AD A056838

Bolt Beranek and Newman Inc.

bbn

Report No. 3878

LEVEL II

12

Research in Natural Language Understanding

Quarterly Progress Report No. 3, 1 March 1978 to 31 May 1978

July 1978

Prepared for:
Advanced Research Projects Agency

Quarterly Progress rept. no. 3,
1 Mar-31 May 78,

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

1. REPORT NUMBER
BBN Report No. 3878

2. GOVT ACCESSION NO.

3. RECIPIENT'S CATALOG NUMBER

4. TITLE (and Subtitle)
RESEARCH IN NATURAL LANGUAGE UNDERSTANDING
~~Quarterly Technical Progress Report No. 3~~
1 March 1978 to 31 May 1978

5. TYPE OF REPORT & PERIOD COVERED
Quarterly Technical Progress Report

6. PERFORMING ORG. REPORT NUMBER
Report No. 3878

7. AUTHOR(s)
B. L. Webber [redacted] R. Bobrow

8. CONTRACT OR GRANT NUMBER(s)
N00014-77-C-0378

9. PERFORMING ORGANIZATION NAME AND ADDRESS
Bolt Beranek and Newman Inc.
50 Moulton Street
Cambridge, MA 02138

10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS

11. CONTROLLING OFFICE NAME AND ADDRESS
Office of Naval Research
Department of the Navy
Arlington, VA 22217

12. REPORT DATE
June 1978

13. NUMBER OF PAGES
58

14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)

15. SECURITY CLASS. (of this report)
Unclassified

15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)
Distribution of this document is unlimited. It may be released to the Clearinghouse, Department of Commerce, for sale to the general public.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)
Anaphora, ATN, augmented transition networks, case grammars, deterministic parsing, discourse, discourse models, natural language understanding, parsing, pronominal reference, reference, semantic grammars, syntactic processing, syntactic-semantic interactions.

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report is the third quarterly progress report of the ARPA-sponsored Natural Language Understanding project at BBN. The goals of the project are to develop techniques required for fluent and effective communication between a decision maker and an intelligent computerized display system in the context of complex decision tasks such as military command and control. This problem is approached as a natural language understanding problem, since most of the techniques required (cont'd)

DD FORM 1473, 1 JAN 73 EDITION OF 1 NOV 65 IS OBSOLETE

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

20. Abstract (cont'd.)

would still be necessary for an artificial language designed specifically for the task. Characteristics that are considered important for such communication are the ability for the user to omit details that can be inferred by the system and to express requests in a form that "comes naturally" without extensive forethought or problem solving. These characteristics lead to the necessity for a language structure that mirrors the user's conceptual model of the task and the equivalents of anaphoric reference, ellipsis, and context-dependent interpretation of requests. These in turn lead to requirements for handling large data bases of general world knowledge to support the necessary inferences. The project is seeking to develop techniques for representing and using real world knowledge in this context, and for combining it efficiently with syntactic and semantic knowledge. This report discusses aspects of our research to date and a general approach to definite anaphoric reference and near-deterministic parsing strategies.

RESEARCH IN NATURAL LANGUAGE UNDERSTANDING

Quarterly Technical Progress Report No. 3

1 March 1978 to 31 May 1978

ARPA Order No. 3414

Contract No. N00014-77-C-0378

Program Code No. 8D30

Contract Expiration Date:
31 August 1978

Name of Contractor:
Bolt Beranek and Newman Inc.

Short Title of Work:
Natural Language Understanding

Effective Date of Contract:
1 September 1977

Principal Investigator:
Dr. William A. Woods
(617) 491-1850 x361

Amount of Contract:
$301,377

Scientific Officer:
Gordon D. Goldstein

Sponsored by
Advanced Research Projects Agency
ARPA Order No. 3414

This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by ONR under Contract No. N00014-77-C-0378.

TABLE OF CONTENTS

	page
Introduction	1
Description Formation and Discourse Model Synthesis	2
1. Introduction	2
2. Indefinite Noun Phrases and Discourse Entities	6
3. Two Interesting Reference Problems	15
3.1 Parameterized Individuals	15
3.2 Disjunction	22
4. Conclusion	24
The RUS System	28
1. Introduction	28
2. Background	29
2.1 The LUNAR Approach	29
2.2 Semantic Grammars	31
2.3 Initial Development of RUS	33
3. Incremental Parsing	34
3.1 Structural Descriptions	34
3.2 Incremental vs. Wholistic Processes	36
3.3 The Basic Operation of RUS	38
3.4 Illustration	41
4. Improvements to Increase Determinism and General Efficiency	45
Abstracts	53
References	56

Introduction

The current BBN ARPA project in Natural Language Understanding has as its primary goal the discovery and development of techniques for dealing with large bodies of information in the kinds of complex decision-making situations that arise in military command and control. In our previous two reports, we discussed our progress in developing appropriate representational conventions for storing natural conceptual information in a machine and algorithms for using that information efficiently. In this report we discuss two aspects of our research aimed towards the creation of a fluent and natural communication interface to such information:

1. The development of a general and effective approach to definite anaphora
2. The design of near-deterministic parsing strategies which can be transported easily to new domains.

Description Formation and Discourse Model Synthesis

Bonnie Lynn Webber

1. Introduction

Previously in Woods & Brachman [1977] we stated that fluent communication with a machine in decision-making tasks can only take place if a speaker has the ability to refer back to any and all the entities introduced into the discourse, as well as any additional ones derivable from them. Because this set of entities is far larger than one might suppose (cf. Webber [1978a]), a speaker would find it awkward, if not impossible, to have to recognize and name upon introduction every entity that s/he may subsequently want to refer to. English however - as well as other natural languages - provides a simple facility for allowing subsequent references without resorting to naming - that is, **definite anaphora.** Expressions which can function as definite anaphors include definite pronouns and definite descriptions. The problem has been to understand how definite anaphora works.

Recently many researchers in linguistics, psychology, philosophy and artificial intelligence have begun to abandon a purely linguistic approach to definite anaphora in favor of a notion of reference into some kind of **model** of the discourse, cf. Karttunen [1976], Levin & Goldman [1978], Lyons [1978]. Stenning [1975]. The research on definite anaphora funded under the current contract (cf. Webber [1978a&b]) follows this latter notion, in particular making the following five assumptions:

1. One objective of discourse is to enable a speaker to communicate to a listener a **model** s/he has of some situation. Thus the ensuing discourse is, on one level, an attempt by the speaker to direct the listener in synthesizing a similar model.

2. Such a **discourse model** can be viewed as a structured collection of entities, organized by the roles they fill with respect to one another, the relations they participate in, etc.

3. The function of a **definite anaphoric expression** is to refer to an entity in the speaker's discourse model (DM{S}). In using a definite anapho., the speaker assumes (a) that on the basis of the discourse thus far, a **similar** entity will be in the listener's model (DM{L}) as well and (b) that the listener will be able to access that entity via the given definite description or definite pronoun.

4. The **referent** of a definite anaphor is thus an entity in DM{S}, which the speaker presumes to have a counterpart in DM{L}. Discourse entities may have the properties of individuals, sets, events, actions, states, facts, beliefs, hypotheses, properties, generic classes, typical set members, stuff, specific quantities of stuff, etc.

5. In deciding which discourse entity a definite anaphor refers to, a listener's judgments stem in part from how the entities in DM{L} are **described**. (When a discourse entity **E** is the referent of a definite anaphor **A**, one might distinguish that description of **E** conveyed to the listener by the immediately preceding text and consider it **A**'s **antecedent**.)

The point of making these assumptions explicit is to stress that insofar as reasoning about discourse entities is mediated by their descriptions, discourse entity descriptions are critical to anaphor resolution.

Now one consequence of these assumptions about discourse models and reference is that the task of understanding definite anaphora can be decomposed into several complementary parts:

1. deciding whether a definite pronoun or definite description is truly anaphoric (i.e., is intended to refer to some entity presumed to already be in DM{L}) or whether the term fills some other role in the discourse;

2. synthesizing a discourse model which is similar to that of the speaker and inhabited by similar discourse entities;

3. constraining the possible referents of a given anaphoric expression down to one possible choice - the "anaphor resolution" problem;

4. determining what other functions a definite description is intended to fill besides enabling the listener to construct or get to its referent.

This note is intended to illustrate briefly how the explicit data - i.e., the actual sentences of the discourse, produced by a particular person (or a particular computer program) in a particular situation - provide material for the model synthesis process. In particular, I shall show (1) how indefinite noun phrases are associated with the evocation of new discourse entities, independently of any higher-level expectations, and

(2) how those new discourse entities will initially be described. I will claim that such an initial description (**ID**) is critical to both model synthesis and anaphor resolution since it allows the listener to reason appropriately about the discourse entity in order to assign it to an appropriate role vis-a-vis his or her higher-level expectations. <*1> Moreover, since it is possible for a discourse entity's current role assignment to be found incorrect, it is the entity's **ID** that allows it to be re-assigned to another role with respect to the listener's revised expectations.

In Section 2 I will consider indefinite noun phrases vis-a-vis the discourse entities they evoke and how those entities are described. I will contrast them briefly with non-anaphoric definite noun phrases and then show that all determined noun phrases, including odd ones like "few orc eggs", "many lemon gum balls", etc. pattern after either definites or indefinites vis-a-vis the discourse entities they evoke and how those entities can be described. In Section 3 I will show how this approach to definite anaphora in terms of discourse entities and their descriptions can accommodate certain problematic cases

---

<*1>. From different points of view, discussions of the relationship between the explicit text and higher-level organizing structures can be found in Collins, Brown & Larkin [1977] and Webber [1978b].

of anaphoric reference that have been discussed in the linguistics and philosophic literatures - the famous "donkey" sentence (cf. Bartsch [1976], Edmundson [1976], Hintikka & Carlson [1977]) and the problem of reference in disjunctive contexts (cf. Karttunen [1977]). Finally, to show that it is not just definite and indefinite noun phrases that can evoke entities in the listener's discourse model, I will illustrate in Section 4 an example of deictically-evoked entities and comment on the problem of describing them appropriately.

2. Indefinite Noun Phrases and Discourse Entities

Except after a copula, indefinite noun phrases <*2> may evoke a new discourse entity into a listener's discourse model. <*3> What I want to focus on here is appropriate **IDs** for them. Consider the following sentences.

1a. Wendy bought a yellow T-shirt that Bruce had liked.
 b. **It** cost twenty dollars.

2a. Each third-grade girl brought a pelican to Wendy's house.
 b. She is roosting **them** on her front lawn.

<*2>. I will often refer to these as "existentials" because of their logical interpretation as existential quantifiers.
<*3>. An indefinite noun phrase following a copula functions together with the copula as a predicate, e.g.

Beverly is a bargain hunter.
Bruce became a librarian.

As such, it is purely descriptive and does not refer to any particular librarian or bargain hunter, cf. Kuno [1970].

3a. If Bruce manages to catch a fish,
b. he will eat it for dinner.

4a. John didn't marry a Swedish woman.
b. She was Norwegian.

5a. Whether Bruce buys a mini-computer or an Advent TV,
b. he will have to do the repairs on it himself.

6. Every man who owns a donkey beats it.

I claimed earlier that the initial description (ID) of a newly-evoked discourse entity is critical for both model synthesis and anaphor resolution, since the ID mediates all reasoning about the entity until its assignment to some role within the model. An entity's ID should imply neither more nor less about it than is appropriate. Now consider what an appropriate description would be for the discourse entity that "it" refers to in sentence 1b. It is not "the yellow T-shirt that Bruce had liked", since sentence 1a. can be uttered truthfully even if Bruce had liked several yellow T-shirts (and both speaker and listener were aware of that fact). Nor is it "the yellow T-shirt that Bruce had liked and that Wendy bought", since sentence 1a. can be truthfully uttered even if Wendy had bought several such T-shirts. What is an appropriate description for the referent of "it" is something like "the yellow T-shirt that Bruce had liked and that Wendy bought and that was mentioned in sentence 1a."

What I am claiming is that in the case of a singular existential that is not within the scope of either negation, a universal quantifier, a hypothetical (e.g. "if", "suppose") or one of several other special contexts (cf. Webber [1978a]), the entity it evokes will be appropriately described via a conjunction of (1) the description inherent in the noun phrase (e.g. "yellow T-shirt that Bruce had admired"); (2) a predicate that embodies the remainder of the sentence (e.g. "which Wendy bought"); and (3) a predicate that relates that entity to the utterance evoking it (e.g. "which was mentioned in (or evoked by) sentence 6a."). This is the description that I am calling the entity's "initial description" or ID. Given how I specified its components then, it should not be surprising that I will claim that the ID of an existentially-evoked discourse entity can be derived from an appropriately structured sentence-level logical representation. Such a representation is independently motivated by its use in regular inference procedures.

Using a somewhat simplified version of the formalism described in Webber [1978a], a simple rule can be stated for forming the ID of an existentially evoked discourse entity - i.e.,

(Ex:C) . F{x} ==> (Ez) . z = ix: Cx & F{x} & **evoke** S,x

Here (Ex:C) is an example of **restricted quantification**, in which **C** represents an arbitrary predicate which x satisfies. **F**{x} represents an arbitrary open sentence in which x is free; i stands for Russell's definit operator, iota; and S is the label assigned to the proposition on the left-hand side of the arrow. Informally, this rule, which I shall call [RW-1], says that if a proposition S states that there is a member x of class C which makes **F**{x} true, then there exists an individual describable as "the C which F's which was evoked by proposition S". This individual is taken to be the discourse entity evoked by the existential noun phrase. For example, let **Y** stand for the predicate corresponding to "yellow T-shirt that Bruce had liked". <*4> Then sentence 1a. can be represented simply as

(Ex:Y) . Bought Wendy, x

Since this matches the left-hand side of the above rule, it follows that

(Ez) . z = ix: **Y** x & Bought Wendy,x & **evoke** S{1a},x

That is, there is an individual describable as "the yellow T-shirt that Bruce had liked, that Wendy bought and that was evoked by sentence 1a." The discourse entity so described is the referent of "it" in sentence 1b.

<*4>. I will soon be more precise about the representation of relative clause containing noun phrases. Here, where the descriptive part of the noun phrase can be treated as an

Examples 2-6 illustrate singular indefinite noun phrases in some of the special contexts noted above. While I will only be discussing examples 5 and 6 in this note, notice that in all five cases, the entity evoked by the indefinite noun phrase is appropriately described by taking into account at least the three factors mentioned above. That is, in example 2 the referent of "them" can be described uniquely as "the set of pelicans, each of which, mentioned in sentence 2a., some third grade girl brought to Wendy's house." <*5> In example 3, the referent of "it" can be described as "the fish mentioned in clause 3a. that Bruce has managed to catch, if Bruce has managed to catch a fish". In example 4, the negation appears intended to scope only "Swedish". Thus the discourse entity referent of "she" can be described as

---

unanalyzed unit, the predicate name Y is an adequate representation.

<*5>. A rule similar to [RW-1] is given in Webber [1978a] for existentials scoped by universals. In all, six such rules are given, covering

1. independent existentials (sg/pl)
   "I saw {a **cat**, **three cats**} on the stoop."
2. definite descriptions (sg/pl)
   "I saw **the {cat, cats} which hate(s) Sam.**"
3. distributives
   "**Each cat on the stoop** hates Sam."
   "**The three cats each** scratched Sam."
4. universally quantified existentials
   "Each boy gave each girl {**a peach, three peaches**}."
5. class dependent definites
   "Each boy gave a woman he knew **the {peach, two peaches} she wanted.**"
6. class dependent distributives
   "Each boy I know loves **every woman he meets.**"

"the woman mentioned in sentence 4a. that John married". (We later learn in sentence 4b. that she is Norwegian rather than Swedish.) IDs for the two other existentially-evoked discourse entities in examples 5 and 6 will be discussed in Section 3.

Notice that a definite noun phrase in the same context as an indefinite noun phrase will also evoke a discourse entity, but one whose ID is somewhat different. To see this, consider the following sentences.

7a. Wendy bought the yellow T-shirt that Bruce had liked.
 b. It cost twenty dollars.

8a. Each third grade girl has seen the pelican on Wendy's lawn.
 b. They prefer it to the plastic flamingo she had there before.

9a. John didn't marry the Swedish woman.
 b. He threw her over for a Welsh ecdysiast.

In each case, an appropriate description for the discourse entity evoked by the singular definite noun phrase is just that singular definite noun phrase itself - "the yellow T-shirt that Bruce had liked", "the pelican on Wendy's lawn", "the Swedish woman". While it is certainly true that the definiteness of these noun phrases may be contingent on context (i.e., identifiability within the speaker's model of the underlying situation), nevertheless unlike entities evoked by indefinite noun phrases, those evoked by definites do not depend for their appropriate IDs on the particular sentences the definite noun phrases appeared in.

The same characteristic behavior of definites and indefinites discussed for singular noun phrases holds for plural noun phrases as well. That is, while both indefinite and definite plural noun phrases evoke discourse entities, the unique initial descriptions that can be assigned to those entities will differ in the two cases. To see this, consider the following example.

10a. I saw **the** guys from "Kiss" on TV today.
b. I saw **the three** guys from "Kiss" on TV today.
c. I saw **all three** guys from "Kiss" on TV today.
d. I saw **some** guys from "Kiss" on TV today.
e. I saw **three** guys from "Kiss" on TV today.

11. **They** were being interviewed by Dick Cavett.

Sentences 10a-c each contains a definite plural noun phrase. That noun phrase should evoke a discourse entity into the listener's model, one appropriately described as "the (set of) guys from 'Kiss'". This can be verified by following either of these sentences by sentence 11 and considering what is the referent of the definite pronoun "they". <*6>

<*6>. While sentences 10b&c. provide the additional information that the number of guys in "Kiss" is three [not actually true - BLW], that information is not needed in order to describe the set uniquely. However, it should not be ignored, as it may be needed later in resolving a definite anaphor like "the three guys".

Sentences 10d&e, on the other hand, each contains an **indefinite** plural noun phrase. That noun phrase will evoke a discourse entity appropriately described as "the (set of) guys from 'Kiss' that I saw on TV today and that was mentioned in Sentence 10d(e)." This is because either sentence is consistent with there being other members of "Kiss" whom I didn't see on TV today, as well as other members whom I did see but whom I don't mean to include in my statement. <*7> Notice again that the set size information provided in sentence 10e. is not necessary for describing that set uniquely. However, it too may be useful later in resolving definite anaphora.

An interesting point is that there seem to be no other patterns that English determiners follow vis-a-vis discourse entity IDs. To see this consider the following sentences.

<*7>. This latter point is a subtle one, and usage may vary from person to person. That is, some people intend an indefinite plural noun phrase contained in a sentence S - "Some <x>s P" - to refer to the **maximal** set - i.e., "the set of <x>s which P". Other people intend it to refer to some subset of that set - "the set of <x>s which P which I (the speaker) intended to mention in sentence S". For a system to cope with this variation in usage, it would be better for procedures to derive the latter, non-maximal set description, which is always appropriate. If a system is sophisticated enough to associate a "belief space" with the speaker (cf. Cohen [1978]), other procedures can later access that belief space (if necessary or desirable) to judge whether the maximal set interpretation might have been intended. (This will again become an issue when I discuss other determiners like "many" and "several".)

12a. Few linguists smoke since **they** know it causes cancer.
b. Few linguists were at the party, but **they** drank more than the whole Army Corps of Engineers.

13a. Many linguists smoke although **they** know it causes cancer.
b. Not many linguists smoke since **they** know it causes cancer.

c. Many linguists don't smoke since they know it causes cancer.

In sentence 12a, the referent of "they" is the discourse entity appropriately described as "(the entire set of) linguists". That is, "few <x>s" can evoke the same discourse entity as the definite noun phrase "the <x>s". However as sentence 12b. shows, "few <x>s" can also pattern after the indefinite plural: the referent of "they" is the entity appropriately described as "the just-mentioned set of linguists who were at the party". (We learn from "few" that this set is small or smaller than the speaker expects.)

"Many", on the other hand, seems to pattern only after the indefinite plural. In sentence 13a., the referent of "they" is appropriately described as "the just-mentioned set of linguists who smoke". (We learn from "many" that this set of linguists is large or larger than the speaker expects.) Sentence 13b. shows that the reverse polarity "not many" acts like "few" vis-a-vis evoking discourse entities: the referent of "they" is the entire set of linguists. However as sentence 13c. shows, a NEG which occurs in the sentence auxiliary does not effect this same change in behavior: "they" refers to the just-mentioned set of linguists who don't smoke.

3. Two Interesting Reference Problems

One reason for this note is to point out the importance of description formation to both discourse model synthesis and reference resolution **and** to show that this process can, to an important degree, be formalized. I have taken as given the notion that a listener is using both the discourse and his or her knowledge of the world to synthesize a model of what s/he believes to underlie the discourse. Definite anaphora are viewed as means by which the speaker refers to entities in DM{S} that are presumed to have counterparts in the listener's model. What I want to show in this section is that this approach to definite anaphora can accommodate not only straight-forward cases as discussed above, but certain problematic cases as well.

3.1 Parameterized Individuals

The problem of formally characterizing the referent of "it" in examples like 6 below has often been discussed in the linguistics and philosophy literatures, cf. Bartsch [1976], Edmundson [1976], Hintikka & Carlson [1977].

6. Every man who owns a donkey beats it.

The problem has been taken to be that while "it" intuitively seems related to the embedded noun phrase "a donkey", there is no way to represent this logically in terms of simple quantifier scoping. What I shall show is that an approach in terms of discourse entities and their IDs makes this intuitive relationship simple both to explain and to represent.

First notice that this problem arises independently of how the matrix noun phrase is determined.

14. A man I know who owns a donkey beats it.
15. The man who owns a donkey beats it.
16. Which man who owns a donkey beats it?
17. No man who owns a donkey beats it.

In all these examples, "it" seems intuitively related to "a donkey". Informally, one might described its referent as "the just-mentioned donkey he owns", where "he" is bound to whatever value that "(each, a, the, which, no) man who owns a donkey" may take. But this is just a discourse entity of a rather special type - one with a parameterized ID, rather than a rigid one. I call such entities "parameterized individuals", borrowing the term from Woods & Brachman [1978]. <*8>

<*8>. The phrase "parameterized individual" is being used somewhat loosely to include "parameterized" sets, stuff, etc. For example,

(i) No man who owns two donkeys beats them.
them = the two donkeys he owns

Notice that parameterized individuals behave somewhat differently from the "actual" discourse entities the sentences evoke. <*9> That is, parameterized individuals all have the **same ID**, independent of how the noun phrase containing the relative clause is determined. On the other hand, the actual discourse entities evoked by these sentences do not. For example,

18a. Each man who owns a donkey beats it.
it = the donkey he owns
b. However, **the donkeys** are planning to get back at **them.**
the donkeys = the set of donkeys, each of which some man who owns a donkey owns
them = the set of men, each of whom owns a donkey

19a. The man I know who owns a donkey beats it.
it = the donkey he owns
b. But **the donkey** is planning to get back at **him.**
the donkey = the just-mentioned donkey that the man I know who owns a donkey owns
him = the man I know who owns a donkey

20a. Which man who owns a donkey beats it?
it = the donkey he owns
-- "None"
b.*Are **the donkeys** planning to get back at {**him**, **them**, ???}?
the donkeys = ???
c.*Is **the donkey** planning to get back at {**him**, **them**, ???}?
the donkey = ???

To show that this approach to definite anaphora in terms of discourse entities and their descriptions can explicate "donkey" sentences as well, I will have to introduce a bit more of the

<*9>. By "actual" discourse entities, I mean ones that can be referred to anaphorically in subsequent sentences.

formalism described in Webber [1978a]. That bit involves an extension of **restricted quantification**, cf. [RW-1] above. In restricted quantification, a quantification operator (e.g. V\-,E), the variable of quantification and the class it ranges over (noted implicitly as a predicate) constitute a structural unit of the representation. For example, "Every boy is happy" can be represented as

    (V\-x:Boy) . Happy x

This is truth functionally equivalent to

    (V\-x) . Boy x ==> Happy x

Similarly "Some boy is happy" can be represented as

    (Ex:Boy) . Happy x

which is truth functionally equivalent to

    (Ex) . Boy x & Happy x

The extension I will introduce will permit the representation of noun phrases with relative clauses as well as simple noun phrases. Semantically, a relative clause can be viewed as a predicate. One way to provide for arbitrary predicates is through the use of the abstraction operator "λ", cf. Hughes & Cresswell [1968]. For example, the noun phrase "a peanut" can be represented as

    (Ex:Peanut)

while the noun phrase "a peanut that Wendy gave to a gorilla" can be represented as

(Ex: λ (u:Peanut)[(Ey:Gorilla) . Gave Wendy,u,y])

In this case

λ(u:Peanut)[(Ey:Gorilla) . Gave Wendy,u,y]

names a unary predicate which is true if its argument is a peanut that Wendy gave to some gorilla.

Using this notation, sentence 6 can be represented as

(V\-x: λ(u:Man)[(Ey:Donkey) . Own u,y]) . Beat x,IT

By applying rule [RW-1] to the embedded clause [(Ey:Donkey) . Own u], the entity evoked by the existential can be identified as

iy: Donkey y & Own u,y & **evoke** S{6.1},u
"the just-mentioned donkey that u owns" <*10>

As I mentioned above, the semantics of restricted quantification is such that the variable of quantification, here x, satisfies the predicate in the restriction. Thus if x satisfies λ(u:Man)[(Ey:Donkey) . Own u,y], there must be an entity identifiable as

---

<*10>. In labeling each clause of a complex sentence, I use the following convention: if the matrix clause is labelled S, its leftmost embedded clause will be labelled S.1, the leftmost embedded clause in S.1 will be labelled S.1.1, etc.

iy: Donkey y & Own x,y & **evoke** S{6.1},y
"the just-mentioned donkey x owns"

This is a parameterized individual - parameterized by the variable in (V\-x:...) - that is a possible referent for "it" in the matrix sentence - i.e.,

(V\-x: λ(u:Man)[(Ey:Donkey) . Own u,y]) . Beat x, iy: Donkey y & Own x,y & **evoke** S{6.1},y
"Every man who owns a donkey beats the just-mentioned donkey he owns"

I noted above that a sentence like "Every man who owns a donkey beats it" could sensibly be followed by a sentence like "However, **the donkeys** are planning to get back at them" (cf. example 18). Given that I have shown how to account for the referent of "it" in the first sentence in terms of discourse entities and their formally derivable descriptions, can the referent of "the donkeys" be account for in the same way? <*11>

To show that it can, I need to present the rule for dealing with class dependent definite descriptions that I mentioned in footnote 5. This rule is motivated by examples such as 21, where the referent of "them" is presumably the discourse entity evoked by the noun phrase "the flower she picked", where "she" stands for the variable bound by "each girl in the class".

<*11>. I shall not take the time here to discuss the path from the phrase "every man who owns a donkey" to the discourse entity informally describable as "the set of men, each of whom owns a

21a. Each girl in the class gave Ivan the flower she picked.
b. He arranged **them** artfully in an empty Glenfiddach bottle.

This is a definite noun phrase, but because of its binding to the distributively quantified noun phrase "each girl", it will evoke a discourse entity with the properties of a set rather than an individual (cf. example 8). In this case, it will be "the set of flowers, each of which was the flower that some girl in the class picked". Simplifying for brevity here, this rule can be written

(Vx:K) . P x,iy:**C** x,y ==> (Ez) . z = {u|(Ex:K) . u = iy:**C** x,y}

where **K** represents an arbitrary unary predicate which x satisfies and both P and C represent arbitrary binary predicates. The right-hand side of this rule implies that in case the left-hand side matches some sentence, there will be a discourse entity roughly describable as "the set of u's, each of which is the thing that stands in relation C to some member of **K**".

Notice now that after the "it" is resolved in "Every man who owns a donkey beats it" (see above), the sentence matches the left-hand side of the above rule - i.e., "Every man who owns a donkey beats **the just-mentioned donkey he owns.** Thus it follows that there is a discourse entity describable as "the set of

---

donkey", since it is rather straightforward, cf. Webber [1978a]. This entity is a possible referent for "them" in sentence 18b.

donkeys, each of which is the just-mentioned donkey that some man who owns a donkey owns" - i.e.,

{w|(Ex:λ(u:Man)[(Ey:Donkey) . Own u,y])
    w = ιz: Donkey z & Own x,z & evoke S{18},z}

This is a possible referent for "them" in sentence 18b.

3.2 Disjunction

The other class of problematic examples that I want to discuss here in terms of discourse entities and their descriptions is one I first encountered in Karttunen [1977]. Karttunen presents examples like the following.

22. If Wendy has a car or Bruce has a bike, **it** will be in the garage.

23. Bruce can have either a bike or a car, but he must keep **it** in the garage.

24. Either Bruce has a new car or he has borrowed his brother's. In any case, **it** is blocking my driveway.

25. Whether Bruce buys a car or his brother buys a bike, he will have to keep it in the garage.

The problem is again to determine just what it is that "it" refers to.

I see two ways of approaching this problem in terms of discourse entities and their **IDs**. One way holds that in each sentence, each term of the disjunction evokes a different discourse entity into DM{L}, each with a different ID:

(22) "the car that Wendy has (if she has a car)"
"the bike that Bruce has (if he has a bike)"

(23) "the bike that Bruce will have (if he chooses a bike)"
"the car that Bruce will have (if he chooses a car)"

(24) "the new car that Bruce has (if Bruce has a new car)"
"Bruce's brother's car"

(25) "the car Bruce will have bought (if he buys a car)"
"the bike Bruce's brother will have bought (if Bruce's brother buys a bike)"

The truth of the disjunction (which seems in each case to be interpreted as exclusive "or") then guarantees there being one and only one entity in the model to which "it" refers. Notice that if the terms were conjoined rather than disjoined, the truth of the conjunction would imply the simultaneous existence of two entities within the model. In that case, either the referent of "it" would be ambiguous or the sentence would just be bizarre.

The other, I think nicer, way of approaching the problem holds that each sentence evokes only a **single** discourse entity into the model, with the indecision (i.e., the disjunction) embodied in its ID. That ID is of the form "A if P, otherwise B". For example, the entity evoked by sentence 22 would be describable as "the car that Wendy has (if she has a car) or the bike that Bruce has otherwise"; that evoked by sentence 23 would be describable as "the bike that Bruce will have (if he chooses a bike) or the car that Bruce will have otherwise"; that evoked by sentence 24, as "the new car that Bruce has (if he has a new car) or Bruce's brother's car otherwise"; and that evoked by sentence 25, as "the car Bruce will have bought (if he buys a car) or the bike Bruce's brother will have bought otherwise".

One advantage to this approach is that additional properties which truthfully follow from either **ID** can be ascribed to the entity without committing oneself to one description or the other. This can be useful in anaphor resolution. For example, in sentence 24, the subject of "block my driveway" must be a physical object, preferably large and somewhat mobile. This condition is satisfied by the discourse entity evoked by sentence 24, independent of which ID is appropriate.

Although there may be other ways to approach the problem of disjunction, the "donkey" problem, and the whole problem of definite reference in general, what I hope to have shown in these two sections is the robustness of an approach based on notions of a discourse model, discourse entities and their formally derived descriptions.

4. Conclusion

In arguing for the importance of description formation to both discourse model synthesis and reference resolution, I concentrated on how indefinite noun phrases evoke new entities into the listener's discourse model and how their appropriate initial descriptions (**IDs**) could be derived from a formal sentence-level representation of the text. There are many other ways in which discourse entities can be evoked, and many interesting problems in forming appropriate descriptions of them. I will conclude therefore with a brief discussion of deictically-evoked discourse entities and the problem of describing them appropriately.

The example comes from the children's book **Babar Loses his Crown** by Laurent de Brunhoff, and involves the following situation: Babar, King of the Elephants, decides to take his wife Celeste and his family on a trip to Paris. In packing for the trip

> "Babar puts his crown in a little red bag." (p.3)

They travel by train and then by taxi to their hotel in Paris, and when they arrive

> "Celeste opens all the bags. Last of all, she opens the little red one. 'Look!' she cries. 'What is this? A flute! Babar! **This is not your bag!**' " (p.10)

Before this point in the story, there should have been one little red bag in DM{L}. Now there should be two. The first is the existentially-evoked discourse entity (say, e{43}) - "the little red bag mentioned in sentence <x> that Babar put his crown in".

However if "this" on page 10 is **not** that entity, then it must be some other one (say, e{48}). How should it be described? Since "this" presumably points to the little red bag Celeste is opening, e{48} can appropriately be described as "the just-mentioned little red bag which Celeste is opening, which contains a flute and not Babar's crown, and which is **not** equivalent to e{43}". <*12>

The problem here is to be able to articulate clearly what each of these properties derives from since they do not come from a single sentence. In this case, one must determine what things relevant to the story do or do not follow from e{48}'s not being Babar's bag.

* * * * * * * *

In this note, I have attempted to explicate briefly an aspect of understanding definite anaphora that precedes the more frequently discussed problem of "anaphor resolution". This aspect involves accounting for what it is that definite anaphors

---

<*12>. Throughout this approach, I am making no assumptions about the separateness of discourse entities. That is, I am not assuming that two discourse entities are necessarily and for all times distinct, unless it is specified explicitly as was done here. Two discourse entities may have incompatible descriptions and as a consequence be **assumed** to be distinct. But I do not view it as impossible for two separately evoked discourse entities with different but compatible descriptions to later be found to be one and the same.

refer to and how such things become available. I moved from the notion of reference into a model to problems of how that model is synthesized, and in particular, how the entities in it are appropriately described. In this endeavor, I focused on the initial descriptions (IDs) ascribed to existentially-evoked entities, briefly touching upon deictically-evoked entities as well. For further discussion, see Webber [1978a&b]. In the next quarter, we will begin implementing this approach to definite anaphora in terms of KLONE structures (cf. Brachman [1978], Woods & Brachman [1978]), beginning with the design of a representation for quantified information which is adequate both for inference and the identification of discourse entities.

The RUS System

Robert J. Bobrow

1. Introduction

It has long been the goal of those writing natural language processing systems to express syntactic constraints in a broad, general way while using tight semantic constraints to guide the parsing and to interpret the resulting structure. Ideally, syntax and semantics (pragmatics) would be so thoroughly integrated in such a system that at any time during the parsing of an input string the component that had the most relevant information would be in control of the processing. That is, syntax and semantics would walk hand in hand - not one before the other, and not one in front occasionally restrained by jerks on a leash between them. On the other hand, in the ideal case both syntax and semantics would be modifiable and extendible individually, without detailed knowledge of the operation of the other. The RUS natural language system described below achieves these goals, providing a framework for close and efficient interaction between a general syntactic processor and a semantic-pragmatic interpreter.

2. Background

The design of the RUS system may be best understood by contrasting it with the design of two earlier natural language systems, the LUNAR system [Woods et al., 1972; Woods, 1977] and the "semantic grammar" system used in SOPHIE [Burton, 1976], which represent two extremes in the way in which syntax and semantics can be combined in a natural language system. The RUS system can be viewed as a synthesis of these two approaches, and in fact its design owes much to our analysis of the advantages and shortcomings of each of them.

2.1 The LUNAR Approach

In the basic LUNAR paradigm, the syntactic and semantic components of the system are completely distinct, with no interactions between them and entirely separate knowledge representations. Syntactic parses are discovered independently of semantic constraints by a parsing procedure that interprets a grammar expressed in Woods' ATN notation [Woods, 1970]. After syntax has produced a syntactic description of the input string, the resulting "parse tree" is passed to a semantic interpreter, which applies a set of semantic rewrite rules to it, using a recursive control structure that ensures that rules are applied to all parts of the tree. There are several advantages to this approach:

1) Because of the complete separation of the two processes, it is possible to modify each separately, simply keeping in mind the allowable structures for the parse trees which form the only communication mechanism between the two components.

2) The ability of the system to capture syntactic regularities is substantial. Thus, if the set of semantic rules is extended to allow the interpretation of one new sentence, a large number of its syntactic variants will also be immediately interpretable without modification of the syntactic processor or the addition of special semantic rules.

On the other hand, a LUNAR-type approach is relatively inefficient, because there is no way for semantic information to be used to guide the production of parse trees. Thus for sentences with more than one acceptable syntactic structure - the vast majority - the semantic interpreter can only be used to accept or reject the entire parse tree and not to indicate how to modify the parsing to produce an acceptable structure. In addition, the semantic interpreter requires its own control structure to scan parse trees, and the applicability of semantic rules must be determined by "pattern matching" on tree fragments. The complexities this introduces are substantially reduced in the RUS system.

2.2 Semantic Grammars

In the other approach to natural language processing mentioned above - "semantic grammars" [Burton, 1976] - the syntactic and semantic processes are so thoroughly integrated that not only is the distinction between them totally erased, there is also no separation between the representation of syntactic regularities and the representation of the semantics of the domain of discourse. That is, a semantic grammar is "a set of grammar rules which characterize, for each concept or relationship, all of the ways of expressing it in terms of other constituent concepts" [Burton, 1976: p. 23]. This representation clearly allows semantic constraints to control the parsing process, and thus can be used to produce extremely efficient systems. In this type of system there is no notion of a "parsable but not interpretable" sentence - if the parser does anything, it produces an interpretation of the meaning of the sentence.

The disadvantages of a semantic grammar approach arise because of the way in which the merged syntactic and semantic knowledge are represented and the completeness of the merger. To produce a semantic grammar for a new domain, it is often the case that no portion of an existing semantic grammar can be used - that is, there is no syntactic knowledge separable from domain-specific semantics. A related problem is that when the semantic grammar for a domain is to be extended to handle a larger semantic domain, the ability to handle syntactic paraphrases in the new area must be explicitly included in the system, even though structurally similar paraphrases may be dealt with in other portions of the grammar. Both of these difficulties arise from the fact that the knowledge representation for the semantic grammar is used to encode the explicit cross-product of syntax and semantics. There is, for example, no notion of the structure "noun phrase" separate from the collection of structures representing "noun phrase which refers to an X" (for each semantic category X in the domain).

The RUS system attempts to steer a course between the Scylla of inefficient but readily generalizable modularity and the Charybdis of efficient but rigid integration. That is, it facilitates separate modification of grammar and semantics, while providing a framework for a strong interaction between them even at the earliest steps of the parsing process. The structure of the resulting system is closely related to the general class of notions referred to as the "case structure" approach to language, and makes use of a valuable property of English (and we suspect all human languages) that is also the basis for the "deterministic parsing" system developed by Marcus at MIT

[Marcus, 1977]. Section 3 of this paper describes this property, which we refer to as the **incremental parsability** of English.

2.3 Initial Development of RUS

The first version of the RUS system was developed in 1976 for the BBN portion of the ARPA Intelligent Terminal (IT) Project. The original intent was to modify the LUNAR system to meet the demands of the IT domain. However, it was determined that the LUNAR system would not be efficient enough -- the primary source of inefficiency lying in its semantic interpreter. Although substantial effort was made to improve the interpreter's performance, the results were unsatisfactory. The initial IT front-end was thus implemented as a "semantic grammar". Although this proved to be quite efficient to run, the investment required to extend and modify it led us to investigate alternative approaches.

The first natural language system using the basic RUS framework was demonstrated in January 1977, at the end of the IT project. It was quite efficient and reasonably easy to modify and extend, although what we consider to be the critical features of the RUS approach were not immediately apparent in its code. Subsequent work to produce a clean and understandable version of the system led to a clearer recognition of the basic ideas embodied in the initial system. These ideas were used in the implementation of natural language front-ends for two other systems: the WHY system, a tutorial program which engages students in an interactive Socratic dialogue about weather processes [Stevens & Collins, 1977], and an industrial management information system. The fact that the implementation of RUS for the WHY system required only four person-months and could handle a vocabulary of over 1000 words indicates that the basic RUS approach is a system with considerable generality and transportability.

3. Incremental Parsing

3.1 Structural Descriptions

In general, syntactic processing can be viewed as a mechanism for providing a structural description of a sequence of words which form a syntactic unit or phrase (e.g. a clause, a noun phrase, a prepositional phrase, etc.). Such a structural description consists of a set of structural descriptions of the phrase's "constituent syntactic units" and a specification of the syntactic relations that hold between the constituent units and the unit as a whole. <*13>

<*13>. Or, alternatively, the syntactic relations that hold between the constituents.

As a simple example, consider the clause

"THE FLEET LAUNCHED TWO SORTIES THIS MORNING"

This might be assigned the somewhat simplified structural description:

```
[CLAUSE
  HEAD = LAUNCH
  TENSE = PAST
  SUBJECT = [NP
               HEAD = FLEET
               DETERMINER = [DET
                              ARTICLE = THE]]
  OBJECT = [NP
              HEAD = SORTIE
              DETERMINER = [DET
                             NUMBER = TWO]]
  TIMEMOD = [NP
               HEAD = MORNING
               DETERMINER = [DET
                              ARTICLE = THIS]]]
```

This indicates that the syntactic unit is a clause with five constituents: a verb, a tense, and three noun phrases. They are related to the the clause as a whole by the syntactic relationships <*14> HEAD (indicating the main verb of the clause), TENSE, SUBJECT (indicating the logical subject of the clause), OBJECT (indicating the logical object of the clause), and TIMEMOD (indicating a modifier of the clause which specifies some aspect of the time of the event). The noun phrases in turn have their own internal structure.

<*14>. Names for syntactic relationships used in RUS will be indicated by capital letters.

3.2 Incremental vs. Wholistic Processes

From the point of view of language understanding, a primary purpose for providing such syntactic descriptions is that semantic interpretation can then be viewed as a process which produces a representation of the interpretation of a syntactic unit as a function of the interpretations of its constituents and the syntactic relations among them. In general, it is possible that either syntax or semantics or both might be "wholistic" rather than "incremental" processes. By this we mean that in the case of syntax, the parser might not be able to assign the set of syntactic relations until it had considered the totality of constituents of a given unit. (In fact, the parser might not even be able to assign the structure of one constituent without considering the structure and contents of all other constituents.) In the case of semantic interpretation, a "wholistic" process would require access to the entire collection of constituent interpretations and their syntactic relations before it could produce any part of the interpretation of the entire unit.

However, the results we have obtained in the RUS system, combined with the results obtained by Marcus, suggest that to a very large extent English can be parsed and interpreted in an incremental, left to right fashion. That is, there is a class of syntactic relations which can be determined incrementally and which is sufficient to provide the information needed by semantics. Such incrementally discovered relations can be associated with constituents using only limited information about the relations assigned to other constituents, and in particular with extremely limited information about the relations of constituents to the right of the given constituent. Corresponding to this incremental assignment of syntactic role to a constituent is the incremental construction of a semantic interpretation which takes place as each constituent is assigned a syntactic relation to the total unit. These two processes are closely linked, but not simply because the semantic process uses the result of the syntactic one in an incremental fashion. The ability of the syntactic process to assign a syntactic role without considering the roles of all other constituents is based critically on its ability to "ask semantics" whether a possible syntactic assignment produces a plausible incremental interpretation. In actual fact, there are phenomena in English that are not amenable to such incremental processing (such as extraposition), but they seem highly constrained in such a way as to permit relatively simple extensions of the incremental scheme to work well.

3.3 The Basic Operation of RUS

The RUS system consists of two components - a syntactic/control component and a semantic interpretation component. As we mentioned above, RUS is currently implemented as a front-end for three different systems. In all three versions of RUS, the interface between the syntactic/control component and the semantic component is well-defined and reasonably simple. In fact, these systems use the same syntactic component and differ only in their specific semantic interpreters.

RUS's syntactic/control component consists of a grammar written in the ATN formalism of Woods [1970], which is translated into an INTERLISP function by the "grammar compiler" of Burton and Woods [1976]. Using standard syntactic categories, the grammar defines a very large subset of English, more extensive than those found in LUNAR, GSP [Kaplan, 1973] or SPEECHLIS [Woods et al., 1975], with the exception of limitations on certain types of conjunction. It can handle complex relative clauses, complements, comparatives, passives, questions, extrapositions, ellipsis and gapping, although not all of them can presently be handled in their most general form.

The semantic component for the WHY system is expressed in the form of a "case frame dictionary", as is the semantic component of the industrial management information system we have designed. The semantic component used in the current ARPA command and control/display-processing domain is expressed in the KLONE implementation of Brachman's Structured Inheritance Networks (S-I Nets) [Brachman, 1977; Woods & Brachman, 1978]. One characteristic that these three versions share lies in their indexing of interpretation rules.

In the earliest version of RUS, semantic interpretation rules were associated with words that could be the HEAD's of phrases (e.g., nouns which could head a noun phrase, verbs which could be the main verb of a clause, etc.). These rules indicated how the interpretation of the phrase was to be modified when a new constituent was added with a specified syntactic relation to the phrase. The most common form of rule indicated (1) the type of syntactic relation it applied to (e.g., SUBJECT) and (2) one or more sub-rules whose applicability depended on the semantic interpretation of the new constituent. For example, the SUBJECT of a clause might be treated as an AGENT if it was interpreted as the description of an animate being. Or it might be treated as an INSTRUMENT if its interpretation was consistent with its being the instrument of the action specified by the clause head.

With the implementation of the WHY system version of RUS, we extended the semantic interpreter to allow "inheritance" of semantic rules by classes of words. For example, since all nouns describing physical objects can take modifiers that specify location - subject to specific plausibility checks on the type of location and the type of object - a single rule for interpreting location modifiers can be associated with "physical object", which can be accessed from all of its denotata. This has led to a substantial reduction in both the size of the semantic rule-set and the effort needed to extend and modify it. Examination of this new inheritance mechanism showed that it was closely related to the concepts being developed by Brachman on Structured Inheritance Networks, which we are using for the representation of world knowledge in this current project. This has led us to develop ways of representing semantic rules for the RUS parser in the KLONE implementation of S-I Nets. We are using KLONE here as the framework for building both the semantic interpreter and the resulting semantic interpretation of sentences.

The existence of three separate systems making use of RUS based on a shared syntactic/control component as provided several benefits:

1) Syntactic capabilities whose need is first noted in one domain have been provided to the other domains with little cost, thus substantially improving the overall competence of the parser.

2) The effort to maintain compatibility among the three systems has proven to be an excellent forcing function for developing and maintaining a clean, well-defined syntactic/semantic communications interface.

3) In general the solutions to problems in building semantic interpretations in any one domain have led to improvements in both the understanding and implementation of the semantic processors in all domains.

3.4 Illustration

To illustrate the basic mode of operation of the parser, we will caricature its operation on a simple sentence using a trivial grammar. Consider the sentence

"DISPLAY ALL CRUISERS WITH A LARGE HEXAGON SYMBOL".

The following abbreviated ATN network can be used to recognize an imperative clause.

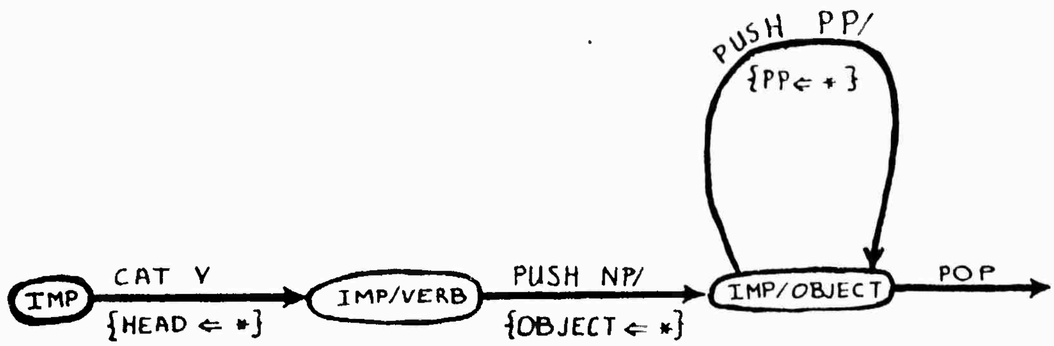

This characterizes an imperative clause as comprising a verb followed by a noun phrase (NP/), optionally followed by one or more prepositional phrases (PP/'s). The points of semantic/syntactic interaction occur where the ATN decides what syntactic relation each syntactic constituent bears to the clause. These are indicated by notations of the form "{rel <= *}" where "rel" is a the name of a syntactic relation. For example, in this network we have indicated that the verb bears the syntactic relation HEAD to the clause, the NP/ is the OBJECT, and the PP/'s bear the relation PP to the clause. The parser searches for these constituents by calling on networks defining NP/ and PP/.

If the desired constituent appears in the correct place (e.g., an NP/ occurs immediately after the verb), the result of such a call is a data structure containing both a syntactic characterization of the unit and its **internal** <*15> semantic interpretation. The parser then interacts with the semantic interpreter again to ascertain whether it is semantically acceptable for the given constituent to have the specified syntactic relation to the phrase - i.e., whether there is a semantic rule which specifies how the interpretation of the current phrase can be extended by adding the new constituent with the specified syntactic relation (e.g., adding the NP/ as the OBJECT of the clause). If the assignment of the new constituent to the indicated syntactic relation fails semantically, then the arc in the ATN is considered to fail. In this case, if the NP/ arc fails because of semantic rejection, the entire ATN for the

---

<*15>. That is, the semantic interpreter operates once to assign a constituent its interpretation, independent of whatever larger unit the constituent may eventually be included in.

imperative clause fails, because there is no syntactic alternative to the arc in the IMP network. <*16> On the other hand, if the PP/ arc fails because of a semantic rejection, then the parser would attempt to POP a completed interpretation of the imperative, since this arc is an alternative to PUSHing for the PP/.

In the example sentence, DISPLAY would be the HEAD of the phrase, and "ALL CRUISERS" would be produced as the noun phrase OBJECT. (While the NP/ network would syntactically allow the prepositional phrase "WITH A LARGE HEXAGON" to modify a noun phrase, the semantic interpreter would reject it as a possible modifier of "ALL CRUISERS". Thus it is the interaction of syntax and semantics that determines that "ALL CRUISERS" rather than "ALL CRUISERS WITH A LARGE HEXAGON" is proposed as the OBJECT of DISPLAY.)

As indicated earlier, once a syntactic constituent has been found, it is usually possible to determine what syntactic relations it bears to the phrase as a whole. However, it is not always possible to identify these syntactic relationships immediately nor to build an incremental interpretation in

---

<*16>. In general, there may be several ways for the ATN to find an NP/ in the desired location because of syntactic or semantic ambiguity. Because of this fact all the possible NP/'s must be checked before the entire IMP network fails.

parallel with syntactically characterizing a constituent. As an example of the first difficulty consider the determination of a critical syntactic relation, the "logical subject" of a clause, which we indicate as SUBJECT. In active declarative clauses, the SUBJECT is the first NP/ in the clause, but in passive clauses, the first NP/ is the "logical object", indicated as OBJECT, of the clause. Thus it is impossible to tell all the syntactic relationships of the first NP/ to the clause until the voice of the clause has been determined. (Note: The first noun phrase does bear an important relation to the clause (which we indicate as FIRSTNP) in all clauses - useful in determining focus and in dealing with verb phrase ellipsis. Thus, the first noun phrase bears two relations to the clause - FIRSTNP and either OBJECT or SUBJECT.)

The second problem mentioned above reflects the privileged status of the HEAD of a phrase. That is, it is difficult to start building the interpretation of a clause until the main verb is found, and it is difficult to start building the interpretation of a noun phrase until the head noun is found. For the latter, this means that the RUS system keeps track of all adjectives, participles and other pre-nominal modifiers and their syntactic relations, but does not attempt to transmit them to the semantic interpreter until the head noun has been discovered and transmitted. In clauses a similar process occurs not only for the FIRSTNP (as SUBJECT or OBJECT), but also for preposed prepositional phrases, as in

"ON THE TOP OF THE SCREEN I WOULD LIKE TO SEE ...".

Note also such potential difficulties as the fact that preposed PP/'s in a sentence may not even be clause-level constituents - e.g.,

"OF THE SHIPS WITHIN 100 MILES OF CYPRUS, WHICH ONES ARE ..."

where the prepositional phrase is actually a constituent of the partitive noun phrase that starts "WHICH ONES". These syntactic variations are among the many handled by RUS, in such a way that the semantic interpreter need not know whether these forms occurred or whether the input was one of the (substantially) equivalent syntactic alternatives

"I WOULD LIKE TO SEE ... ON THE TOP OF THE SCREEN" and

"WHICH OF THE SHIPS WITHIN 100 MILES OF CYPRUS ARE ...".

4. Improvements to Increase Determinism and General Efficiency

This quarter a substantial effort was made to improve the efficiency of the parser, and in particular to make the parsing process as deterministic as possible. There were several motivations for this:

1) In the ARPA version of RUS, the semantic processor uses a preliminary implementation of S-I Nets which is relatively space inefficient. Thus there is a heavy space penalty for calls to the semantic processor which do not produce structure that is actually used in the final parsing.

2) Although the results obtained by Marcus [1977] indicate that a substantial part of English can be dealt with in a deterministic fashion without backup, it was unclear how much his results depended on the characteristics of his non-ATN parsing mechanism and how much could be captured in an ATN parser.

3) General efficiency improvements, especially ones leading to greater determinism, are important not only for "production versions" of the system in which rapid response to users is vital, but also for experimental development of the type undertaken in the ARPA contract, since a great deal can be learned if one can afford to parse a large number of sentences to investigate various linguistic problems.

The amount of non-determinism in parsing a sentence is reflected in part in the amount of back-up that occurs. The basic ATN is a non-deterministic parsing mechanism: when more than one arc leaves a state in the ATN the parser must treat that state as a potential **branch-point**. That is, the parser must select an arc to follow, and if its path from that arc becomes blocked, it must be prepared to back-up to previous branch-points and try alternative arcs. A deterministic parser on the other hand must be able to treat a state with many arcs as a **choice-point** and make the correct choice of which arc to follow, without allowing for any back-up to that state.

In order to see how much of the non-deterministic behavior of the RUS ATN was unavoidable, we ran it on a collection of test sentences which included many linguistic phenomena normally considered to require non-deterministic parsing - ambiguous words (e.g., words like "flow" which can be nouns or verbs), structurally ambiguous noun-noun modifiers (e.g., missile cruiser radar systems), reduced relative clauses (e.g., "all missiles cruisers launched this year" versus "all missile cruisers launched this year"), etc. We then performed a detailed analysis of the types of back-up that occurred. This indicated three major causes for back-up:

1) the existence of unnecessary branch-points in the ATN;
2) the preponderance of "hypothesis-driven" (as opposed to "data-driven") characterizations of English grammar found in the ATN;
3) the interaction of the normal depth-first control structure of the ATN with the capability for semantic rejection of constituents.

In a typical ATN, there are many states that are not true non-deterministic branch-points **for any one sentence.** That is, for any given sentence there is at most one acceptable arc from any given state. In those cases, the parser should be able to take the correct arc and not have to provide for back-up to that state. Fortunately, the Burton ATN compiler has extended the normal ATN notation in a way that allows us to indicate such states to the compiler. That is, it permits any set of arcs from a single state to be combined into one "GROUP" arc. In a GROUP arc, the first arc whose test succeeds is the only arc which can be taken from the GROUP. Analysis of the RUS ATN showed that many arcs could be collected into GROUPs without any difficulty. Moreover, in states where arcs could not be GROUPed immediately, it was often the case that GROUPing was facilitated by allowing a small amount of look-ahead in the tests on the arcs. That is, such arcs were permitted to notice not only the characteristics of the current word, but also characteristics of the next one or two.

Finally, in this vein of eliminating unnecessary branch-points, we introduced the notion of an "almost-GROUP". This captured our intuition that most sentences could pass deterministically through a given state, and moreover, it would be easy to distinguish the sentences which had to be treated non-deterministically. Thus a single node could effectively be split in two, with one GROUP that split the situation into deterministic and non-deterministic cases and another GROUP for the deterministic case.

The second cause of back-up mentioned above has been pointed out by Marcus [1977] - the typical use of ATN's as a top-down, hypothesis-driven parsing mechanism. That is, when a point in the parsing is reached where it is possible for a constituent of type X to appear, the parser PUSHes to a network which actually looks for X. Normally this is done purely on the basis of the structure which has been found up to this point in the sentence, with no checking being done on the actual words. We would like to avoid such a PUSH arc if it is clear that the current word (or the next few words) precludes such a constituent. For example, there are places in the analysis of a clause where a noun phrase might or might not occur - e.g., after a main verb which can be used either transitively or intransitively. We do not want to PUSH for a noun phrase there if the next word clearly precludes its presence - e.g., if the next word is a preposition or a conjunction.

After analyzing where the RUS ATN PUSHed for constituents which we felt were "obviously" not present, we were able to put tests on the offending PUSH arcs which blocked them when the next words were obviously inconsistent with the hypotheses they embodied. We found that this required looking no further than the next three words, and often no further than the next one or two words. This is consistent with, and a slight strengthening of Marcus's "three chunk" look-ahead. Although there are cases where the three word look-ahead is insufficient, they seem to be relatively rare.

Finally, the third source of back-up lay in the very heart of the RUS approach, namely the incremental semantic testing of constituents, coupled with the ATN's standard "depth-first" control structure. As we indicated in Section 3.4, if we PUSH for a constituent and the result is semantically unacceptable, we must still allow for the possibility that a semantically acceptable constituent exists, but that the one discovered first by the PUSH was not it. Moreover, given the ATN's "depth-first" control structure, all alternative possible constituents of the desired category will be found **before** any alternatives to the PUSH are tried. This is equivalent to saying that there is more likely to be a constituent of the specified type here which will fit into the current phrase than that the first semantically acceptable constituent of that type found will fit somewhere, but **not** in the current phrase.

However, as the parser becomes more nearly deterministic, the latter is more likely to be the case. That is, the first semantically meaningful result returned from a PUSH is likely to be the best description of what actually occurs at that position. This is particularly true for optional constituents, such as prepositional phrase modifiers (particularly those specifying location or time). A frequent case is where an embedded noun phrase (NP/) will PUSH for a prepositional phrase. Although the parser will find one, the semantic interpreter will reject it as a modifier of the embedded NP/. The point is not that the PP/ is incorrectly formed, but that it belongs to the matrix clause or NP/.

For example, consider the sentence "DISPLAY THE CRUISER ON THE UPPER HALF OF THE SCREEN". When the parser is processing the embedded noun phrase that starts "THE CRUISER", it will hypothesize a prepositional phrase and subsequently find "ON THE UPPER HALF OF THE SCREEN". This is indeed the correct prepositional phrase to find at this point, as opposed to "ON THE UPPER HALF", but it is **not** one which modifies "THE CRUISER". If the control structure of the ATN forces us to explore all other possibilities for finding a prepositional phrase at this point, we will generate useless parses of acceptable but irrelevant prepositional phrases before we eventually decide that the noun phrase "ALL CRUISERS" has no prepositional phrase modifiers and that the prepositional phrase "ON THE UPPER HALF OF THE SCREEN" is a modifier of the clause.

In order to avoid this difficulty we have changed the control structure of the parser so that it will not back-up to produce alternative parsings of a constituent until all other alternatives have been investigated. That is, the parser will produce just the first semantically acceptable result of a PUSH.

If that is rejected semantically as a constituent of the current phrase, the parser will postpone trying branch points which would produce alternative results for the PUSH until all other alternatives have been tried. When this modification is combined with the well-formed-substring facility (WFS) which is a normal part of the parser, we have the result that if a PUSH is semantically rejected essentially because the correct syntactic structure was hypothesized by the wrong level of network, the result is stored in the WFS. If some higher-level phrase then PUSHes for the same type of constituent at the same place in the string, it will get the previously found result without further parsing.

The net effect of these changes has been better than expected. The new version of the parser handles about 70% of the test sentences completely deterministically, with no backup whatsoever, and another roughly 15% with only the back-up caused by semantic rejection and placement of a constituent at a higher level. Almost all the cases where the parser actually has to back up are ones which cannot be resolved locally and which lead people to "garden path". Since the sentences we chose as examples include some of the worst situations for deterministic parsing, we feel that the parser will have a very high rate of deterministic parsing in normal use.

Abstracts

On The Epistemological Status of Semantic Networks

Ronald J. Brachman
BBN Report No. 3807*

April 1978

This paper examines in detail the history of a set of network-structured formalisms for knowledge representation -- the so-called, "semantic networks". Semantic nets were introduced around 1966 as a representation for the concepts underlying English words, and since then have become an increasingly popular type of language for representing concepts of a widely varying sort. While these nets have for the most part retained their basic associative nature, their primitive representational elements have differed significantly from one project to the next. These differences in underlying primitives are symptomatic of deeper philosophical disparities, and I discuss a set of five significantly different "levels" at which networks can be understood. One of these levels, the "epistemological", or "knowledge-structuring, level, has played an important implicit part in all previous notations, and is here made explicit in a way that allows a new type of network formalism to be specified. This new type of formalism accounts precisely for operations like individuation of description, internal concept structure in terms of roles and interrelations between them, and structured inheritance. In the final section of the paper, I present a brief sketch of an example of a particular type of formalism ("Structured Inheritance Networks") that was designed expressly to treat concepts as formal representational objects. This language, currently under development, is called "KLONE", and it allows the explicit expression of epistemological level relationships as network links.

*To appear in **Associative Networks - The Representation and Use of Knowledge in Computers.** Nicholas V. indler, ed. New York: Academic Press, 1978.

A Formal Approach to Discourse Anaphora

Bonnie Lynn Webber
BBN Report No. 3761

May 1978

Extended natural language communication between a person engaged in solving a problem or seeking information and a machine providing assistance requires the machine to be able to deal with anaphoric language in a perspicuous, transportable non-ad hoc way. This report takes the view that dealing with anaphoric language can be decomposed into two complementary tasks: (1) identifying what a text potentially makes available for anaphoric reference and (2) constraining the candidate set of a given anaphoric expression down to one possible choice. The second task has been called the "anaphor resolution" problem and, to date, has stimulated much research in psychology and artificial intelligence natural language understanding.

The focus of this report is the first task - that of identifying what a text makes available for anaphoric reference and how it does so. Evidence is given to back up two strong claims:

1. None of the three types of anaphoric expressions that I have studied - definite anaphora, "one"-anaphora and verb phrase deletion - can be understood in purely linguistic terms. That is, none of them can be explained without stepping out of the language into the conceptual model each participant is synthesizing from the discourse.

2. On the other hand, if a discourse participant does not assign to each new utterance in the discourse a formal representation in which, inter alia,

   a. quantifiers are indicated, along with their scopes;
   b. main clauses are distinguished from relative clauses and subordinate clauses;
   c. clausal subjects are separated from clausal predicates;

   then s/he will not be able to identify all of what is being made available for anaphoric reference.

Building on these claims, I show that there is an intimate connection between such a formal sentential analysis and the

synthesis of an appropriate conceptual model of the discourse. The computational implications of this research are discussed, primarily in terms of possible implementations within current levels of technology.

Discourse Model Synthesis:
Preliminaries to Reference

Bonnie Lynn Webber

(Paper presented at the Sloan Workshop on Computational Aspects of Linguistic Structure and Discourse Function, University of Pennsylvania, 24-27 May 1978)

This paper starts from the point of view that a speaker's successful use of a definite anaphoric expression relies upon a listener's ability to synthesize a discourse model similar to that of the speaker and inhabited by similar discourse entities. That is, it is those discourse entities which are accessible by definite anaphora.

Now contributing towards the process of synthesizing a discourse model are two different types of input: one is the actual text and the other is the listener's broader expectations about what that text is about. This paper just focuses on how the determiner structure of an explicit noun phrase is taken in part to govern what discourse entities should be available in the listener's discourse model. The paper then speculates on several interrelationships between the text's contributions to the synthesis of an appropriate discourse model and those of the listener's contentful expectations.

References

Bartsch, R. Syntax and Semantics of Relative Clauses. In R. Bartsch, J. Groenendijk & M. Stokhof (Eds.), **Amsterdam Papers on Formal Grammars.** The Netherlands: University of Amsterdam, 1976.

Brachman, R. On the Epistemological Status of Semantic Networks. (Technical Report 3807) Cambridge MA: Bolt Beranek and Newman Inc., 1978.

Burton, R. Semantic Grammar: An Engineering Technique for Constructing Natural Language Understanding Systems. (Technical Report 3453) Cambridge MA: Bolt Beranek and Newman Inc., 1976.

Burton, R. & Woods, W.A. A Compiling System for Augmented Transition Networks. Preprints of Coling-76. Ottawa, Canada: University of Ottawa, June 1976.

Cohen, P. On Knowing What to Say: Planning Speech Acts. Technical Report No. 118. Department of Computer Science, University of Toronto, 1978.

Collins, A., Brown, J.S. and Larkin, K. Inference in Text Understanding. (CSR-40) Center for the Study of Reading, U. of Illinois and Bolt Beranek and Newman Inc., 1977. (Also in R. Spiro, B. Bruce and W. Brewer (Eds.), **Theoretical Issues in Reading Comprehension.** New Jersey: Lawrence Erlbaum Associates, 1978.)

Edmondson, J.A. Semantics, Games and Anaphoric Chains. In R. Bartsch, J. Groenendijk & M. Stokhof (Eds.), **Amsterdam Papers on Formal Grammars.** The Netherlands: University of Amsterdam, 1976.

Hintikka, J. & Carlson, L. Pronouns of Laziness in Game-theoretical Semantics. **Theoretical Linguistics,** 1977, **4**(1/2), 1-30.

Hughes & Cresswell, M. **An Introduction to Modal Logic.** London: Methuen Publishing Co., 1968.

Kaplan, R.M. A General Syntactic Processor. In R. Rustin (Ed.), **Natural Language Processing.** New York: Algorithmics Press, 1973.

Karttunen, L. Discourse Referents. In J. McCawley (Ed.), **Syntax and Semantics** (Volume 7). New York: Academic Press, 1976.

Karttunen, L. Whichever Antecedent. Squib Anthology, Chicago Linguistics Society, 1977.

Kuno, S. Some Properties of Non-Referential Noun Phrases. In R. Jakobson and S. Kawamoto (Eds.), **Studies in General and Oriental Linguistics.** Tokyo, Japan: TEC Company Ltd., 1970.

Levin, J. & Goldman, N. Process Models of Reference. Unpublished ms. Marina del Rey CA: Information Sciences Institute, 1978.

Lyons, J. **Semantics.** England: Cambridge University Press, 1977.

Marcus, M. A Theory of Syntactic Recognition for Natural Language. Unpublished Ph.D. thesis, Dept. of Electrical Engineering, MIT, 1977.

Stenning, K. Understanding English Articles and Quantifiers. Unpublished doctoral dissertation. New York: The Rockefeller University, 1975.

Stevens, A. & Collins, A. The Goal Structure of a Socratic Tutor. In **Proceedings of the ACM Annual Conference,** 1977. Also Technical Report 3578 Cambridge, MA: Bolt Beranek and Newman Inc., 1977.

Webber, B.L. A Formal Approach to Discourse Anaphora. (Technical Report 3761) Cambridge MA: Bolt Beranek and Newman Inc., 1978a.

Webber, B.L. Discourse Model Synthesis: Preliminaries to Reference. Presented at the Sloan Workshop on Computational Aspects of "Computational Aspects of Linguistic Structure and Discourse Setting", University of Pennsylvania, 25-28 May 1978b.

Woods, W.A. Transition Network Grammars for Natural Language Analysis. **Communications of the ACM,** 1970, 13(10), 591-606.

Woods, W.A. Semantics and Quantification in Natural Language Question Answering. (Technical Report 3687) Cambridge MA: Bolt Beranek and Newman Inc., 1977.

Woods, W.A., Kaplan, R.M. & Nash-Webber, B.L. The Lunar Sciences Natural Language Information System: Final Report. (Technical Report 2378) Cambridge MA: Bolt Beranek and Newman Inc., 1972.

Woods, W.A. & Brachman, R.J. Research in Natural Language Understanding - Quarterly Technical Progress Report No. 1. (Technical Report 3742) Cambridge MA: Bolt Beranek and Newman Inc., 1978.

Woods, W.A., Bates, M., Brown, G., Bruce, B., Cook, C., Klovstad, J., Makhoul, J., Nash-Webber, B., Schwartz, R., Wolf, J., & Zue, V. Speech Understanding Systems - Final Report. (Technical Report 3438) Volume 4. Cambridge MA: Bolt Beranek and Newman Inc., 1976.

Official Distribution List
Contract N00014-77-C-0378

	Copies
Defense Documentation Center Cameron Station Alexandria, VA 22314	12
Office of Naval Research Information Systems Program Code 437 Arlington, VA 22217	2
Office of Naval Research Code 200 Arlington, VA 22217	1
Office of Naval Research Code 455 Arlington, VA 22217	1
Office of Naval Research Code 458 Arlington, VA 22217	1
Office of Naval Research Branch Office, Boston 495 Summer Street Boston, MA 02210	1
Office of Naval Research Branch Office, Chicago 536 South Clark Street Chicago, IL 60605	1
Office of Naval Research Branch Office, Pasadena 1030 East Green Street Pasadena, CA 91106	1
Office of Naval Research New York Area Office 715 Broadway - 5th Floor New York, NY 10003	1

Naval Research Laboratory 6
Technical Information Division
Code 2627
Wasnington, D.C. 20380

Naval Ocean Systems Center 1
Advanced Software Technology Division
Code 5200
San Diego, CA 92152

Dr. A.L. Slafkosky 1
Scientific Advisor
Commandant of the Marine Corps (Code RD-1)
Washington, D.C. 20380

Mr. E.H. Gleissner 1
Naval Ship Research & Development Center.
Computation & Mathematics Dept.
Bethesda, MD 20084

Captain Grace M. Hopper 1
NAICOM/MIS Planning Branch (OP-916D)
Office of Chief of Naval Operations
Washington, D.C. 20350

Mr. Kin B. Thompson 1
NAVDAC 33
Washington Navy Yard
Washington, D.C. 20374

Advanced Research Projects Agency 1
Information Processing Techniques
1400 Wilson Boulevard
Arlington, VA 22209