

CONTEXTUAL REFERENCE RESOLUTION

DAVID KLAPPHOLZ AND ABE LOCKMAN

*Department of Electrical Engineering  
and Computer Science  
Columbia University  
New York, New York 10027*

ABSTRACT

With the exception of pronomial reference, little, has been written (in the field of computational linguistics) about the phenomenon of reference in natural language. This paper investigates the power and use of reference in natural language, and the problems involved in its resolution. An algorithm is sketched for accomplishing reference resolution using a notion of cross-sentential focus, a mechanism for hypothesizing all possible contextual references, and a judgment mechanism for discriminating among the hypotheses.

### The reference resolution problem

The present work began as an attempt to develop a set of algorithms and/or heuristics to enable a primitive-based, inference-driven model of a natural language user (Schank 1972 Rieger 1974) to properly resolve pronomial references across sentence boundaries. The authors quickly realized, however, that the problem of pronomial reference resolution is only a small aspect of a problem which might be termed nominal reference resolution, itself but a small aspect of the problem of the coherence of a text, (or conversation) i. e. the manner in which it "means" more than the logical conjunction of the meanings of its individual constituent sentences.

Examples of the first problem, i. e. pronomial reference resolution are given in sentence sequences 1-4 below.

1. Yesterday some boys from our village chased a pack of wild dogs; the largest one fell into a ditch.
2. The wild dogs which forage just outside our village suffer from a strange bone-weakening disease. Yesterday some boys from our village chased a pack of wild dogs; the largest one broke a leg and fell into a ditch.
3. Yesterday John chased Bill half a block; he was soon out of breath.
4. My friend Bill has an extremely severe case of asthma. Yesterday John chased Bill half a block; he was soon out of breath.

The problem in utterance (text, conversation etc.) excerpts of the above type is that of determining the referents of the various occurrences

of the pronouns "one," and "he"

For the moment we simply note that usually preferred referents of the two occurrences of "one" are "boy" and "dog", (examples 1 and 2 respectively) and those of the two occurrences of "he" are "John" and Bill (examples 3 and 4 respectively.)

The more general problem of nominal reference resolution is exhibited in the following annotated excerpt from a recent newspaper article (N. Y. Times 7/15/75, byline Arnold Lubasch); subscripted bracketing of the excerpt is intended only to enable later reference to specific parts of the text.

<sub>1</sub>[Some of the major provisions of <sub>2</sub>[the state's Fair Campaign Code]<sub>2</sub> <sub>1</sub> were declared unconstitutional here yesterday by <sub>3</sub>[a special Federal court]<sub>3</sub> that assailed <sub>4</sub>[the restrictions on election campaigning] as "repugnant to the right of freedom of speech."

<sub>5</sub>[The three-judge court, <sub>5</sub> which was convened to consider a constitutional challenge by three State Assembly candidates last year threw out <sub>6</sub>[<sub>7</sub>[the 'code's]<sub>7</sub> prohibition against attacking any political candidate's race, sex, religion or ethnic background]<sub>6</sub>

<sub>8</sub>[It]<sub>8</sub> also overturned <sub>9</sub>[<sub>10</sub>[<sub>11</sub>[the code's]<sub>11</sub> ban]<sub>10</sub> on any misrepresentation of a candidate's party affiliation, position on political issues and personal qualifications, including the use of "character defamation" and scurrilous attacks." ]<sub>9</sub>

According to <sub>12</sub>[the court's]<sub>12</sub> 38-page decision, written by <sub>13</sub>[Judge Henry F. Werber]<sub>13</sub> with the concurrence of <sub>14</sub>[Judges Leonard P. Moore, and Mark A. Constantino]<sub>14</sub>. <sub>15</sub>[<sub>16</sub>[the provisions]<sub>16</sub> banning misrepresentation]<sub>15</sub> "cast a substantial chill on the expression of protected speech that are unconstitutionally overbroad and vague."

If newspaper reporters had a bit more sympathy for those of us concerned with natural language processing, the above excerpt might have read as follows:

The state has a Fair Campaign Code.

Some of the major provisions of the state's Fair Campaign Code are provisions which restrict something.

Some of the things restricted by some of the major provisions of the state's Fair Campaign Code which restrict something are activities having to do with election campaigning.

Some of the activities having to do with election campaigning which are restricted by some of the major provisions of the state's Fair Campaign Code which restrict something are attacking a political candidate's race, sex, religions or ethnic background and misrepresenting a candidate's party affiliation, position on political issues ...

Last year three state assembly candidates filed a constitutional challenge to some of the major provisions of the state's Fair Campaign Code which restrict something.

Yesterday a special Federal court declared unconstitutional those of the major provisions of the state's Fair Campaign Code which restrict something ...

The point is that in order for a machine or a human to validly claim to have "understood" the original excerpt he/she/it must be able at the very least to demonstrate that he/she/it has established the following relationships between various items occurring in the excerpt.

(Integers represent subscripted bracketed regments of the original excerpt.)

(i) The identity of 2, 7, and 11

(ii) The identity of 3, 5, 8, and 12

(iii) The fact that 4, 6, 9, and 15 are elements, subsets or parts of 1

(iv) The fact that 13 and 14 are members of 3

and on and on and on. (I. e. a closer analysis of the original excerpt reveals many more relationships which must be established before "understanding" may be claimed.)

If people actually wrote/spoke in the style of the somewhat facetious paraphrase of the original excerpt, the nominal reference problem would be reduced to one of matching lexical patterns and recognizing a few syntactic cues; to state the obvious, the necessity for more succinct linguistic communication has forced the development of elliptical devices which shift the burden of nominal reference resolution from syntactic analysis to an analysis of the "semantics" of sentences in context. More specifically, nominal references cannot in general be resolved without the use of general semantic information as well as specific world knowledge.

While the fact that syntactic analysis alone is insufficient for understanding is anything but novel, the question of the magnitude of the nominal reference problem and of its solution's crucial dependence upon

local context seems to have been little commented upon. (Clark (1975) discusses the problem from a viewpoint different from that of this paper.)

The reader who remains unconvinced by the examples above that local context (and specific world knowledge relating to local context) must play a crucial role in reference resolution is asked to consider the two sentence sequences 5a, 6, and 5b, 6.

5. a. The founding fathers had a difficult time agreeing on how the basic laws governing our country should be framed.
- b. Those foolish people at the country club have spent an incredible amount of time arguing about club rules.
6. The second article of the constitution, for example, was argued about for months before agreement was reached.

In sentence sequence 5a, 6, "the second article" clearly refers to the second article of the constitution of the United States, while in sentence sequence 5b, 6, the reference is to the second article of the constitution of the country club. In each case the only factor involved in resolving the reference is the semantic content of its local context- in this case the meaning of the sentence preceding the one in which the reference occurs.

Since the lexical item "the constitution" appears in the example just considered, a word concerning such proper-noun-like objects is in order. In any language there are lexical items and phrases such as those appearing in 7 below, which, in the absence of compelling alternative, have standard default referents; for example the standard

default referents of the items in 7 are the corresponding items in 8 below.

7. a. The constitution
  - b. The founding fathers
  - c. Wall Street
  - d. The establishment
  - e. The president
  - f. Madison Avenue
8. a. The constitution of the U.S.
  - b. The founding fathers of the U.S.
  - c. The U.S. business community (or that part of it residing in New York City.)
  - d. Those people who have the power to influence the course of events in the nation etc. etc.
  - e. The president of the U.S.
  - f. The advertising industry.

In order for textual occurrence of such proper-noun-like objects to be properly handled, their standard default referents must be listed in the lexicon. This is not to say that occurrences of proper-noun-like objects cannot be references to objects occurring previously in the text; rather it is the case that their default options must also be considered as possible referents.

As final examples of the reference resolution problem let us consider sentence sequences 9 and 10 below.

9. The president was shot while riding in a motorcade down one of the major boulevards of Dallas yesterday; it caused a panic on Wall Street.
10. John was invited to tea at the Quimby's last Saturday; he would have loved to go, but he knew he'd be busy then.

In example 9, while the first sentence of the sequence contains a number of noun objects (president, motorcade, boulevards, Dallas) which are potential referents for the occurrence of "it" in the second sentence, none of these is in fact, the proper referent; rather, the proper referent of "it" is the event (or fact) that "The president was shot while ... ."

In example 10 we have an instance of an adverbial reference ("then") which must be recognized as referring to "yesterday" rather than to some non adverbial object occurring in the first sentence of that example.

#### Sketch of a Solution

From the point of view of computer implementation, the problem of nominal reference resolution is one of creating tokens for noun objects mentioned in a text, and discovering and encoding the relations, alluded to in the text, which hold between them and various other tokens in memory.

This problem, though certainly not its magnitude or ramifications, was noticed by Rieger (1974) in his pioneering implementation of a primitive-based model of a natural language user. Rieger's system, however, suffers from the incredible inefficiency resulting from its need to search all of memory in order to attempt any reference resolution; in addition it will often miss a quite obvious referent entirely, and, in fact, resolves non-pronomial references only accidentally if at all.



Before presenting a sketch of a proposed solution to the nominal reference resolution problem, it would be well to detail more precisely the overall language processing environment within which it is meant to operate and of which it is a most necessary part.

First, we assume that a relatively small set,  $S$ , of semantic primitives and a logical-calculus-like language,  $L$ , for expressing "meanings" are available. The set  $S$  and language  $L$  must satisfy the following two conditions.

- (i) The predicate, function, and constant symbols of  $L$  are members of  $S$ .
- (ii) There is a one-to-one mapping from meanings of (natural language) sentences to formulas of  $L$ .

While a set of primitives and a meaning representation language even demonstrably close to satisfying the above conditions have yet to be produced, we will, in examples to follow, make use of meaning representations; the only claim we will make for them is that the functions served by their constituent constructs must be served by the elements of any adequate system.

In addition to a meaning representation scheme we will assume an encoding of world knowledge of the sort which a "typical" adult might possess, again with the same obvious caveat.

While the question of translation from natural language sentences to 'meaning representations will not be touched upon here, we will assume sentence-by-sentence translation of the sort exhibited in various

examples to follow.

The solution to the reference resolution problem rests in recognizing the fact that reference is an elliptical device, and that the human understander of natural language cannot recapture that which was elided once he is too far from it in the text; in fact, he cannot resolve a reference to a point in the text more than a few-sentences back without going back and pondering it (if he can do so at all). We should note that this is true even in the case in which the referent doesn't actually appear in the text, but appears only in an inference from some statement made in the text. In this latter case - a case which we will discuss only at the very end of this paper the reference is not resolvable (and would not therefore have been made by the creator of the text in the first place), unless the statement from which the inference is made appears shortly before in the text. Though we cannot say precisely how far back is meant by "shortly before," it is certainly no more than a few sentences. For a given sentence, S, appearing in a text we will refer to the sequence of sentences preceding S by no more than the intended distance as the focus of S.

In terms of computer implementation, we will, in the processing of a text (which we conceive of as proceeding sentence-by-sentence), maintain the following focus sets.

- (i) The noun-object focus - the set of tokens of all noun <sup>objects in The</sup> meaning re-  
presentations of the focus of S (where S is the sentence currently being processed)

- (ii) The event focus - a set containing, for every sentence  $W$  in the focus of  $S$ , the object  $EVENT(F)$ , where  $F$  is the meaning representation of  $W$ , and  $EVENT$  is a function which maps the meaning of a formula,  $F$ , into a noun-like object whose meaning is "the event (or fact) that  $F$ "
- (iii) The time focus - a set containing takens for all time references (e.g. yesterday, five o'clock, etc.) occurring in the meaning representation of the focus of  $S$ .

The reader may question our inclusion of every object appearing in the meaning representation of the focus of  $S$  in one of the above focus sets, i. e. in the set of potential referents. In fact, however, it seems to be the case that any object (of one of the above-mentioned types) occurring in the meaning representation of the focus of  $S$  may be the referent of an object occurring later in the meaning representation of  $S$ . Consider, for example, the sentence sequences formed by taking each of the sentences of 12 below - in turn - as an immediate continuation of a text containing sentence 11 below.

11. Stan argued with his sister Fran in an attempt to convince her that she should bring Mary, whom he would like to get to know, on their planned trip to the San Diego Zoo tomorrow.
12. a. He was really insistent.  
 b. She was hard to convince.  
 c. It was useless.  
 d. He thinks she's the prettiest one of all Fran's friends.  
 e. The prospect really excites him.  
 f. He argued that it wouldn't tie Mary up for more than half a day.  
 g. It's the best one in the country, you know.

- h. She thought it was a terrible idea.
- i. She happened to be busy then, but expressed an interest in coming along another time.

Each of the underlined items in sentences 12a-12i references some object in sentence 11. (For the sake of clarity we present in 13 below the referents as we understand them.)

- 13. a. Stan
- b. Fran
- c. The attempt (to convince . . .)
- d. Mary
- e. EVENT (Stan will get to know Mary)
- f. The trip
- g. The San Diego Zoo
- h. Both she and it are ambiguous; if she is taken to be "Fran," then it refers to EVENT (Fran will bring Mary . . .); if she is taken to be "Mary", then it refers to EVENT (Mary will come. . .)
- i. Tomorrow

The point is, of course, that any item in (the meaning representation of) a sentence, S, may be referenced by some item in (the meaning representation of) a latter sentence.

On the other side of the coin the question of identifying potential references is just as important as that of identifying the set of all possible referents for an object which is known to reference something. If we were concerned only with pronomial reference resolution, the problem would have a simple solution; every pronoun is a reference. For nominal items other than pronouns the problem is far less simple;

if a noun occurs in a text just how do we know if there is a previously occurring nominal item to which it refers? As much as we would like there to be algorithmically testable criteria, i. e. recognizable syntactic and/or semantic cues, for making the decision, there seem to be none.

Thus, the mechanism we propose considers every token appearing in the translation of a sentence as a possible reference.

At present, we hypothesize the existence of a small set, R, of relations which are sufficient to account for all instances of nominal reference. Included in this set are, at the very least, the relations identity, member of, subset of, and part of. Note that although this list of relations is quite small, it suffices to handle all the examples of reference presented thus far (i. e. those occurring in sentence sequences 1-6 and 9-12 as well as those occurring in the excerpted newspaper article above).

All of the above observations taken together lead to the following sketch of an algorithm for reference resolution.

- I. As each new sentence, S, is translated into its meaning representation, the various focus sets (noun-object, event, time) are updated.
- II. A set, H, is formed containing all tuples of the form  $(N_1, N_2, \rho)$  such that  $N_1$  is a nominal item occurring in (the meaning representation of S,  $N_2$  is an object occurring in the focus set (noun-object, event, or time) appropriate to  $N_1$ , and  $\rho$  is a member of R; H is the set of all current reference hypotheses arising from S.
- III. A "judgment mechanism," discussed below, is invoked to determine the likelihoods of the correctness of the various members of H.

It is clear that following step II any further processing of reference hypotheses requires that all members of H be considered relative to one another, since the correctness or incorrectness of one may depend crucially upon that of others. In the general case not all hypotheses will turn out to be correct, and in fact some may contradict others - for instance in the case of two hypothesis-triples with identical first and second elements and different third elements.

Once it has been created, the set H is submitted to a "judgment mechanism" whose task it is to choose some of the hypotheses as valid and others as invalid. The judgement mechanism must clearly have access to the world knowledge stored in memory, and must be capable of performing inferencing of a sort which produces decisions as to the relative likelihoods of the various hypotheses.

Before giving examples of just how such a judgment mechanism might work, we should make it clear that our sense of "inferencing" is very different from Rieger's (1974). In Rieger's sense inferencing is undirected, while ours is directed toward the goal of validating hypotheses. There is, in addition, another sense in which the sort of inferencing to be done by the judgment mechanism is directed. The fact that the reasons for validating or throwing out a particular reference hypothesis (on the part of human natural language users) involve the information conveyed in local context as well as world knowledge relating to items contained in that information (and world knowledge relating to items contained in world

knowledge relating to items contained in that information, etc.) constitutes a good guess as to the particular pieces of world knowledge and the rules of inference which must be involved in judging that hypothesis.

#### Examples of reference resolution

14 and 15 below contain components of possible meaning representations of the two sentences of sentence sequence 1 at the beginning of this paper.

14. C1: CHASED ( $x_1, x_2$ )  
 C2: TIME (C1, YESTERDAY)  
 C3: SUBSET ( $x_1$ , [BOYS])  
 C4: SUBSET ( $x_2$ , [DOGS])  
 C5: GREATER (SIZE ( $x_1$ ), 1)  
 C6: GREATER (SIZE ( $x_2$ ), 1)

15. C7: FALL INTO ( $y_1, y_2$ )  
 C8: TIME (C7, PAST)  
 C9: MEMBER ( $y_2$ , [DITCH])  
 C10: MEMBER ( $y_1, y_3$ )  
 C11: LARGEST ( $y_1, y_3$ )

The meaning representations proposed for the two sentences are

$C_1 \wedge C_2 \wedge C_3 \wedge C_4 \wedge C_5 \wedge C_6$  and  $C_7 \wedge C_8 \wedge C_9 \wedge C_{10} \wedge C_{11}$  respectively. Note that we are not claiming that the predicates CHASED, and FALL INTO and the constants YESTERDAY, BOY, DOG, PAST and DITCH are at the level of semantic primitives; rather, the above analyses are at just the level which we need to illustrate the operation of the reference resolution mechanism. Furthermore, the symbols YESTERDAY, BOY, DOG, PAST and DITCH should be taken as pointers to the definitions of the appropriate items encoded in memory in whatever fashion. The bracketing in the notation [A], where A is a pointer to a definition, is meant to be a function which takes A


into an object whose meaning is the class of items satisfying the meaning pointed to by A.

Once the translation of the first sentence of sequence 1 into its meaning representation has been completed - on the assumption that that sentence is at the beginning of the text being processed - the various focus sets will contain the following:

noun object focus:  $\{x_1, x_2\}$ ; event focus:  $\{(C_1 \wedge C_2 \wedge C_3 \wedge C_4 \wedge C_5 \wedge C_6) \}$ ;

time focus {YESTERDAY}.

After the second sentence is translated the set, H, of reference - triple hypotheses presented to the judgment mechanism will then be the following:



Note that no member of the event focus occurs in H because the translation of the second sentence contains no term of the form EVENT(y); for simplicity we omit the question of time referencing.

All of the relations between  $y_2$  and  $x_1$  or  $x_2$  can be ruled out on the basis of SUBSET ( $x_2$ , [DOG]) SUBSET ( $x_1$ , [BOY]), MEMBER ( $y_2$ , [DITCH]) and of the world knowledge to the effect that boys/dogs cannot be identical to, members of, subsets of or parts of ditches (of course in some weird fairy tale setting one of these might be possible and shouldn't be thrown out; but in such a case local context would inform us of the "weird" situation and the appropriate one wouldn't be thrown out.)

The hypothesis that  $y_1$  or  $y_3$  is a part of either  $x_1$  or  $x_2$  can be



ruled out on the basis of SUBSET ( $x_1$ , [BOY]) and SUBSET ( $x_2$ , [DOG]), which tell us that  $x_1$  and  $x_2$  are sets of objects, and the world knowledge that sets don't have "parts" in the sense of the "part of" relation.

Identify between  $y_1$  and either  $x_1$  or  $x_2$  can be ruled out on the basis of MEMBER ( $y_1$ ,  $y_3$ ) which tells us that  $y_1$  is an individual and SUBSET ( $x_1$ , [BOY]), SUBSET ( $x_2$ , [DOG]), GREATER (SIZE ( $x_1$ ), 1), and GREATER (SIZE ( $x_2$ ), 1), which tell us that  $x_1$  and  $x_2$  are sets containing more than one object. (Remember that we're not doing axiomatic set theory in which there are no "individuals" in our sense and in which the sort of "individual" which is dealt with can be a subset of some set.)

Finally, the "member of" relation between  $y_3$  and either  $x_1$  or  $x_2$  can be ruled out on the basis of MEMBER ( $y_1$ ,  $y_3$ ) which requires that  $y_3$  be a set, SUBSET ( $x_1$ , [BOY]), SUBSET ( $x_2$ , [DOG]), GREATER (SIZE ( $x_1$ ), 1), and GREATER (SIZE ( $x_2$ ), 1), which tell us that  $x_1$  and  $x_2$  are sets containing more than one element each, and the fact that sets are not members of sets. (Again, we're not dealing with set theory; if in fact, we were talking about axiomatic set theory in English, then local context would contain that information, and different inferences would come into play.)

This leaves us with the following hypotheses :

$$y_3 \left\{ \begin{array}{l} \text{is identical to} \\ \text{is a subset of} \end{array} \right\} \left\{ \begin{array}{l} x_1 \\ x_2 \end{array} \right\}$$

$$y_1 \text{ is a member of } \left\{ \begin{array}{l} x_1 \\ x_2 \end{array} \right\}$$

But some of these hypotheses are consistent with one another; in fact the hypotheses

$$y_3 \left\{ \begin{array}{l} \text{is identical to} \\ \text{is a subset of} \end{array} \right\} x_i \quad i = 1, 2$$

imply the hypotheses

$$y_1 \text{ is a member of } x_i \quad i = 1, 2$$

respectively because of MEMBER ( $y_1, y_3$ ). At any rate, the judgment mechanism assumes at this point that either  $y_1$  is a member of  $x_1$  or  $y_1$  is a member of  $x_2$ . The reader is asked to recall at this point that in presenting the usually preferred referents for references in sentence sequences 1-4 the claim was made that in sentence sequence 1, the usually preferred referent for "one" is "boys." The reason for this claim is the author's observation. that, when such a pronomial reference occurs as the surface subject of a sentence, in the absence of semantic content which discriminates among the various possible referents, most people seem to take the surface subject of the last sentence in the focus as the intended referent. The reason for this human judgment is probably that the reader/hearer takes the surface subject to be the "topic" of a sentence. If this observation is correct, the judgment mechanism should, in the current example, simply choose "one of the boys" ( $y_1$  is a member

of  $x_1$ ) as the proper referent. If this observation is incorrect, the judgment mechanism should judge that there is ambiguity in the reference "one "

Sentence sequence 2 at the beginning of this paper would be handled in precisely the same manner as sentence sequence 1 up to the point at which " $y_3$  is a member of  $x_1$ " and " $y_3$  is a member of  $x_2$ " were the remaining hypotheses. The knowledge that "the dogs" referred to suffer from a strange bone-weakening disease would then cause the judgment mechanism to strengthen the likelihood that "one" refers to "dogs," thus causing " $y_1$  is a member of  $x_2$ " to be the preferred judgment.

Sentence sequence 16 below contains an example of EVENT reference.

16. The president was shot yesterday. It caused a panic on Wall Street. Omitting all other details of the translation into meaning representation we simply note that the primitive-level predicate into which "cause" is translated requires an object of the form EVENT (F) as its subject (i. e. if we say something like "John caused a stir" what we mean is that John did something and the event (or fact) that he did that caused a stir.) Thus, when the 2nd sentence is handled, the only possible referents for "it" will be the objects contained in the EVENT focus, namely just EVENT (the president was shot yesterday). The judgment mechanism thus must simply decide if the event (or fact) that the president was shot yesterday was likely to have caused a panic on Wall Street, a judgment which, with adequate world knowledge, should certainly be confirmed.

Sentence sequence 17 is a very similar case.

17. The president was shot yesterday. Bill told me all about it. It caused a panic on Wall Street.

In order to resolve the reference "it" in the last sentence of 17, the judgment mechanism would have to decide on the relative likelihoods of i and ii below

- (i) The event (or fact) that the president was shot yesterday caused a panic on Wall Street.
- (ii) The event (or fact) that Bill told me about the president being shot yesterday caused a panic on Wall Street.

Again, with the availability of reasonable world knowledge about such things as presidents, their being shot and panics, the judgment mechanism should be able to choose the proper referent for "it"

While a fully detailed specification of the judgment mechanism must await further investigation, the above examples should illustrate, at least in part, the manner in which we conceive of its operation.

### Conclusions

The phenomenon with which we have been dealing is one example of what we would like to call the "creative" aspect of language use; more specifically, reference of the sort we have described - and attempted to handle - is an elliptical device necessary for effective communication; moreover, it is a device which exhibits the ability of language to "change the ground rules" in a very flexible and fluid manner in response to context.

At this point we must admit that there is an even more creative type of reference than the sort we have dealt with. 18 below is an example of this type of reference.

18. Last week I caught a cold while visiting my mother in Chicago; as usual, the chicken soup had too much pepper in it.

The interesting reference in the above example is "chicken soup." There is no item in the first sentence to which it is directly related; on the other hand, few people have any trouble resolving it by interpolating between the two sentences of example 18 the idea expressed in sentence 19 below:

19. When I get sick my mother makes me chicken soup.

If sentence 19 were available, our reference resolution mechanism would easily come up with an identity relation between the two occurrences of "chicken soup." Obviously, for our proposed mechanism to resolve this reference, some sort of inferencing must first work on the 1st sentence of 18 to produce the meaning of 19 as an inference. Thus it is clear that reference resolution and general inferencing must be interleaved.

The mechanism proposed above does not handle the entire problem. It does, however, seem to be a minimal model of reference resolution (minimal in the sense that at least this much must be going on). In addition, it provides for that control over the use of general inferencing which is required to avoid a combinatorial explosion (BOOM).

#### References

Clark, H.H. (1975), "Bridging", Conference on Theoretical Issues in Natural Language Processing, 10-13 June 1975, Cambridge, Mass.

Rieger, C.J. (1974), Conceptual Memory: A Theory and Computer Program for Processing the Meaning Content of Natural Language Utterances, Ph.D. Thesis, Stanford University, 1974.

Schank, R. (1972), "Conceptual Dependency: A Theory of Natural Language Understanding," Cognitive Psychology 3(4), 1972.