the pronoun was introduced.

In a declarative framework, an explanation in terms of processing order is impermissible hence we represent left-to-right dependencies by explicit equations. Although these equations are in principle non-directional, it can be helpful to think of them as providing a means for transmitting information from one node in the syntactic structure to another.

Bottom-up information flow is central to syntax-driven compositional semantics of the familiar sort: semantic values are associated with the leaves of the syntax tree, and the semantic value of a complex constituent is determined as a function of the semantic values of the constituents daughters. The diagram in (13) shows this direction of information flow.



Although this approach has proven to be extremely powerful, it is awkward and intuitively unsatisfactory as a means for dealing with anaphoric dependencies. Even if much semantic information is indeed composed on a bottom-up regime, it seems highly plausible that anaphoric information – that is, information about the set of available antecedents – flows in a left-to-right direction. We have already seen that a simple left-to-right model of this information flow can be constructed by regarding meaning as a relation between contexts, but we have also seen that such a model is inadequate for dealing with the facts of bound anaphora. A more satisfactory model can be constructed by reflecting on the principles involved in constructing Discourse Representations. As we pointed out in the previous section, Kamp's construction rules centre on the determiners *a* and *every*, since they trigger the introduction of reference markers, the binding of argument positions, and the introduction of sub-spaces. What we shall suggest, therefore, is that information about possible antecedents flows from a determiner to the determiner's restrictor, and from the restrictor to the determiner's scope. The following diagram (14) illustrates how this top-down, left-to-right flow is integrated with the orthodox phrase marker of *a girl kissed a boy*.




⁶ This term is intended to be reminiscent of work by Fauconnier (1985) on mental spaces, and by Reichman-Adar (1984) on context spaces, though considerable work needs to be done in showing that these ideas are in fact compatible.

The light, incoming lines on the left-hand side of a node indicate incoming information about the set of possible antecedents. This set will be encoded in something we call the "in-list". The light lines on the right-hand side of a node indicate outgoing information about antecedents, encoded in the form of an "out-list". In general, the out-list of any node will be its in-list plus any additional information added by that node. Circled nodes mark constituents that supplement their in-list with new reference markers. The in-list and the out-list together form a difference list, in that the content added by any item is the difference between its in-list and out-list.

Alternatively, one can view the in-lists and the out-lists of nodes as streams along which information about antecedents flows: this anaphoric information is threaded through the syntactic tree structure. Notice that we assume the sentence as a whole will be fed an in-list which is supplied by the preceding discourse. Moreover, the sentence as a whole will also produce an out-list, which will provide potential antecedents for following discourse.

The next diagram (15) illustrates the flow of information for *every girl kissed a boy*.



By contrast with (14), the out-list from the VP, containing reference markers for *girl* and *boy*, is "trapped" at that level rather than percolating up to the S node. The out-list for the sentence as a whole is just the sentence's in-list. This captures the idea from binding theory that the scope of a quantifier is normally limited to its c-command domain (Reinhart 1981, Reinhart 1983); in terms of DRT, it corresponds to the closed subspace that is associated with universal NPs.

Let us summarize our claims so far. We have suggested that there is a contrast between the bottom-up information flow of compositional semantics, on the one hand, and the top-down flow that is naturally associated with anaphoric information. We have also suggested that top-down flow is largely determined, according to the principles of DRT, by the lexical properties of determiners and their structural position in the sentence.

One possible implementation of this analysis would be to factor out anaphoric, contextual information from the rest of semantics, and to use two distinct mechanisms for building the two kinds of representation. However, such an approach fails to explain why the spaces in a DR, and the list of contextually-given antecedents always covary; that is, when a new DR subspace is opened, a new context list begins, and when a DR subspace is closed, a context list is simply "dropped", i.e. it does not serve as the in-list to any other expression. Indeed, the fact that a DRS in Kamp's theory consists of a universe, corresponding to our context list, and a set of conditions, corresponding roughly to compositional semantic information, suggests that it out to be possible to enrich the notion of a context from being just a list of antecedents to being a whole DR structure.

In our analysis, then, we thread a list through the syntactic structure which contains both conventional semantic information and information about available antecedents, so that an expression mapping an incoming context into an outgoing context does more than incrementing the set of possible antecedents: it also adds conditions to the context that correspond to its truth-conditional semantics.

It is necessary that the context be structured, rather than a simple list, as it was in the naive model, and as discussed above. This is because we need to be able to incorporate the semantic structures associated with all expressions, even those that are anaphorically opaque to following anaphora. In the model described immediately above, we accounted for the anaphoric opacity of an expression by "dropping" its context list after it had been processed, but such "dropping" in a system where the context lists also contain "compositional" semantic information would result in that semantic information also being lost.

Rather, we structure the context list as an ordered list of the currently open DR spaces, starting at the most embedded space, and working upward through the superordinate spaces. For example, the context list for an item located in DR space K_1 in (11) would be [K_1, K_2, K_3, K_4], where each K_i is a set of reference markers and conditions, the current contents of the corresponding space. The first space on the context list is the most embedded space, i.e. the *current space*, and identifies the place where new conditions and reference markers are to be added. Since the context list consists of the active space plus all of the spaces superordinate to it, any reference markers contained in these spaces are possible antecedents for anaphora in the active space.

5. The Grammar

We turn now to considering the induction of DRSs. In this section we describe a simplified version of the grammar that we have implemented. The grammar presented here is the actual input to the proof procedure: the parser is nothing more than a declarative statement of the well-formedness conditions of an utterance, plus a proof procedure capable of determining whether or not these conditions actually hold of a given utterance.

The rules are written in DCG format (Clocksin and Mellish 1984) in a superset of Prolog that we developed in this project. This language, which we have dubbed PrAtt (for Prolog with Attributes), allows an attribute-value notation as well as the standard position-value notation of Prolog. For example, the expression "N:syn:index" refers to the value of the *index* attribute of the *syn* attribute of the variable N.

We make heavy use of the attribute-value notation to represent feature bundles associated with constituents. Two attributes that are present on every constituent are *syn* (for "syntax") and *sem* (for "semantics"). The *sem:in* and *sem:out* attributes contain the context in-lists and out-lists respectively, while the *syn* attribute holds information used to construct the function-argument structure of the clause.

Expressions act on the context list by opening or closing spaces (i.e. pushing or popping spaces from the context list), adding reference markers and conditions to the active space, and looking through all of the spaces in the context list for antecedents for anaphora.

Consider, for example, the common noun *woman*. It inserts a reference marker *f* and a condition *woman(f)* into the active space. Using our earlier relational notation, we can express its meaning as follows:⁷

$$(16) \quad [\text{Active|Super}] \mid \text{woman} \mid [[f, \text{woman}(f) \mid \text{Active}]]\text{Super}]$$

In our implementation, this would be written as in (17).

$$(17) \quad n(N) \text{ --> } [\text{woman}], \\ \{ N:\text{syn:index} = w, \\ N:\text{sem:in} = [\text{Current} \mid \text{Super}], \\ N:\text{sem:out} = [[w, \text{woman}(w) \mid \text{Current}] \mid \text{Super}] \}.$$

⁷ We use standard Prolog notation here: variables begin with a capital letter, constants with a lower-case letter, "[x,y]" is the list that contains x and y, and "[x|y]" is the list that consists of x CONSed onto y.

The bracketed equations are conditions that must be satisfied in rewriting an N to the lexical item *woman*. The first equation assigns a reference marker to the lexical item,⁸ the second equation analyses the incoming context list into two parts, the current space (Current) and a list of the superordinate spaces (Super), while the third equation requires the active space of the outgoing context list to contain the reference marker and the condition associated with the noun.

A sample entry for a verb is shown in (18). Again, the equations associated with the lexical entry dissect the incoming context into the current space and a list of superordinate spaces, and place the condition associated with the verb into the outgoing context.

$$(18) \quad v(V) \text{ --> } [\text{saw}], \\ \{ V:\text{sem:in} = [\text{Current} \mid \text{Super}], \\ V:\text{sem:out} = \\ [[\text{saw}(V:\text{syn:arg1}, V:\text{syn:arg2}) \mid \text{Current}] \mid \text{Super}] \}.$$

One interesting property of this rule is that it is responsible for placing a condition into the context that in essence represents the compositional semantics of the entire clause. The *syn* attributes of constituents are used to connect the NP arguments of the verb with the verb itself; thus the necessary information to build the condition associated with the entire clause is available at the verb. One can view the equations in the phrase structure rules associated with the *syn* attribute as directing information from the NP arguments inward and downward to the verb.

The crux of the grammar is located in the lexical entries for determiners, as hinted earlier. (19) contains the lexical entry for the indefinite article *a*.

$$(19) \quad \text{det}(\text{Det}) \text{ --> } [a], \\ \{ \text{Det:sem:res:in} = \text{Det:sem:in}, \\ \text{Det:sem:scope:in} = \text{Det:sem:res:out}, \\ \text{Det:sem:out} = \text{Det:sem:scope:out} \}.$$

As we shall see later, the phrase structure rules are written in such a way that the value of the clauses's *sem* attribute is equal to its subject's determiner's *sem* attribute, and the *semantics* attribute of the restrictor and the scope of a clause are placed in that determiner's *sem:res* and *sem:scope* attributes respectively. As noted earlier, an indefinite determiner does not cause the creation of any additional subspaces, rather the restrictor and the scope are simply placed into the current active space. Therefore, the equations associated with the indefinite determiner simply connect the in-list associated with the sentence to the restrictor's in-list, feed the restrictor's out-list to the scope's in-list, and take the out-list from the scope as the out-list for the clause as a whole.

The lexical entry associated with the universal quantifier *every* is a little more complicated. It must create two new spaces, one for the restrictor, the other for the scope, and the finally close off both spaces, and build the structure associated with the clause as a whole.

$$(20) \quad \text{det}(\text{Det}) \text{ --> } [\text{every}], \\ \{ \text{Det:sem:res:in} = \\ [[] \mid \text{Det:sem:in}], \\ \text{Det:sem:scope:in} = \\ [[] \mid \text{Det:sem:res:out}], \\ \text{Det:sem:scope:out} = \\ [\text{Scope}, \text{Res} \mid [\text{Current} \mid \text{Super}]], \\ \text{Det:sem:out} = \\ [[(\text{Res} ==> \text{Scope}) \mid \text{Current}] \mid \text{Super}] \}.$$

⁸ For simplicity here we have reference markers directly to lexical entries; however more correctly the reference markers should be assigned to lexical tokens, allowing two occurrences of the same lexical entry to refer to different objects in the world.

The first equation in (20) pushes a new, empty space onto the determiner's in-list as the active space, and makes that list the restrictor's in-list. The second equation takes the restrictor's out-list pushes another new, empty space onto it, and makes the resulting list the scope's in-list. The final equation takes the scope's out-list, removes the two spaces that were added for the restrictor and the scope, and produces a new list in which the original active space has a complex condition added to it representing the whole universally quantified expression. This last list serves as the out-list for the determiner, and hence for the clause as a whole.

Below are the phrase structure rules responsible for connecting the various attributes of the constituents as described above.

- (21) $np(NP) \rightarrow \{ NP:sem = Det:sem, \\ Det:sem:res = N:sem, \\ NP:syn = N:syn \}, \\ det(Det), n(N).$
- (22) $vp(VP) \rightarrow \{ VP:sem = NP:sem, \\ NP:sem:scope = V:sem, \\ VP:syn:arg2 = NP:syn:index, \\ VP:syn = V:syn \\ \}, \\ v(V), np(NP).$
- (23) $s(S) \rightarrow \{ S:sem = NP:sem, \\ S:syn = VP:syn, \\ NP:sem:scope = VP:sem, \\ VP:syn:arg1 = NP:syn:index \\ \}, \\ np(NP), vp(VP).$

It remains only to give the lexical entry associated with pronouns, and our fragment is complete. This is given in (24).

- (24) $np(NP) \rightarrow [her], \\ \{ member(Space, NP:sem:in), \\ member(NP:syn:index, Space), \\ type(NP:syn:index, feminine), \\ NP:sem:scope:in = NP:sem:in, \\ NP:sem:out = NP:sem:scope:out \}.$

The first three equations require that there be some space containing a reference marker of feminine type with which the pronoun's reference marker can unify;⁹ the last two equations take account of the fact that an anaphoric pronoun, while not adding any conditions of its own to the context, can appear in subject position, and thus can have a scope expression.

We have now completely described our declarative formulation of DRS theory. This formulation (together with phrase structure rules that analyse a discourse as a series of sentences) suffices to obtain the analyses shown below.¹⁰

- (25) *Every woman chased a donkey.*
 $DRS = [[w, woman(w)] = => [chased(w, d), \\ d, donkey(d)]]$
- (26) *A woman chased a donkey. Every boy saw her.*
 $DRS = [[b, boy(b)] = => [saw(b, w)], \\ [d, donkey(d)] = => [chased(w, d)], \\ w, woman(w)]$

⁹ The definition of *member* used here is the conventional one used in Prolog (albeit interpreted by the PrAtt interpreter, while the *type* predicate is a set of clauses of the form $type(w, feminine)$, etc.

¹⁰ Note that because later elements are pushed onto the front of a DR space, the order of the elements in the DR spaces is the reverse of their "normal" presentation. This does not affect their truth conditional semantics, however.

We have also implemented a more complex version of this grammar incorporating a treatment of unbounded dependencies, and obtained analyses like the following:

- (27) *Every man who owns a donkey beats it.*
 $DRS = [[man(m), owns(m, d), d, donkey(d), m] \\ = => [beats(m, d)]]$

The parser indicates ill-formedness of its input in the standard Prolog fashion, viz. it fails to find a well-formed DRS for the input sentence.

- (28) *A woman who loves every man kissed him.*
 no

6. Conclusion

The declarative reformulation of DRS theory proposed here is relatively faithful to Kamp's original formulation, but has the advantage that it inherits a fully specified declarative and procedural semantics from the underlying Prolog system. It emphasises the view that expressions of the language can be viewed as relations between preceding and following contexts, and shows how these relations can be specified in a formally precise way.

This model opens up several important questions. Kamp showed that the treatment of anaphoric dependencies, normally viewed as a left to right phenomenon, can be integrated with the treatment of the "conventional" truth-conditional semantics of clauses: we have shown that both of these can be integrated into an extended unification-based model of grammar. This integration allows one to be precise about the nature of the syntax/semantics/discourse interface(s), and also allows experimentation with respect to the analysis of specific linguistic phenomena. For example, in our larger grammar (not presented here) we capture strong and weak cross-over phenomena by introducing the reference marker associated with a relative clause NP when the corresponding gap is reached. We are thus analysing what is usually thought of as a syntactic phenomenon in terms of the accessibility of reference markers, a discourse property.

From a computational point of view, there is a delicate interaction between the specific rules adopted in declarative formulation of the theory and the "power" of the inference procedure needed to determine the well-formedness of a particular utterance with respect to them. The top-down left-to-right inference procedure inherited from Prolog suffices for the grammar presented here, but one can easily write grammars in PrAtt for which this inference procedure may fail to terminate. We are investigating other inference procedures, such as Earley Deduction (Pereira and Warren 1983) and Left Corner parsing to see if they have better termination properties. Essentially, the problem is one of arranging the equations in the grammar to be applied in an order such that the search space is finite: thus research on various corouting strategies, such as the use of the *freeze* predicate is relevant here.

7. References

- Bach, E. and Partee, B. (1981) Anaphora and semantic structure. In Masek, C. S., Hendrick, R. A. and Miller, M. F. (eds.) *Papers from the Parasession on Language and Behavior at the Seventeenth Regional Meeting of the Chicago Linguistics Society*, Chicago, May 1-2, 1981, pp1-28.
- Barwise, J. and Perry, J. (1983) *Situations and Attitudes*. Cambridge, Mass.: MIT Press.
- Fauconnier, G. (1985) *Mental Spaces: Aspects of Meaning Construction in Natural Language*. MIT.
- Haik, I. (1984) Indirect Binding. *Linguistic Inquiry*, 15, 185-223.
- Hein, I. (1982) The Semantics of Definite and Indefinite Noun Phrases. PhD Thesis, University of Massachusetts. Distributed by Graduate Linguistics Student Association.
- Kamp, H. (1981) A Theory of Truth and Semantic Representation. In Groenendijk, J. A. G., Janssen, T. M. V. and Stokhof, M. B. J. (eds.) *Formal Methods in the Study of Language*, Volume 136, pp277-322. Amsterdam: Mathematical Centre Tracts.
- Klein, E. (1985) VP Ellipsis in DR Theory. In Groenendijk, J. and Stokhof, M. (eds.) *Fifth Amsterdam Colloquium*, Amsterdam, 1985. To appear.
- Kowalski, R. A. (1979) *Logic for Problem Solving*. Amsterdam: North Holland.
- Kowalski, R. (1979) Algorithm = Logic + Control. *Communications of the ACM*, 22, 424-436.
- Ladusaw, W. A. (1980) Polarity Sensitivity as Inherent Scope Relations.
- Pereira, F. C. N. and Warren, D. H. D. (1983) Parsing as Deduction. In *Proceedings of the 21st Annual Meeting of the Association for Computational Linguistics*, Massachusetts Institute of Technology, Cambridge, Mass., June, 1983, pp137-144.
- Pereira, F. C. N. (1985) A Structure-Sharing Representation for Unification-Based Grammar Formalisms. In *Proceedings of the 23rd Annual Meeting of the Association for Computational Linguistics*, University of Chicago, Chicago, Illinois, 8-12 July, 1985, pp137-144.
- Reichman-Adar, R. (1984) Extended Person-Machine Interface. *Artificial Intelligence*, 22, 157-218.
- Reinhart, T. (1981) Definite NP anaphora and C-Command Domains. *Linguistic Inquiry*, 12, 605-631.
- Reinhart, T. (1983) Coreference and bound anaphora: A restatement of the anaphora questions. *Linguistics and Philosophy*, 6, 47-88.
- Roberts, C. (1984) Anaphora, coreference and the binding theory. Unpublished paper, University of Massachusetts, Amherst.
- Stalnaker, R. C. (1972) Pragmatics. In Davidson, D. and Harman, G. (eds.) *Semantics of Natural Language*, Synthese Library, pp380-397. Dordrecht: D. Reidel.
- van Eijck, J. (1985) Aspects of Quantification in Natural Language. PhD Thesis, University of Groningen. Unpublished PhD Thesis.
- Webber, B. L. (1979) *A Formal Approach to Discourse Anaphora*. London: Garland Publishing.