

ANAPHORA AND LOGICAL FORM:
ON FORMAL MEANING REPRESENTATIONS
FOR NATURAL LANGUAGE

Bonnie Nash-Webber
Bolt Beranek and Newman Inc.
Cambridge MA 02138

Raymond Reiter
Department of Computer Science
University of British Columbia
Vancouver, B.C., Canada

Abstract

We argue, on favor of representations We then adopt "function" for representation, identifying the of anaphoric suggests c properties of a facilitate the possible antecedents appropriate re with such pro possible to de rich class of a

1. to provide a computational approach to certain problems in anaphora in natural language;
2. to argue in favor of formal meaning representation languages (MRLs) for natural language.

These two objectives are not independent. It appears that the solutions to certain problems in anaphora are best formulated with respect to an appropriately structured logical MRL, so that the structural entities out of which such an MRL is composed suggest possible antecedents for anaphor resolution.

More specifically, we have set ourselves the following problem: what form should a meaning representation assume in order to facilitate the identification of possible antecedents of anaphoric expressions, and what computational mechanisms does this task require? Moreover, we have chosen to investigate this problem of identifying a set of possible antecedents without

invoking general world knowledge. The separate issue of choosing the most appropriate antecedent from this set will, in general, require plausible reasoning based on such general world knowledge. We are also aware of instances where such knowledge is required even to propose possible antecedents. Nevertheless," In this paper, our concern is to explore the implications of a purely syntactic approach as well as to ascertain its limitations. It turns out that a surprisingly rich class of anaphora, both pronouns and ellipses, is amenable to such an approach, provided that an appropriately structured logical MRL is used. We shall find that the use of such an MRL leads to particularly simple rules for identifying possible antecedents, and that the structure of the MRL can be exploited computationally to preclude certain inappropriate ones. We shall also find that this task of identifying possible referents is intimately bound up with an ability to form appropriate descriptions of them, and that these descriptions are, in turn, intimately related to logical form.

11. Why Logical Meaning Representations?

Although there is universal agreement within the AI community that natural language understanding systems must provide some underlying meaning representation onto which surface strings are mapped, the nature of this representation remains a contentious issue. One aspect of this debate has to do with the form that this representation should take. There appear to be two points of view: logical forms [e.g., Sandewall, 1971; Woods, et al, 1972] and structured networks [e.g., Wilks, 1975; Schank, 1975; Simmons, 1973],

Despite this lack of any formal distinction between networks and logical forms, there is a widespread bias within the AI community against logical MRLs for natural language. [See for example, Charniak and Wilks, 1976]. We suspect

that there are two implicit assumptions underlying this anti-formal point of view:

1. that the choice of a logical form necessarily implies a commitment to a corresponding proof theory as one's sole computation mechanism;
2. that logical forms must have their "natural" representation at the implementation level, e.g., that $(x)(Ey) . Px,y \ \& \ Qx,y$ must be represented by the S-expression $((X)(E Y)(AND (P X Y)(Q X Y)))$.

Neither of these assumptions is justified. We have already observed that networks can be best viewed as implementation level representations for logical forms, and as we shall show in Section IV., the computations that we propose for anaphor resolution within a logical MRL are in no way based on any kind of proof theory.

If it is the case then that logical forms and networks are one and the same, why prefer one over the other? We favor logical form on the following grounds:

A. Semantics

By virtue of its being a formal language, a logical form inherits a well defined semantics, namely, its Tarskian semantics. This is not the case for network representations presented in vacuo (i.e., without a translation mechanism mapping the network to a logic). As Woods [1975] points out, networks often fall short of this requirement.

B. Computation

The high level conceptual representation of meaning provided by logical forms encourages the formulation of appropriate processing algorithms at an equally high conceptual level, independent of how these logical forms are represented at the implementation level. This provides for perspicuous descriptions of algorithms, without specifying the irrelevant, CONS cell level, pointer chasing details required by network representations. The examples of Section IV illustrate the ease with which such rules can be formulated, as well as their conceptual clarity.

There are two issues here: representational perspicuity and representational adequacy. The first is largely a subjective matter. We believe logical forms to be more readable and comprehensible than their corresponding network forms, especially when the usual

network primitives are considerably augmented in order to correctly represent logical connectives and quantifiers and their scopes [Schubert, 1975; Hendrix, 1975].

The second issue representational adequacy - is far more important, largely dealing with the ability of a given MKL to express the meaning of surface strings. A related issue is that of representational closure. Can one tell, from the given specification of an MRL, what can and what cannot be expressed within it? Because any logical MRL has both a well defined syntax and a well defined semantics it necessarily exhibits a high degree of closure. This is not the case for network representations presented in vacuo, precisely because they have no semantics. Many of the network based meaning representations in current natural language systems [Schank, 1975; Wilks, 1975; Norman and Rumelhart, 1975] suffer from this defect, a fact that makes it extremely difficult to assess their content.

It is instructive in connection with the above discussion to note that these very same issues were hotly debated within the data base management community during the early 1970's. There too, the basic choice was between a network view of data [CODASYL, 1971] and a logical or so-called relational view [Codd, 1970]. Moreover, the arguments advanced in favor of the relational view were in many ways isomorphic to those we have made favoring logical form for meaning representations. At least within the data base community, the logical view currently prevails, primarily because its high conceptual level provides so-called "data independence," i.e., one's view of the data is independent of implementation details.

III. On Appropriate Logical MRLs

In the previous section, we argued on fairly general grounds in favor of formal meaning representation languages for natural language. Of course, not just any logical MRL will do. At the very least, any such formal language must provide for quantification and the usual logical connectives, but even under these requirements there remains a broad spectrum of possible logical representations. There are at least two dimensions to this spectrum corresponding to representational level and representational structure. With respect to level, representations in current systems range from very "surface" [Simmons, 1970] to very "deep", primitive-based ones [Wilks, 1975;

Schank, 1975]. Wilks, in Computational Semantics [Charniak and Wilks, 1976; 1761 provides a good discussion of these issues. In this paper, we use as illustration an MRL that keeps very close to the surface syntax and lexicon. We do so because the computational task that we have taken as a "forcing function" for an adequate representation, namely, identifying possible antecedents for anaphor resolution, seems not to require a deep level and is moreover facilitated by a "surfacy" one, at least for a broad and interesting class of phenomena.

Our focus in discussing logical MRLs is on their representational structure. We have found that the need to provide appropriate antecedents for anaphor resolution suggests certain structural constraints on possible MRLs which greatly facilitate this process. In this connection, we emphasize that we are not here proposing a fully developed logical MRL. To do so would require, at the very least, adequate representations for tense, modality, mass terms, events, etc. issues which we have so far completely ignored. The MRL used in this paper is merely a vehicle for displaying certain formal structural properties which we have found necessary for the identification of antecedents. Our belief is that any fully articulated logical MRL will have to provide these structural units if it is to deal effectively with anaphora. Accordingly, one way of viewing our proposal is as a set of design constraints on the structure of possible logical MRLs for natural language. The remainder of this section deals with these structural properties.

A* Lambda-expressions

For a formal MRL to be adequate for the resolution of verb phrase ellipsis, it must provide for constructions equivalent to lambda-expressions. (This has also been observed by several linguists recently, including Sag [1976] and Williams [1977].) For example, the sentence pair

- 1a. John loves Mary.
- 2. So does Bill.

requires, as the antecedent of the ellipsed verb phrase, the formal construct Mx [Love x, Mary] corresponding to "loving Mary", whence the resolved sentence 1b. becomes

Bill, $\>s(x)$ [Love x, Mary]

which simplifies to

Love Bill, Mary

(Note that our preferred notation for applying a lambda-expression to an argument is to follow the argument by the lambda expression, corresponding to normal subject-predicate word order in English.)

B. Separation of Descriptive and Assertional Information: ~Types

Since the antecedents of many anaphoric expressions are descriptions, an adequate formalism must be so organized that these descriptions stand out clearly. For example, consider the pair of sentences

- 2a. Some cotton T-shirts are expensive,
- b. but not the one Mary gave John.

In a "flat" predicate calculus MRL (ignoring the distinction between "some" plural and "some" singular), sentence 2a. might be represented by

(Ex) . Cotton x & T-shirt x
& Expensive x

Now intuitively, the antecedent of "one" in sentence 2b. is something like "cotton T-shirt", but from the flat predicate calculus representation, there is no more reason to suppose that Cotton and T-shirt form a possible antecedent than Cotton and Expensive, or T-shirt and Expensive, or any one or all three. That is, there is no structural indication that Cotton T-shirt is a referenceable unit. We believe such an indication is necessary in any formalism adequate for anaphor resolution.

Using the structure of a typed logic, predicates that constrain the range of a quantified variable - i.e., types - (like T-shirt here) can be structurally distinguished from predicates that assert things (as "Expensive" does here). Moreover, using the lambda operator, the notion of type can be extended from simple one-place predicates to more complex ones to yield all and only the allowable referenceable entities. For example, we can represent

"T-shirt"
as T-shirt
"cotton T-shirt"
as $X(u:T\text{-shirt})[\text{Cotton } u]$
"T-shirt that Mary gave Fred"
as $\>,(u:T\text{-shirt})[\text{Gave Mary, Fred, } u]$

(The first is merely a shorthand for $\>,(u:T\text{-shirt})[\text{True}]$.) Notice that we are postulating a representation for "cotton T-shirt" that is more highly structured than a simple conjunction of Cotton and T-shirt, i.e. (T-shirt x & Cotton x). Specifically, we are separating that part of the noun phrase denoting the primary

class to which an entity belongs (usually the head noun) from those parts denoting restrictions on that class (conveyed by adjectives and relative clauses).

This provides yet another structural property that a logical MRL should possess in order to facilitate the identification of antecedents for anaphor resolution. Consider

3. Mary bought a tie-dyed cotton T-shirt and Fred bought an embroidered one.

Whether intuitively "one" refers to "cotton T-shirt" or "tie-dyed cotton T-shirt" or merely "T-shirt", it must refer at least to "T-shirt", the primary class denoted by the noun phrase. There is no way (pronominally) in English to refer to a restriction without also referring to the primary class. This is our main reason for keeping them distinct in our logical MRL.

Another consequence of this separation of descriptive from assertional information is that it avoids problems that Woods [1975] discusses with respect to adequate representations for relative clauses. First, sentences like "A dog that had rabies bit a man" and "A dog that bit a man had rabies" can be assigned distinct representations, for example

(Ex: X(u:Dog) [Have u, Rabies]) (Ey:Man)
 . Bit x, y
 (Ex: MurDog) [(Ey:Man) . Bit u, y]
 . Have x, Rabies

More importantly, processing rules such as those proposed in Section IV, can treat these two representations differently. As Woods point out, conventional semantic networks fail to represent the distinction between these two sentences. From the perspective of the above discussion, one reason for this is clear: a conventional network is a representation, at the implementation level, of a "flat" predicate logic.

IV* Examples

In this section, we present several fragments of discourse, each containing anaphoric expressions - pronouns and/or ellipses. Recall that we are not concerned here with the kinds of external knowledge needed to choose among possible antecedents for an anaphoric expression. What we are concerned with is insuring that

1. in cases where the antecedent of an anaphoric expression is not explicitly present, it can often be derived through purely syntactic

manipulations of an appropriately structured MRL;

2. the properties we have proposed for a logical MRL make such manipulations simple to express and apply.

Since developing our approach to anaphora and logical form, we have discovered that it is compatible with one major line of thought in transformational linguistics today (For a survey of current ideas on anaphora in linguistics and psychology, see [Nash-Webber, 1977].) That line of thought is illustrated both by Chomsky's recent remark that the "general principles of anaphora apply to logical forms rather than to surface structures directly" [Chomsky, 1975: 241 ft. 31], as well as by recent research on logical theories of verb phrase deletion and gapping [Sag, 1976; Williams, 1977].

With our examples, we give a small set of manipulation rules which yield the needed antecedents. We make no claims for the completeness of these rules; there obviously remains a great deal of work to be done along these lines (see Section VI). We do believe, however, that the examples indicate the utility of our basic approach, and that this approach provides a promising direction for further research.

A. Implicit Sets

Our first example illustrates one way of deriving a set as a candidate antecedent for "they". Consider the sentences

- 4a. Mary gave each boy a T-shirt,
- b. She bought tjiem at Filene's.

The first may be represented as

4c. (Vx:Boy)(Ey:T-shirt)
 . Gave Mary,x,y

(For simplicity, we will ignore the fact that "each boy" is probably anaphoric, referring to each boy in some previously mentioned set or one implicitly defined by context, and treat it rather as a universally quantified noun phrase.)

Notice that we are considering each sentence individually since we want to assign it a representation that is correct, but which does not depend on what may follow. The result will often be a reading that is in some sense noncommittal: it will be vague but true. If subsequently we learn more about the situation, we will refine this representation to reflect our new knowledge state.

The second sentence we represent initially with its anaphoric elements overtly marked, that is,

4d. Bought SHE^ THEY , Filene's

(We subscript the pronoun symbols merely to keep several instances of the same one distinct, as would be the case in "They thanked her for them".) Next, we identify possible referents for the anaphoric terms. Since Mary is the only female around, we trivially assign her as the referent of SHE.

Regarding candidate antecedents for THEY, we postulate two ways of deriving possible sets from sentences like 4a.

1. Form the set description of any type restricting a universally quantified variable. (toe represent the set description of type C by {x|Cx}.)
2. Let to be a (prior) formula not containing the anaphoric element THEY, nor any negation in the main clause. (In the current example, 4c. plays the role of W.) Suppose to has an existentially quantified variable y that lies within the scope of a universally quantified variable. Form the set description of the set of y's satisfying W. This is a straightforward procedure, involving the type of y restricted by an expression deriving from W. Details are given in [Nash-Webber, forthcoming], but the example should suggest its basic outline.

Thus, sentence 4a. yields {x|Boy x}, the set of boys, via the first procedure, and {v|T-shirt v & (Ew:Boy) . Gave Mary,w,v} via the second one, i.e. the set of T-shirts, for each of which there is some boy to whom Mary gave it. Substituting each of these sets in turn for THEY^, yields

- 4e. Bought Mary, {x|Boy x}, Filene's
- f. Bought Mary, (v|T-shirt v & (Ew:Boy) . Gave Mary,w,v}, Filene's

That is, either Mary bought all the boys at Filene's or she bought there all the T-shirts she gave out. Real world knowledge would now be needed to choose the more plausible reading.

Notice that in English the pronoun "they", as well as many plural noun phrases, are ambiguous between a collective reading ("all together") and a distributive one ("taken one at a time"). Sometimes, a lexical item will indicate that a plural should be understood distributively, as "each" does in "Mary's suitcases were each weighed at the

airport". Sometimes, semantic **selectional restrictions** can be used to choose between the two. For example, "pile" requires a collective interpretation of its object: "She piled them into a heap" **doesn't mean** "for each one, she piled it into a heap". But often, only unknown aspects of the situation can furnish the appropriate information. For example, if **we learn** that "Mary's suitcases were weighed at the airport", **we cannot say for certain** whether each one was weighed there separately or just the whole lot together.

In line then with the policy described above of always opting for a vague **but** true interpretation, rather than making unsubstantiated choices, we will interpret plurals non-distributively, unless or until additional information would lead us to an alternate decision.

B. Type Antecedents

Our next example illustrates the identification of descriptions as candidate antecedents for anaphoric "one".

- 5a. Mary gave each boy a green T-shirt,
- b. She gave Sue a red one.

toe interpret sentence 5a. like 4a. above, except for the additional modifier "green" on T-shirt.

(tfx:Boy) (Ey: X(u:T-shirt) [Green u])
 . Gave Mary, x, y

Sentence 5b. can be represented initially as

(Ez: A(u:P?) [Red u]) . Gave **SHE₁Sue,z**

That is, there is something of unknown type P? that should be derivable from context, which we know explicitly is red, which some known female SHE, gave Sue. Our task is now to identify possible antecedents for SHE, and P?.

There are simple syntactic criteria for rejecting Sue as an antecedent **for** SHE,. So again by default, there being no other females around, we assign Mary as **the** referent for SHE,.

As for P?, its possible antecedents include all "recently" mentioned types, independent of the particular quantifiers. ("Recent" seems to mean here **the** current sentence, **the** previous **one**, and **perhaps** the one before that. It does **not seem** to be affected by **task structure** [Deutsch, 1975] or **story structure**, or any of the other **factors** that **seem to change the set of available antecedents for definite pronouns**, "he", "it", etc.)

The types explicitly given in example 5 are: Mu:T-shirt [Green u], T-shirt and Boy. Notice that when one type is constructed out of other types via the lambda operator, we include them all as possible candidate antecedents. Prescribing exactly what criteria one would use to identify the most plausible antecedent for P?, or in what way one would apply them, is not within the scope of this paper. But they would include the semantic criterion that one be able to predicate Red of an entity of type P? This would eliminate v(u:T-shirt) (Green u) through application of a "clashing color" axiom: if something is green, it is not red. (Notice that if sentence 5b. had been "Fred, she gave a large one." there would be no reason to eliminate this description as a plausible antecedent.) Under rhetorical criteria, we would expect parallelism to argue for plausibility. That is, if two successive sentences are structurally similar ("parallel") and in the latter, anaphoric "one" helps to fill role R (here, the object), then it has a very plausible antecedent in the noun phrase filling role R in the previous sentence (here, the previous object "a green T-shirt"). But our point here is not to specify procedures for choosing among candidate antecedents; it is rather to show how a suitable logical framework provides in a straightforward way all and only the appropriate possibilities.

c* Predicate Antecedents

1• Simple Verb Phrase Deletion

The next few examples illustrate some problems involving verb phrase ellipsis, which are handled rather neatly within our framework.

- 6a. Mary gave Sue a T-shirt,
- b. Jane did too.

The representation that we assign to sentence 6a. is

$(\text{Ex:T-shirt}) . \text{Gave Mary, Sue, x}$

Sentence 6b., we interpret as predicating something (P?) of Jane that had previously been predicated of someone else:

P? Jane

To identify possible antecedents for P?, we find the one-place predicates that either are given explicitly or can be derived via lambda abstraction on the subject position. (Again, one probably need only search for such predicates in the current sentence if it has several clauses or in the one or two sentences

immediately preceding it, as the half-life of predicate antecedents, like that of type antecedents, seems to be very short. Note that we are viewing the first argument place of a predicate as corresponding to surface subject position. Though this requires a different representation for active and passive sentences, we see the need for this on other grounds, for example, their difference with respect to simple verb phrase deletion:

John hit a linguist
Fred did too.

as opposed to

John was hit by a linguist.
Fred was too.
*Fred~~dTd too.

This example is simple in that there is only one such one-place predicate abstractable off a subject:

$X(r)[(\text{Ex:T-shirt}) . \text{Gave } r, \text{ Sue, } x]$

that is, giving Sue a T-shirt. Substituting for P? yields

Jane, $X(r)I(\text{Ex:T-shirt}) . \text{Gave } r, \text{ Sue, } xJ$

which is equivalent to

$(\text{Ex:T-shirt}) . \text{Gave Jane, Sue } x$

Note that this representation does not commit us to both girls having given Sue the same T-shirt, nor need they be different. The description of the first one is

$\text{flz:T-shirt } z \ \& \ \text{Gave Mary, Sue, } z$

"a T-shirt that Mary gave Sue", where \wedge indicates the indefinite operator. (z might be called in English "the T-shirt which Mary gave Sue" if no other T-shirt in the discourse meets this description.) The second T-shirt is describable as

$\text{Vlw: T-shirt } w \ \& \ \text{Gave Jane, Sue } w$

"a T-shirt that Jane gave Sue".

It is important to be able to derive such descriptions, since the entities they describe may serve as antecedents for later anaphoric expressions, for example,

6c. Neither of them fit her.

where "them" refers to the implicit set of T-shirts given to Sue, who is also the most plausible antecedent of "her".

2* sloppy identity

Our next example illustrates a phenomenon that has been called the "sloppy identity problem" [Ross, 1967]. It involves accounting for the appearance of an additional reading for sentences containing deleted verb phrases. That is, while sentence 7a seems unambiguous, sentence 7b might mean either that Fred beats Garth's wife or that he beats his own. How do we account for this?

- 7a. Garth beats his wife,
- b. Fred does too.

We assign 7a. an initial representation in which its anaphoric term is overtly marked.

- 7c. Beat Garth, 'S(Wife)HE₁

(where 'S is defined to be a function that takes a unary predicate like Wife, School, etc., and returns a function like "wife-of", "school-of", etc. 'S(Wife), for example, is a function that takes a man as its argument and returns his wife: 'S(Wife)John is John's wife. Having a function like 'S eliminates the need to postulate a separate "Y-of" function for every unary predicate Y.)

With no other male around, we can assign HE₁ to Garth by default, that is,

- 7d. Beat Garth, 'S(Wife)Garth

We assign sentence 7b. the representation

- 7e. P? Fred

Now, while there are no explicit one-place predicates around to serve as an antecedent for P?, there are two ways in which to abstract one from 7d.

- (i) >(r)[Beat r, 'S(Wife)Garth]
- (ii) y(r) [Beat r, 'S(Wife)r]

The first represents beating Garth's wife and the second, beating one's own. Substituting for P?, we get the two plausible readings

- Fred, >(r)[Beat r, 'S(Wife)Garth]
- Fred, y(r)[Beat r, 'S(Wife)r]

That is, either Fred beats Garth's wife or he beats his own.

(While we have not done so here, it is actually important to distinguish whether an argument place was originally filled by an anaphoric expression or by a full noun phrase. Observe that if sentence 7a. had read "Garth beats Garth's wife", which would be represented simply as 7d., the

following sentence, "Fred does too", could only mean that Fred beats Garth's wife. Thus in an actual implementation, we would have to indicate "he = Garth", rather than simply replacing "he" with "Garth", in order to derive all and only the correct lambda abstractions.)

3' Abstracting Predicates from Non-Subject Position

The point of the next example is to illustrate abstracting one-place predicates from positions other than the surface subject. In English, "likewise", "too", "similarly", etc., followed optionally by a preposition and then by a noun phrase indicate that the role filled by the new noun phrase in a previously mentioned predicate is a non-subject one. (As illustrated earlier, a noun phrase followed by an auxiliary requires that the noun phrase fills the subject role of a previous predicate.)

- 3a. John hit a cop.
- b. Likewise, a CIA agent.

The first sentence may be represented as

- 8c. (Ex-Cop) . Hit John, x

and the second one as

- 8d. (Ey:CIA-agent) . Q? y

where Q? stands for an anaphoric predicate like P?, but one whose argument fills a non-subject role.

To resolve Q?, we must identify the one-place predicates that can be abstracted from non-subject positions. From 8c., we get

- >(r)[Hit John, r]

which, substituted for Q? in 8d., yields

- (Ey:CIA-agent) . y, >,(r)[Hit John, r]

or more simply,

- (Ey:CIA-agent) . Hit John, y

that is, "Likewise, John hit a CIA agent."

It might appear that one could resolve "likewise"-ellipses at the level of the surface string alone, but this is not the case. Consider the following example:

- 9a. John gave Sally her present,
- b. Likewise, Fred.

Obviously, while substituting "Fred" for "Sally" in the surface string would yield an interpretable sentence, "John gave Fred

her present", this is not the only, nor the most plausible reading of the ellipsis in 9b.

We represent 9a. as

9c. Gave John, Sally, 'S(Present) SHE

which we interpret as

9d. Gave John, Sally, 'S(Present)Sally

she being the only female around. Sentence 9b., we represent as

9e. $\exists?$ Fred

To resolve $\exists?$, we identify the one-place predicates that may be abstracted from non-subject positions. From 9d., we get

- i. $\lambda(r)[\text{Gave John, } r, \text{'S(Present)Sally}]$
- ii. $\lambda(r)[\text{Gave John, } r, \text{'S(Present)r}]$
- iii. $\lambda(r)[\text{Gave John, Sally, } r]$

substituting for $\exists?$ and flattening for clarity, we get

-Gave John, Fred, 'S(Present)Sally
"Likewise, John gave Fred Sally's present."

-Gave John, Fred, 'S(Present)Fred
"Likewise, John gave Fred Fred's present."

-Gave John, Sally, Fred
"Likewise, John gave Sally Fred."

Again, the preferred interpretation would be chosen by using world knowledge.

C. "Donkeys"

As our final example of how an appropriate logical representation of a sentence can yield antecedents necessary for anaphor resolution, we will consider a particularly bothersome class of sentences, illustrated by example 10.

10. Every man who owns a donkey beats it.

The problem lies in identifying the antecedent of "it". It is not "a donkey". The sentence does not mean that every man who owns a donkey beats a donkey, but rather that he beats any donkey that he owns. Moreover, there is no way of construing the existential quantifier associated with "a donkey" such that "it" falls within its scope. How does the correct antecedent for "it" emerge from our framework?

We first assign sentence 10 the interpretation

$$(\forall x:\lambda(u:\text{Man})[(\exists y:\text{Donkey}) . \text{Own } u, y]) . \text{Beat } x, IT_1$$

That is, for every man for whom there is some donkey that he owns, he beats it. Now while there is nothing explicit to serve as the antecedent for "it", it turns out that "it" can also reference a certain kind of functional entity which arises from existentials.

We postulate the following rule for identifying a possible antecedent for IT.

1. Find a type restriction which contains an existentially quantified variable not within the scope of either a universal quantifier or negation.
2. Determine the description of y with respect to this type restriction: any entity which satisfies this description is a possible antecedent for IT. (Again, we omit the specification of the rule for determining y 's description, although one should be clear from the example.)

For i., there is one such type restriction - $[(\exists y:\text{Donkey}) . \text{Own } u, y]$. The description of the existentially quantified y is

ii. $\lambda(u)[\eta y: \text{Donkey } y \ \& \ \text{Own } u, y]$

That is, it is a function which, given a u , returns a donkey that u owns if u owns a donkey. For a given x then,

$\eta y: \text{Donkey } y \ \& \ \text{Own } x, y$

is a donkey that x owns, Substituting into i. yields

iii. $(\forall x:\lambda(u:\text{Man})[(\exists y:\text{Donkey}) \text{Own } u, y]) . \text{Beat } x, \eta y: \text{Donkey } y \ \& \ \text{Own } x, y$

Notice that this rule is independent of how the type containing the existential has been quantified. Thus, in

11. Some man who owns a donkey beats it.
12. Which man who owns a donkey beats It?

the antecedent of "it" is the donkey obtained by applying function (ii) to the quantified variable associated with "man".

* Discussion

The examples of Section IV were designed to illustrate the feasibility of deriving possible antecedents for anaphoric expressions directly from an appropriately structured logical representation. Notice that basic to this representation is an adequate indication of the scope of logical operators -

quantifiers, conjunction and negation - for otherwise, we could not deal correctly with antecedents arising from existentials (e.g., the examples in Sections IV.A and IV.D). Also basic is the recognition and correct attachment of modifiers - relative clauses, prepositional phrases, pronominal modifiers, etc. - necessary for correctly handling "one" anaphora (e.g., Section IV.B). Taken together, these impose the requirement of a pre-processor for mapping surface strings onto logical forms at least as powerful as that of the LUNAR system [Woods et al, 1972]. It follows that much of the burden of antecedent identification is actually being placed upon this pre-processor, given the need for an appropriate logical form before our approach can be applied.

It should also be clear that what we are describing in this paper is essentially a competence model for anaphor resolution. In its crudest implementation, one would first generate a set of possible antecedents, and then test each of these by plausible reasoning using general world knowledge. Of course, we are not seriously proposing such a generate and test implementation. There are a variety of heuristics that can be invoked to aid the choice of a most probable antecedent and any performance model must make use of such knowledge. (Heuristics for assigning antecedents have been proposed throughout the linguistics, psychology and AI literatures. See [Nash-Webber, 1977] for a discussion of many of them.) Nevertheless, even a performance model must have the ability to determine the space of possible alternatives from which such heuristics are to make their choice. Since some of these alternatives may not be present explicitly, it is here that the approach of this paper becomes relevant.

VI. Further Problems

As this paper is necessarily brief, we do not have the space to discuss at length such interesting issues as the effects of negation or various opaque contexts on the kinds of antecedents evoked. These are discussed in [Nash-Webber, forthcoming]. We will, however, mention one such issue - the problem of existential noun phrases in negative contexts.

The scope of negation is inherently ambiguous, and as with quantifiers, different scope interpretations yield different antecedents. Moreover, some interpretations may yield no antecedent at all. For example, we know that in a positive context, an existential noun phrase will always result in an entity which can be described in terms of that

context. So, if John married a Swedish girl, "she" can refer to the Swedish girl that John married. However, in a negative context, "she" may refer to other entities depending on how the scope of negation is interpreted.

Negated Vejr

- 13a. John didn't marry a Swedish girl,
b. He lived with her for three years.

"she" = the Swedish girl John was involved with

Negated Modifier

- 14a. John didn't marry a Swedish girl,
b. She was from Denmark.

"she" = the girl John married

Negated Main Desc

- 15a. John didn't marry a Swedish girl,
b. She was at least 15 years HTs senior.

"she" = the Swedish female John married

Negated Proposition

- 16a. John didn't catch a trout,
b. *He ate it for dinner.

Again, we would want to postulate a neutral initial representation for negation, one that might be vague, but would nevertheless be true. Only when we were required to e.g., in order to resolve an anaphoric expression - would we then attempt to make a commitment to the scope of negation, (Note that a belief context poses much the same problem as negation, i.e. that of determining the scope of belief. For example, in

- 17a. John thought he married a Swedish girl,
b. but she was really from Denmark.

as in Example 14 above, "she" is the girl that John married. Here again, it is only the modifier "Swedish" that should be taken as falling within the scope of belief.)

As we mentioned earlier, before one can fix on a particular MRL, one must provide for mass concepts; for tense; for quantifiers other than universals and existentials; for facts, events, states or acts; and for generics, among other things, as the following examples illustrate.

- 18a. When John spills beer,
 b. his dog licks it: up.

"it" = the bit of beer that John spills

19. John drinks beer because it tastes good.

"it" = beer

- 20a. ^{b*} ^{tnev} Mary linguists smoke, although
 know it causes cancer.

"they" = the linguists who smoke

"it" = smoking

- 21a. Few linguists smoke, since
 b. they know it causes cancer.

"they" = linguists

- 22a. A beagle smiled at me yesterday,
 b. They are very friendly dogs.

"they" = the generic class of beagles

- 23a. John dunked Mary's braids in the
 inkwell.
 b. it made her cry.

"it" = the event of John's dunking
 Mary's braids in the inkwell

- 24a. John dunked Mary's braids in the
 inkwell.
 b. He didn't regret doing it.

"it" = the act of dunking Mary's
 braids in the inkwell

Finally, although we have indicated the need for plausible inference for choosing an appropriate candidate from a set of possible antecedents (e.g., Section IV.B.), it is also the case that such inferring may be needed to derive possible antecedents. That is, not all possible antecedents are structurally derivable.

- 25a. Yesterday I saw a couple in the park.
 b. He was wearing shorts and she had on a dashiki.

Clearly, what is required is some sort of general knowledge of the form: "A couple usually consists of two individuals, one male and one female."

Although we can see no a priori reasons why a formal approach could not accommodate the use of plausible reasoning in the derivation of possible antecedents, we have chosen not to explore these issues in this paper. Rather, our intention in this work is to first determine just how far an essentially syntactic approach can be pushed.

Notice that our treatment of all of the examples of Section IV has a decidedly syntactic character: descriptions of those entities proposed as possible antecedents are either explicitly present in some formula of the MRL, or can be derived from such a formula by appropriate local operations on its structure, independent of the availability of general world knowledge. While there are cases where world knowledge is required even to derive possible antecedents for pronouns, we believe we have shown that the vast majority can be derived by purely syntactic considerations and that world knowledge is used (perhaps heuristically) to select from among these. The determination of possible antecedents based on such purely syntactic considerations and the formulation of design constraints on MRLs to facilitate this process best describe the objectives of this paper.

Acknowledgements

This research was supported by the National Institute of Education under Grant MS-NIE-C-400-76-0116 and by the National Research Council of Canada under Grant A-7642. The authors wish to thank Dr. Alan Mackworth and Dr. William Woods for reading and suggesting improvements to earlier drafts of this paper.

References

- [1] Charniak, E. and J Wilks (eds.) (1976) Computational Semantics. North-Holland Publishing Co., Amsterdam.
- [2] Chomsky, N. (1975) Reflections on Language. Pantheon Books, New York.
- [3] CODASYL (1971) Data Base Task Group of CODASYL Programming Language Committee Report, ACM Publications, New York.
- [5] Codd, T. (1970) "A Relational Model of Data for Large Shared Data Banks," CACM .L3, No. 6.
- [4] Deutsch, B. (1975) "Establishing Context in Task Oriented Dialogs," American Journal of Computational Linguistics' No.~4
- [6] Hendrix, G. (1975) "Expanding the Utility of Semantic Networks through Partitioning," Advanced Papers of 4IJCAI, Tbilisi, USSR, 3-8 September, 1975.

- [7] Nash-Webber, B.L. (1977)
 "Anaphora: A Cross-disciplinary Survey",
 Technical Report CSR-31. Center for the
 Study of Reading, University of Illinois
 and Bolt Beranek and Newman, Inc.
- [8] Nash-Webber, B.L. (forthcoming)
 Forthcoming doctoral dissertation,
 Department of Applied Mathematics, Harvard
 University.
- [9] Norman, D. & D. Rumelhart (1975)
Explorations in Cognition, W.H. Freeman &
 Company, San Francisco.
- [10] Ross, J.R. (1967)
 "Constraints on Variables in Syntax,"
 unpublished doctoral dissertation, MIT
 Department of Foreign Literatures and
 Linguistics.
- [11] Sag, I. (1976)
 "Deletion and Logical Form," unpublished
 doctoral dissertation, MIT Department of
 Foreign Literatures and Linguistics.
- [12] Sandewall, E.J. (1971)
 "Representing Natural Language Information
 in Predicate Calculus," in Machine
 Intelligence 6, Edinburgh University
 Press.
- [13] Schank, R. (ed.) (1975)
Conceptual Information Processing,
 North-Holland Publishing Co., Amsterdam.
- [14] Schubert, L. (1976)
 "Extending the Expressive Power of
 Semantic Networks," Artificial
 Intelligence 7(2), pp. 163-198.
- [15] Simmons, R. (1973)
 "Semantic Networks: Computation and Use
 for Understanding English Sentences," in
Computer Models of Thought and Language,
 Schaffner and Colby (eds.), W.H. Freeman &
 Co., San Francisco.
- [16] Simmons, R.F. and B.C. Bruce (1971)
 "Some Relations between Predicate Calculus
 and Semantic Net Representations of
 Discourse," in Proc. 2nd IJCAI, London.
- [17] Wilks, Y. (1975)
 "A Preferential, Pattern-seeking Semantics
 for Natural Language," Artificial
 Intelligence 6, pp. 53-74.
- [18] Williams, E. (1977)
 "Discourse and Logical Form", Linguistic
 Inquiry 8(1) pp. 101-140.
- [19] Woods, W.A. (1975)
 "What's in a Link: Foundations for
 Semantic Networks," in Bobrow and Collins
 (eds.), Representation and Understanding:
 Studies in Cognitive Science. New York:
 Academic Press.
- [20] Woods, W.A., R.M. Kaplan and
 B.L. Nash-Webber (1972)
 "The Lunar Sciences Natural Language
 Information System: Final Report." BBN
 Report 2378, Bolt Beranek and Newman Inc.,
 Cambridge MA.