

BUILDERS: AN IMPLEMENTATION OF DR THEORY AND LFG

Hajime Wada
Department of Linguistics
The University of Texas at Austin

Nicholas Asher
Department of Philosophy
Center for Cognitive Science
The University of Texas at Austin

ABSTRACT

This paper examines a particular PROLOG implementation of Discourse Representation theory (DR theory) constructed at the University of Texas. The implementation also contains a Lexical Functional Grammar parser that provides f-structures; these f-structures are then translated into the semantic representations posited by DR theory, structures which are known as Discourse Representation Structures (DRSs). Our program handles some linguistically interesting phenomena in English such as (i) scope ambiguities of singular quantifiers, (ii) functional control phenomena, and (iii) long distance dependencies. Finally, we have implemented an algorithm for anaphora resolution. Our goal is to use purely linguistically available information in constructing a semantic representation of discourse as far as is feasible and to forego appeals to world knowledge.

Introduction : DR Theory

DR theory, versions of which have been developed independently by Hans Kamp (1981) and Irene Heim (1982), has several motivations, but certainly one of the principal ones was to examine the anaphoric properties of definite and indefinite noun phrases. Out of this study emerged a novel theory of indefinites and definites that has provided, among other things, a semantic theory of intersentential anaphora. One of our goals in implementing DR theory for a fragment of English was to investigate how a purely linguistic theory of anaphora might help in the real time processing of anaphors by automated natural language understanding systems.

DR theory has two basic components: the **DRS construction algorithm** and the **correctness definition**. The DRS construction algorithm is a mapping from natural language discourses to DRSs, the semantic representations that the theory posits. The correctness definition maps DRSs into models that (i) provide the truth conditions of one sentence in isolation and (ii) show how the content of one sentence contributes to the content of the antecedent discourse it follows. We have implemented only a fragment of the presently developed DRS construction algorithm.

We want to argue that for philosophical reasons it makes little sense to try to implement on a computer not only the DRS construction algorithm but also the correctness definition. To understand why we hold this view, however, we have to provide at least a very sketchy overview of these two components of DR theory.

The language in which DRSs are described is to begin with quite simple. Its vocabulary consists of a set of individual **reference markers** ($x, y, z, x_1, \text{etc.}$), a set of mental state markers ($p, p_1, p_2, \text{etc.}$), and a set of n -ary **predicates**, for which English nouns, verbs and intersective adjectives will serve fine. We also have certain logical symbols in the DRS language: \neg, \vee, \Rightarrow .

Next, we define **conditions** and DRSs by a simultaneous recursion. (I will use boldfaced letters of the appropriate type as metalinguistic variables for reference markers.)

Definition 1:

1. Suppose that ϕ is an n -ary predicate and x_1, \dots, x_n are reference markers. Then $\phi(x_1, \dots, x_n)$ is an atomic condition.
2. Suppose x_1 and x_2 are reference markers. Then $x_1 = x_2$ is an atomic condition.
3. A DRS K is a pair of sets $\langle U, \text{Con} \rangle$, where U is a set of reference markers and Con a set of conditions.
4. Let K_1 and K_2 be DRSs and let p be a mental state reference marker. Then $\neg K_1, K_1 \vee K_2, K_1 \Rightarrow K_2$, and $p:K_1$ are conditions.

This framework yields a treatment of indefinite and definite noun phrases that has made an important contribution to understanding the anaphoric and "pseudo referential" behavior of indefinites. Let us first briefly sketch the theory's treatment of indefinites. When processed by the DRS construction algorithm, a singular indefinite introduces into a DRS a reference marker that functions essentially as a free variable. This reference marker can be identified with other reference markers that are introduced by anaphoric pronouns at potentially unlimited distances from the original indefinite NP. Indefinites get their existential force by the way a DRS is assigned truth conditions.

To take a simple example that uses only indefinites and anaphoric pronouns, the DRS construction algorithm yields the DRS in (2) for

(1) A man loves a woman. She is beautiful.

(2)

<1> $x, <2> y, <5> z$
<0> [A man loves a woman]
<1> [loves a woman(x)]
<2> woman(y)
<3> loves(x, y)
<4> [She is beautiful]
<5> $z = y$
<5> [is beautiful(z)]
<6> beautiful(z)

The informal interpretation of (2) is that there are three objects (corresponding to x, y and z) that have the respective properties ascribed to x, y and z in the conditions <1> - <6>. The numbers to the left of the reference markers and conditions indicate how the construction algorithm might proceed from step to step in a top down algorithm on some sort of standard parse tree like an LFG c-structure. Initially, we begin with the unanalyzed sentence within a DRS K . The subject N(oun) P(hrase) node in the parse tree introduces a condition into the DRS, the condition $\text{man}(x)$, into Con_K and a reference marker x into U_K . That exhausts the content of the indefinite noun phrase. The

V(erb) P(hrase) node introduces an intermediate step in the DRS that is further broken down by the algorithm as it goes down the parse tree; we shall write this intermediate step as an as yet not fully analyzed condition -- 'loves a woman(x)'. The reference marker introduced by the NP already processed (the subject NP) is an argument of that condition. The unanalyzed condition is then broken down by the algorithm as it goes down the parse tree; the NP that is a constituent of the VP introduces another reference marker and a condition 'woman(y)' into K; and finally, the verb itself introduces a condition into Con_K.

After step <3>, the algorithm has finished with the first sentence in (1). It now processes the second sentence, using the background DRS constructed from the first as a context for the interpretation of the new input. The subject noun phrase of the second sentence is simply an anaphoric pronoun. Because the pronoun is an NP, the algorithm requires that it introduce a new reference marker into the DRS. Because the pronoun is anaphoric, the reference marker it introduces must be linked to some already introduced reference marker in U_K. Thus, U_K provides a set of contextually "discourse individuals" (objects that have been talked about in the discourse) that can be referred to by NPs in subsequent discourse. Once, the algorithm is finished with 'she', the rest of the second sentence is processed in the same sequence of steps used to process the first.

Our implementation in general follows the left to right process described above. Using f-structures as inputs to the construction algorithm allows us to bypass the steps that are enclosed within the square brackets, since we have already available in the f-structure the basic predicate argument structure of this simple sentence. The f-structures play a more important role in more complex sentences like those involving relative clauses; there they make transparent the way in which the predicate argument structure of subordinate clauses fits together with the predicate argument structure of the main clause.

Accessibility


The simple example in (1) already introduces the problem of anaphoric coreference. DR theory, like all linguistic theories, provides constraints on what are the possible, antecedent, coreferential NPs for an anaphoric, pronominal NP. We shall say that a reference marker *u* introduced by an antecedent NP can be linked anaphorically to a reference marker *v* introduced by a pronominal NP just in case *u* is *accessible* to *v*. *u* is accessible to *v* just in case <*u,v*> belongs to the transitive closure of the relation that is defined as follows: a) *u* has already been introduced into U_K prior to the introduction of *v*, b) *u* ∈ U_K,

and there is a K" such that K' ⇒ K is a condition in K", c) K occurs as a component of some condition in K' and *u* has already been introduced into U_K prior to the introduction of the condition containing K.¹ Our implementation uses this notion of accessibility to constrain the process of anaphora resolution.

While singular indefinite NPs simply introduce, when processed by the construction algorithm, reference markers and atomic conditions into the DRS, "quantificational" NPs (those involving determiners like 'every', 'each', and also 'many') will introduce logically complex conditions (see Kamp, (1981), Frey & Kamp (1985) for details).

(3) Every man loves a woman

yields on the "default" left-right scope reading the following DRS:




As we shall see shortly, the truth conditions of such a DRS are essentially this: pick any man; then there's a woman that loves him. The DRS in (4) differs from the one that is the result of processing the first sentence in (1), in that the determiner 'every' makes a decidedly different contribution from a pure indefinite. 'Every' introduces a logically complex structure between two DRSs; it corresponds in first order logic to a universal conditional.

Our implementation also countenances relative clauses and the truth functional sentential conjunctions that correspond to the DRS connectives already mentioned. Given the original design of the algorithm reproduced in our implementation and the notion of accessibility, the following two sentences turn out to have almost exactly the same DRS:

- (5) Every farmer who owns a donkey beats it.
 (6) If a farmer owns a donkey he beats it.

The DRS in (7) is what (5) yields. The DRS for (6) is almost identical, except that the DRS in the consequent of the conditional ⇒ contains two new reference markers in its universe *z* and *w* and the conditions *z* = *x*₁ and *w* = *x*₂.




The construction algorithm thus shows how to arrive by a completely effective means at the same DRS and the same truth conditions for (5) and (6) and in so doing solves a longstanding puzzle concerning the "donkey sentences." See Kamp (1981), Heim (1982).


So far we have seen only left-right scope assignments to sentences with indefinites and true quantificational NPs. But of course, there are other possible scope assignments, which our implementation also generates. We generate the left-right scope assignment as a default, but our implementation can generate the other scope assignments as well. Sometimes these are needed in resolving anaphoric links. For instance, consider the discourse in (8):

(8) Every man loves a woman. She is beautiful.

On the default left-right scope assignment, (8) yields:



which is a DRS that is incoherent; *z* cannot be identified with a reference marker, since neither *x*₁ nor *x*₂ are accessible to *z*. If 'a woman' takes wide scope, however, an anaphoric link is possible, and it is one that our implementation finds:



Definite noun phrases, which include definite descriptions and directly referential expressions like proper names, demonstratives, and indexicals, have a quite complex role in DR theory. All definites generate *adequacy* or *felicity* conditions for discourses containing them. When a definite NP

is used in a discourse, the discourse will be adequate if and only if the recipient has sufficient information given by the context and background assumptions to link the reference marker introduced with some contextually available individual and/or some reference marker already introduced in the discourse. This thesis is developed and defended at length in Heim (1982). If such a link is not possible, then the discourse as a whole may lack a determinate truth value. The intuition behind this requirement is that whenever the speaker correctly employs a definite NP, he has a "definite individual in mind" that the recipient must in at least some minimal way be able to isolate.

These claims lead to a special role for definites in the construction algorithm. They always have in effect wide scope over those logical operators that serve as barriers to anaphoric links, since their denotation must, if the felicity conditions are to be accommodated, already have been introduced in the discourse. To take an example, consider the case of a proper name. Insofar as a proper name is a definite, the recipient of an utterance containing one must have some information enabling him to link the reference marker introduced with the appropriate, contextually available individual. The sort of information at issue may be often quite minimal and the set of contextually available individuals quite large (it includes at least potentially everything that can be named). But once this link is made, the logical barriers to further anaphoric links between the reference markers introduced by name and those introduced by anaphoric pronouns that underlie the constraint of accessibility have no effect. That is, the reference marker introduced by a proper name is always accessible to a subsequently introduced reference marker. The reason for explaining the way definites work is to make a case for distinguishing reference markers introduced by definites from those introduced by indefinites as we have done in our implementation.

Correctness Definition

The model theory or **correctness definition** for the extensional fragment surveyed is very simple. We define a *model* for the DRS "language" to be an ordered pair $\langle D, [] \rangle$, where D is a set of entities (the "domain" of the model), and $[]$ an interpretation function that takes n -ary predicates into sets of n -tuples in $\mathcal{P}(\cup(D^n))$. Further we define an **embedding function** f for a DRS K in a model M to be a map from reference markers in U_K into the domain of M . We will also define an **extension** of an embedding function f to an embedding function g for a DRS K_1 to be the function: $g: (\text{Dom}(f) \cup U_{K_1}) \rightarrow D$. We shall abbreviate 'g extends f to an embedding of K as $f \sqsubseteq_K g$ '. We now define, again by simultaneous recursion, the notions of a **proper embedding** of a DRS K in the model M and the **satisfaction** of a condition C in a DRS K in the model M by an embedding function. I will abbreviate 'f is a proper embedding of K in M' by ' $[f, K]^M = 1$ ' and 'the model M satisfies C under an embedding function f for K' as $M \models_{f, K} C$.

Definition 2: Let x_1, \dots, x_n be reference markers, ψ an n -ary DRS predicate, K, K_1 and K_2 DRSS, and let $[\psi]^M$ be the extension of ψ in M . Then

1. If ϕ is an atomic condition of the form $\psi(x_1, \dots, x_n)$, $M \models_{f, K} \phi$ iff $\langle f(x_1), \dots, f(x_n) \rangle \in [\psi]^M$.
2. If ϕ is an atomic condition of the form $x_1 = x_2$, $M \models_{f, K} \phi$ iff $f(x_1) = f(x_2)$.

3. If $\phi \in \text{Con}_K$ is a complex condition of the form $\neg K_1$, then $M \models_{f, K} \phi$ iff $\neg \exists g \supseteq_{K_1} f [g, K_1]^M = 1$

4. If $\phi \in \text{Con}_K$ is a complex condition of the form $K_1 \vee K_2$, $M \models_{f, K} \phi$ iff $\exists g \supseteq_{K_1} f [g, K_1]^M = 1 \vee \exists g \supseteq_{K_2} f [g, K_2]^M = 1$

5. If $\phi \in \text{Con}_K$ is a complex condition of the form $K_1 \Rightarrow K_2$, then $M \models_{f, K} \phi$ iff $\forall g \supseteq_{K_1} f (([g, K_1]^M = 1 \rightarrow \exists h \supseteq_{K_2} g [h, K_2]^M = 1))$

6. $[f, K]^M = 1$ iff f is an embedding function such that: (i) $g \subseteq f$; (ii) for every condition ξ in Con_K $M \models_f \xi$; (iii) if K has an external anchor A , then $A \subseteq f$.

7. $[f, K]^M = 1$ iff $[f, K]_\Lambda^M = 1$ where Λ is the empty function.²

A DRS K is *true* in a model M just in case there is a proper embedding of K in M .

From the standpoint of the theory of information processing, a DRS represents the result of a recipient's processing of a verbal input. The DRS then as it stands captures the information that the recipient has gleaned from the sentence. The correctness definition evaluates that information content. The two components of DR theory have thus distinctly different tasks: anaphora resolution and other phenomena like scope disambiguation that are necessary for discourse understanding must take place at the level of DRS construction; the assignment of truth conditions is provided by the correctness definition. We want to emphasize that the mapping characterizing the correctness definition is not something that the recipient in general has access to or can construct. For a knowledge of such a mapping involves at least on occasion the knowledge of what are the denotations of directly referential expressions (since these provide the constraints that we have called external anchors on the embedding functions); and as almost two decades of philosophical argument have made plain, the recipient of a discourse containing directly referential expressions need not and generally does not know in any interesting sense what the denotations of those expressions are in order for the discourse to be comprehensible to him. In general, we know that we must distinguish between the truth conditional content of an expression and the content of the expression that is available to the interpreter. This is a thesis that follows directly from accepting the basic principles of the theory of direct reference, which DR theory at least to some extent incorporates. So insofar as we wish to model or to mimic the way humans process verbal inputs, then we must limit ourselves in implementations of natural language understanding to the information provided by the DRS construction algorithm.

In effect this is a useful constraint, for it forces the DR theoretician to make a principled division of labor between these two components of the theory. It also allows us to provide a much more clean cut approach to the implementation of semantic theories than is possible with theories that do not invite this sort of division. For we are able to separate tasks of discourse understanding, which, one feels, ought to be tasks that can be accomplished by means of effective algorithms, from the evaluation of a discourse at a set of indices-- which in general has no effective solution. This may be only a small and rather obvious, philosophical point, but it is one which is worth keeping in mind when one is trying to come up with tractable analyses of various aspects of natural language understanding.

The Actual Implementation : BUILDERS

The implementation we have constructed has three distinct modules-- an LFG parser, a DRS constructor, and finally an anaphoric resolution component. The inputs to BUILDERS are multisentential discourses, and the outputs are scope disambiguated semantic structures, i.e., DRSs, with reference markers introduced by anaphoric pronouns identified with the appropriate reference marker introduced by some previously processed NP.

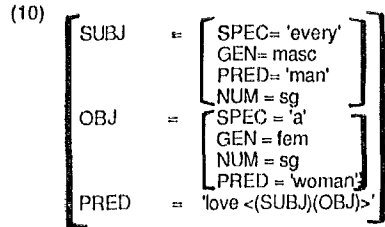
The LFG parser of the kind used in BUILDERS has already been discussed in detail by Frey and Reyle (1983), Frey (1985) and Reyle (1985). We have augmented their LFG parser to handle compound sentences (those containing conjunctions like 'and', 'or' or 'if... then'). The parser provides a separate f-structure for each of the sentential components of a compound sentence. These then serve as arguments to connectives. We also allow for the possibility of several scope assignments, when more than one operator is involved. We have also augmented the parser by attaching to each value of the specifier in an f-structure a unique identifier; in view of the fact that the LFG parser is a front end to our DRS maker, the identifiers we have chosen are new reference markers (i.e., reference markers that have not already appeared in previously processed discourse).

Once the parser has finished its task, the DRS constructor begins its work. The overall structure of the constructor is similar to that of the informal algorithm discussed earlier and also is based on the work of Frey (1985). We first translate the semantically relevant entries in SUBJ into DRS conditions and partial DRS structures, then those in OBJ and finally PRED of the f-structure as a whole (the main verb) is translated into a DRS condition. The translations for the lexical items in an f-structure are stored in a separate database and can be easily augmented as the grammar and the lexicon grow. The translations of the determiners or specifiers in an f-structure yield structures that we shall call *partial DRSs*. In the fragment treated by BUILDERS³, we deal only with singular determiners; 'a' (or its equivalents 'some' and the like) and 'every' introduce different partial DRSs. They have the form of a triplet. In the case of 'every' it is $\langle x, \lambda P \lambda Q, ([] \Rightarrow ([x], [P]), ([], [Q])) \rangle$, where x is the reference marker introduced by the specifier in the f-structure, P and Q are variables for sets of properties (i.e. sets of conditions), and the structure $([] \Rightarrow ([x], [P]), ([], [Q]))$ is the schematic form of a DRS.

In the case of 'a' it is $\langle x, \lambda P \lambda Q, ([x], [P], [Q]) \rangle$. The translations of a common noun phrase are also triplets, but they do not contain abstraction over properties or property sets. Instead, they supply the properties that are to be filled in for P and Q. So, for instance, the translation for the predicate 'man' is the triple $\langle \lambda x, 0, [man(x)] \rangle$. Again following Frey (1985), we shall call such triples *predicative DRSs*. We have a special translation for the main predicate of the sentence's f-structure; it is an ordered triplet of the form $\langle 0, 0, \beta(X, Y) \rangle$. X and Y are to be filled in eventually by the reference markers introduced by the specifiers in SUBJ and OBJ respectively. The translations of the entries in the f-structures are then combined together by means of a process called *conversion*. Conversion is like the application of a λ -abstract to an argument. Following the general path of the construction algorithm, we begin by introducing a partial DRS with the specifier of SUBJ. Suppose for example that SUBJ contains 'every' as a specifier. Its translation is the partial DRS (i) $\langle x, \lambda P \lambda Q, ([] \Rightarrow ([x], [P]), ([], [Q])) \rangle$. The common noun phrase in SUBJ (the head noun + any modifiers) yields a predicative DRS, which will contain a complex property if the common noun phrase is itself complex-- i.e. contains modifiers like possessives or relative clauses. So in general it will be of the form (ii) $\langle \lambda x, 0, (U_{CN}, CON_{CN}[x]) \rangle$, where U_{CN} is the set of reference markers and $CON_{CN}[x]$ is the set of conditions derived from the common noun phrase and where 'CON_{CN}[x]' denotes the fact that at least one condition in CON_{CN} contains 'x' as an argument. Converting (i) with (ii) yields the partial DRS $\langle x, \lambda Q, ([] \Rightarrow ([x \cup U_{CN}], [CON_{CN}[x/x]]), ([], [Q])) \rangle$. '[x/x]' denotes the replacement of every occurrence of 'x' in the conditions in CON_{CN} with the reference marker x.


The process of conversion turns to OBJ and processes it in the same way. This yields another partial DRS that, like the one above, contains a property abstract, e.g., $\langle y, \lambda Q, \phi \rangle$. The first strategy is to construct a DRS exhibiting the left-right scope of quantifiers as they occur in the discourse. So in this case, we first combine the partial DRS from OBJ with the translation of the main predicate, which is of the form $\langle 0, 0, \beta(X, Y) \rangle$. The structure $\langle y, \lambda Q, \phi \rangle$ combines with $\langle 0, 0, \beta(X, Y) \rangle$ to yield a *complete* DRS of the form $\langle 0, 0, \phi'(X) \rangle$. $\phi'(X)$ is the result of the conversion of the partial DRS ϕ with the predicative DRS $\langle 0, 0, \beta(X, y) \rangle$. Finally, conversion is applied again using the partial DRS derived from SUBJ, yielding another complete DRS and the desired result. To get alternative scope assignments for quantifiers, the program backtracks and tries to do conversion in a different way.

To take an example, let us see how the DRS constructor would handle the sentence, 'every man loves a woman.' The output of the parser yields the structure:



The DRS constructor yields the partial DRS (i) $\langle u_1, \lambda Q, ([] \Rightarrow ([u_1], [man(u_1)]), ([], [Q])) \rangle$ for SUBJ and the partial DRS (ii) $\langle u_2, \lambda Q, ([u_2], [woman(u_2), Q]) \rangle$ for OBJ. The translation of the main PRED of the whole f-structure is: (iii) loves(X,Y). Conversion of (ii) with (iii) yields the complete DRS (iv): $\langle 0, 0, [u_2] [woman(u_2), loves(X, u_2)] \rangle$, and the conversion of (i) with (iv) yields $\langle 0, 0, ([] \Rightarrow ([u_1], [man(u_1)]), ([u_2], [woman(u_2), loves(u_1, u_2)])) \rangle$, which is the desired result. To get the alternative scope reading, the successive conversion of (iii) with (i) and then with (ii) yields the DRS: $\langle 0, 0, ([u_2], [woman(u_2), ([u_1], [man(u_1)], ([], [loves(u_1, u_2)]))]) \rangle$.

This is the basic part of the DRS constructor. The two embellishments we have made to this basic part concern the treatment of relatives and possessives. The treatments of possessives and relatives are quite similar, so we shall describe here just one-- the treatment of relative clauses. The syntactic treatment of an NP containing a relative clause of the form 'DET N that S' yields following kind of f-structure:⁴



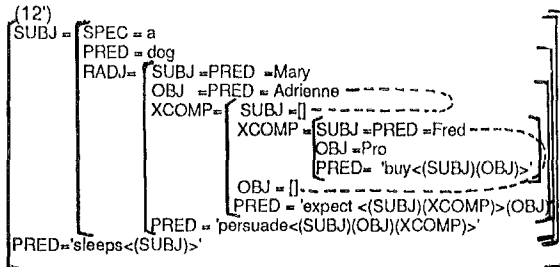
Note that pro carries an identifier identical to that of the specifier of the head.

Suppose that the translation of the specifier introduces a new reference marker x. Then the translation of 'pro' yields the following partial DRS: (i) $\langle x, \lambda P, ([], [P]) \rangle$. Since that is all that is in the SUBJ position of RADJ, we get (i) as the partial DRS associated with SUBJ. Now we use conversion to construct the partial DRS for ϕ (call this (ii)). To get the default scope

assignment we convert (ii) with the translation for β and then convert that result again with (i). The result is a *quasi-complete DRS* of the form $\langle iv \rangle \langle 0, 0, (URADJ, CONRADJ) \rangle$, where $URADJ$ is the set of reference markers and $CONRADJ$ is the set of conditions derived from the processing of $RADJ$. Quasi-complete DRSs are distinguished from complete DRSs, in that they are derived from f-structures containing 'pro'. To use (iv) in building up the partial DRS for the main NP, the program converts the structure for $RADJ$ (iv) into a λ -abstract term-- viz. $\langle v \rangle \langle 0, \lambda P, (URADJ, [CONRADJ, P]) \rangle$. We are now able to combine (v) with the translation of $PRED$ to get the translation of a complex common noun phrase-- viz. $\langle vi \rangle \langle \lambda x, 0, (URADJ, [CONRADJ, N(x)]) \rangle$. Now this predicative DRS may be combined with the partial DRS produced by the specifier in the usual way to yield the partial DRS for the complete NP.

Processing relative clauses in this way allows us to handle long distance dependencies. Consider the sentence,

(12) A dog that Mary persuaded Adrienne to expect Fred to buy sleeps.




The DRS constructor begins with the specifier of the subj f-structure in the usual fashion and then begins to construct the common noun phrase from the inside out. So then we begin with the innermost XCOMP. Suppose that the specifier introduces a reference marker x . How, we have to ask, is the translation of an identification between two f-structures to be done here? We suppose that such an identification is very similar to *pro* and that it should be translated in exactly the same way. So the translation and conversion of the innermost XCOMP yields the quasi-complete DRS $\langle 0, 0, ([a, p], [expect(a,p), p: ([f], [buy(f,x)])]) \rangle$. In a similar fashion we handle the construction of the entire $RADJ$; translation and conversion yield once again the quasi-DRS $\langle 0, 0, ([m, a, q], [persuade(m, a, q) q: ([p], [expect(a,p), p: ([f], [buy(f,x)])])]) \rangle$. Now we convert the result into a λ -abstract since we are done with $RADJ$ and we can now form the complex noun phrase translation in the familiar fashion.

Anaphora Resolution


The pronoun resolver is the last module in the program and perhaps the most complex. It operates on a complete DRS as an input and searches for conditions of the form 'pronoun(x) = *'. The reason we need to operate on complete DRSs is that the constraints on anaphoric relations imposed by the accessibility relation require that the logical structure of the sentence be determined. The scope of various quantifiers and truth functional operators, however, is only fully determined at the level of complete DRSs. Once a DRS for a discourse D has been constructed, we go back and examine it-- looking for pronouns. We also construct a database in which all the reference markers introduced in the DRS for D are stored in a tree structure, the transitive closure of which defines the accessibility relation on the universe of the DRS for D . Associated with each reference marker in this database are the gender and number of the NP that introduced it. When we

come to a reference marker introduced by a pronoun, it will occur on a certain node n_m in the tree. The reference markers accessible to it are all those on nodes n_j such that there is a path from the root to n_m passing through each n_j . At this point the program searches back to find the first available reference marker whose associated gender and number agrees with the gender and number of the pronoun, and which satisfies certain other constraints that the pronoun might have. One constraint is that a reference marker x in a DRS K introduced by a non-reflexive pronoun cannot be identified with a reference marker y , if K already contains a condition of the form $\phi(x, y, z_1, \dots, z_n)$. The opposite is true for a reference marker introduced by a reflexive pronoun. If these constraints are met, the program then replaces the condition 'pronoun(x) = *' with a condition of the form $x = y$, where y is the reference marker that was found to match x . Let us take a look at a typical example.

(13) A man loves a dog. He feeds it.



The pronoun resolver now takes over and produces a tree structure of available discourse referents as it goes through the CON list of the DRS. In this case we have a simple tree of the form $u_1, \langle \text{sing, masc} \rangle \dots u_2, \langle \text{sing, neut} \rangle \dots x_1, \langle \text{sing, masc} \rangle \dots x_2, \langle \text{sing, neut} \rangle$. As it is constructing such a tree, it also looks for reference markers introduced by pronouns. The condition 'pronoun(x₁) = *' tells it that x_1 is such a reference marker. At this point, it now searches back for the reference marker in the database with the appropriate number and gender. After rewriting the conditions, the pronoun resolver prints out the DRS:



There are, however, several problems with the resolver as described so far. One has to do with the anaphoric behavior of definites. Many English speakers find (14) acceptable but (15) bad.

- (14) If Mary likes every one who likes John_i, then she likes him_i.
- (15) *If Mary likes every one who likes someone_i, then she likes him_i.

Definites like proper names seem to be available for anaphoric linkage despite the presence of logical barriers to anaphora. But given our description of the pronoun resolver, we have not said anything that would distinguish the anaphoric behavior of definites as opposed to indefinites. We will follow the suggestion of Kamp (1983) and treat definites as having wide scope over the logical barriers to anaphora available in the present fragment.⁵ Thus, while the program constructs the accessibility tree of reference markers, it places each reference marker introduced by a proper name or other definite at the root of the tree as well as in its normal position. This creates a certain amount of duplication but allows us to get the right reading for (14) while still getting the preferred "most recent NP" readings first for the majority of anaphoric discourses. Thus, our program succeeds in finding the right DRS for (14) but predicts (15) to be bad. It also predicts that a sentence like (16) will fail, which seems marginal to at least some English speakers.

(16) ? If John likes her_j, Fred likes Mary_j.

There are still more complexities, however, to the pronoun resolver. Though presently not implemented, we see a need to distinguish the resolution strategy for pronouns that occur in subordinate clauses from the standard one. The need to do this is made evident by the apparent acceptability of Bach-Peters type sentences with indefinites like the one in (17):

(17) A man_j who hardly knows her_j loves a woman_j who scorns him_j.

The processing story that must be told to make (17) acceptable is quite complex. For instance, it cannot be that relative and other subordinate clauses are simply not processed until the main DRS is already completed. Though this processing strategy would in conjunction with our pronoun resolver predict (17) to be good, it would also predict the marginal (18) as equally acceptable:

(18) ? A man who hardly knows her_j loves Mary_j.

Pronouns in subordinate clauses seem to pose an additional complication for the pronoun resolver. We believe pronouns in subordinate clauses should be handled slightly differently than those occurring in main clauses. Like other pronouns, when the resolver finds a pronoun in a subordinate clause α it attempts to find the appropriate antecedent. But if it fails, it leaves the pronoun unresolved until it has processed the rest of the DFIS. If it finds a condition later on that is linked to the NP in which α occurs, then it will try again to find antecedents for all the unresolved pronouns in α . This sort of strategy would make (17) acceptable but not (18), which is what is desired. We hope to incorporate this into the pronoun resolver soon.

Conclusion

The pronoun resolver is an attempt to take the resolution of pronouns as far as is possible on a purely grammatical level. We realize that there are some cases where resolution depends on world knowledge and guesses about the speaker's intentions and the like. But using world knowledge is often computationally expensive. Our goal was to develop a program that would minimize appeals to world knowledge in building a semantic representation of the content of a discourse and would instead use to maximum advantage the information that can be gleaned from the words themselves.

References

- Asher, N. (1982) "Truth Conditions and Semantic Competence: Toward a Theory of Linguistic Understanding," Ph.D. dissertation, Yale University.
- Asher, N. (forthcoming) "Belief in Discourse Representation Theory," *Journal of Philosophical Logic*.
- Chierchia & M. Rooth (1984) "Configurational Notions in Discourse Representation Theory", in C. Jones & P. Sells, eds. *Proceedings of N.E.L.S 14, Amherst GLSA (1984)* pp. 49-63.
- Fenstad, J.E., P.K. Halvorsen, T. Langholm and J. van Benthem (1985) "Equations, Schemata and Situations: A framework for linguistic semantics", *CSLI Report, Stanford University*.
- Frey, W. (1985) "Syntax and Semantics of Some Noun Phrases", in Laubsch (ed): *Proceedings of the German Workshop on Artificial Intelligence 1984*.
- Frey, W. & H. Kamp (1985) "Distributive and Collective Plurals", Talk Presented at University of Texas Conference on Syntax and Semantics, March 1985.
- Frey, W., and U. Reyle, (1983) "A PROLOG Implementation of Lexical Functional Grammar as a Base for a Natural Language Processing System." in *Proceedings of the First Meeting of the ACL. Europe*.
- Halvorsen, P. (1983) "Semantics for Lexical Functional Grammar" in *Linguistic Inquiry* Vol. 14, Num. 4. MIT Press.
- Heim, I. (1982) *The semantics of Definite and Indefinite Noun Phrases*, Ph.D. dissertation, Univ. of Mass.
- Johnson M. & E. Klein (1985) "A Declarative Form of DR Theory," talk presented at Summer Association of Symbolic Logic & Computational Logic Meeting, Stanford CA, 1985.
- Kamp, H (1985) "Context, Thought, and Communication," in *Proceedings of the Aristotelian Society* 85, 239-261.
- Kamp, H. (1981) "A Theory of Truth and Semantic Representation", in Groenendijk, J., Janssen, Th. & Stokhof, M., eds., *Formal Methods in the Study of Language*, Mathematisch Centrum Tracts, Amsterdam.
- Reyle, U. (1985) "Grammatical Functions, Quantification and Discourse Referents." in *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*. 829-831.

NOTES

¹We should note that there are other ways of defining accessibility. One might be tempted to do it wholly in terms of the possibility of assigning the DRS containing the identification coherent truth conditions. Chierchia and Rooth (1984) investigate this possibility. We do not find their approach computationally useful, however, insofar as checking for truth conditional consistency is any more expensive than following the accessibility constraints on DRSs.

²The truth conditions for attitude reports are too complex to give here. For details see Asher (forthcoming), Kamp (1985).

³We have plans to expand the constructor to handle plurals as well.

⁴The other alternative of course is that the pro occur in the object position. Our algorithm handles that similarly.

⁵The question is more delicate within the context of propositional attitude verbs. There it seems not all definites function in the same way. Some permit anaphoric linkage across propositional attitude contexts and some do not.