A SEMANTIC EXPERT USING AN ONLINE STANDARD DICTIONARY

Jcan-Iouis Binot
B.I.M.
Kwikstraat, 4
B3078 Evcrberg, Belgium

Karen Jensen
Computer Science Department
IBM Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, New York 10598, 11 S A

ABSTRACT

A system has been developed to find the most likely attachment for prepositional phrases in English sentences in a fairly unrestricted way. The system receives as input a syntactic sentence parse provided by a general-purpose computational grammar called PEG (PLNLP English Grammar) The semantic decision that is necessary to make the right attachments is made (a) by parsing (also with PEG) the natural language definitions of an online standard dictionary, in this case Webster's Seventh New Collegiate Dictionary; (b) by relating words to other words in the dictionary; and (c) by reasoning heuristically about the comparative likelihood of different possible attachments. The basic assumption of this research is that natural language itself is a knowledge representation language that can be conveniently accessed and richly exploited. Techniques such as those presented here offer hope for eliminating the time-consuming and often incomplete hand coding of semantic information that has been conventional in natural language understanding systems.

I. INTRODUCTION

Programming a computer to make a reasonable syntactic analysis of a sentence is possible. Programming a computer to understand that sentence is more difficult, largely because the computer lacks the kind of prior experience that would enable it to "know" what it is supposed to "understand* There have been various frontal assaults on overcoming this difficulty These have generally involved the manual encoding of information in specific systems: "semantic lexicons" (e.g., Johnson 1984); "deep semantics" for a subdomain (e.g., Sager 1981); or attempts to specify world knowledge in some more exhaustive way, usually by reducing experience to formulations in predicate logic (e.g.,Lenat et al., 1986).

This line of work has produced interesting results but has also raised difficult problems. On the one hand, to encode manually the knowledge necessary for any large scale application is an enormous and possibly self-defeating task. On the other hand, encoding means choosing a specific representation, which is by no means easy given that nobody has yet a complete picture of the kind of information that must be represented. Some researchers have even challenged recently the notion that an encoding of any type could lead to any real form of natural language understanding (Winograd and Elores 1986).

We suggest another approach based on the fact that natural language words are also symbols and, as such, can be handled very well by computers. Most knowledge can be directly expressed in natural language, a fact which has already been recognized (Kayser 1982, Zadrozny forthcoming). Moreover, large bodies of knowledge, such as human dictionaries and encyclopedias, arc already available online. If computers can access that information, and use it to solve natural language processing ambiguities, they (and we) would be much further ahead.

The research presented here is aimed toward that goal As a first step, we arc concentrating on the disambiguation of prepositional phrase attachments, a difficult problem which is of considerable current interest (Lytincn 1986, Dahlgren and McDowell 1986, Schubert 1986) Wc are developing and implementing techniques for processing the definitions of an online standard dictionary in order to extract from them the semantic information necessary to resolve such ambiguities. Possible extensions include the processing of other kinds of attachment ambiguities and word meaning disambiguation

'This work addresses the areas of natural language understanding, reasoning, knowledge representation, and, briefly, learning Understanding is approached by treating online reference works as knowledge bases. A heuristic reasoning component gives robust results. One of the most interesting directions of this research is the suggestion that natural language itself is a knowledge representation language that can be conveniently accessed and richly exploited At the end of the paper, wc make a small incursion into the domain of learning, suggesting that formal knowledge bases could be progressively built by automatically extracting information from natural language sources.

The computational principles discussed here offer hope for eliminating the time-consuming, and often incomplete, hand coding of semantic information that has been conventional in natural language understanding systems.

II. PREPOSITIONAL PHRASES ATTACHMENT AMBIGUITIES

Consider the following sentences:

(1) *I ate a fish with a fork*
(2) *I ate a fish with many bones.*
(3) *I ale a fish with my fingers.*

(4) *I went to Boston by bus.*
(5) *I will have finished this work by noon.*

At the syntactic level, all these examples offer attachment ambiguities Sentences (1) to (3) were discussed in (Binot 1985); sentence (4) is discussed in (Sowa and Way 1986). However, in both cases, the attachment ambiguity was solved by hand-coded knowledge. We are designing a semantic expert which is able to decide by consulting automatically a standard dictionary what is the correct attachment for the prepositional phrase in each case, and also what is the intended meaning of the preposition. The input to this system is a syntactic parse tree for the ambiguous sentence, produced by a grammar called PEG.

The computational grammar PEG (PLNLP English Grammar) analyzes English sentences by using syntactic information in augmented phrase structure rules with a bottom-up, parallel processor (Jensen forthcoming). The system that provides these facilities is PLNLP, or the Programming language for Natural language Processing (Heidorn 1975). PEG produces useful sentence parses for a large amount, and a wide variety, of English text. These parses contain syntactic and functional (subject, object, etc.) information,

i This research was entirely performed at the IBM Thomas J. Watson Research Center.

but, as yet, no semantic information, or other information beyond the functional level.

PEG's method for dealing with attachment ambiguities is to attach modifiers in a single arbitrary pattern (usually to the closest possible head), but to mark other possible attachment sites so that they can be referred to for later semantic processing. A question mark indicates these possible attachment sites (sec Figure 1) This mark triggers the intervention of the semantic expert, which attempts to solve the ambiguity.

```
DECL    NP      PRON*   "i"
        VERB*   "ate"
        NP      DET     ADJ*    "a"
                NOUN*   "fish"
          ?     PP      PREP    "with"
                        QUANT   ADJ*    "many"
                        NOUN*   "bones"
```

Figure 1. PEG parse tree for a syntactically ambiguous sentence

///. ESTABLISHING SEMANTIC CONNECTIONS

Disambiguation is not a matter of exact deduction, but rather of finding the 'most likely* interpretation in a given context Our system does that by consulting dictionary entries for specific words in the context in order to assign to each possible interpretation a likelihood factor.*[2] These factors, like the certainty factors in MYCIN (Shortliffe 1976), may range from -I (absolute disbelief) to + 1 (absolute belief), with 0 denoting complete uncertainty. This process of assigning likelihood factors is directed by declarative rules which, making use of expert system technology, try to capture in some sense the semantic expertise brought to bear by a person while consulting a dictionary. Figure 2 shows the answer that the system gives to the problem posed by the ambiguity in Figure 1.

```
((WITH <EAT> <FORK>)
   (INSTRUMENT 0.72 )(MANNER 0.252 )
   (COAGENT -0.3XPARTOF -1. )(ASS0CIATI0N -1. ))
((WITH <FISH> <F0RK>)
   (ASSOCIATION 0.1 )(PART0F -0.3 )
   (COAGENT -I.)(INSTRUMENT -1. )(MANNER -1. ))
```

Figure 2. Solution of the semantic expert for sentence (1)

This answer is an ordered list of the possible constructs, each construct being provided with an ordered list of the possible interpretations of the preposition for that construct. The likelihood factor of a construct is defined as the likelihood factor of its best interpretation. Translated into English, Figure 2 says that there are two possible attachments for the with-PP: to the main verb "eat* (WITH <EAT> <FORK>), or to the immediately preceding noun "fish" (WITH < FISH > < FORK >). For the first construct, the INSTRUMENT meaning of the preposition receives a rating of 0.72, the MANNER meaning a rating of 0.252, etc. Thus, according to the system, the most likely interpretation of the ambiguity in example (1) is to attach 'fork" to "cat* with an instrumental meaning.

The system works by evaluating successively each possible construct and each possible meaning of the involved preposition. Due to space constraints, we shall only discuss here the INSTRUMENT and PARTOF meanings of "with,' and the INS T RUMENTI and I IMF, meanings of "by," although other meanings are handled by the system, as Figure 2 shows.

The basic strategy for evaluating a construct is to use the dictionary in order to find some semantic connection between the two terms of the proposed construct (the head and the complement). Dictionary entries contain several kinds of information that arc of interest for this aim the definition(s), example sentences and phrases, synonyms, and usage notes and other comments. In the first stage of this research, we are using only the definitions provided by Webster's Seventh New Collegiate Dictionary (W7).

Finding semantic connections in the dictionary is possible because there arc specific words and phrases in definitions, forming what we shall here call *patterns,* which arc almost systematically used to express specific semantic relations (Markowitz ct al. 1986). For the two meanings of "with" illustrated in the previous examples, some of these patterns are

PARTOF part of, arises from, end of, member of, etc
INSTRUMENT* for, used for, used to, a means for, etc

I*hesc patterns generally take, as their objects, some central term (or terms) in the definition of the complement word An attempt can then be made to link that term with the head of the construct that is being studied. We shall illustrate this with sentence (1) The system starts by looking up in W7 the definitions of the ambiguous prepositional complement "fork," and by parsing each of them The first definition reads:

fork. 1 *An implement with two or more prongs used exp for taking up (as in eating), pitching or digging.*

Fhe corresponding parse tree as produced by PEG is shown in Figure 3

The system will apply to such parse trees a pattern-matching mechanism looking, in this case, for one of the INSTRUMIvNT patterns mentioned above. Each pattern is formally represented as a recursive formula like the one shown in Figure 4:

```
(CTAKE HEAD ((SEGTYPE PP)
              (HAS PRP ((BASE FOR)))
              (CHAS HEAD ((PRESPART)))
             ))
```

Figure 4. A pattern for INSTRUMENT indications in parse trees.

Such formulas use the fact that every parse tree node and every word is represented in PENLP as a record with attributes and values. This specific pattern instructs the system to look in the parse tree for a node having a type of PP (Prepositional Phrase), having a PRP (PRcPosition) attribute pointing toward a record representing the preposition Tor," and having a HEAD attribute pointing toward a term marked as a present participle (PRESPART). If such a record is found, its head should be returned as value by the pattern-matching mechanism (or, if it is a coordinate structure, the heads of all conjuncts should be returned).

Applying this pattern to the parse tree of Figure 3, the system will find one PP node satisfying the description: "especially for taking..." The heads of the VP conjuncts inside that PP then yield three possible instrument indications: "take," "pitch," and "dig." The system will try to relate these terms with "eat" by using synonym or taxonym relations also extracted from the dictionary. (For the present, we deliberately avoid the definition phrase "as in eating" -- which offers a direct match --in order to show that our approach does not rely on such lucky coincidences.)

/ For now, we arc using only the context of the immediate sentence; but we have plans lo move to intcrscntcnlial context.

```
--------------------------------------------------------------------
NP*     DET       ADJ*     "an"
        NOUN*     "implement"
        PP        PREP     "with"
                  QUANT    AJP      ADJ*     "two"
                           CONJ     "or"
                           ADJ*     "more"
                  NOUN*    "prongs"
                  PTPRTCL  VERB*    "used"
  ?       ?       PP       AVP      ADV*     "especially"
                           PREP     "for"
                           VP       VERB*    "taking"
                                    PREP     "up"
  ?       ?       ?                 PP       PUNC     "("
                                             PREP     "as in"
                                             VERB*    "eating"
                                             PUNC     ")"
                           CONJ     ","
                           VP       VERB*    "pitching"
                           CONJ*    ", or"
                           VP       VERB*    "digging"
```

Figure 3. PFG parse tree for the first definition of fork.

Basically, the taxonym relations arc established by extracting the heads of the dictionary definitions, using the same techniques as above. In this case, 'take" is defined as a direct hypcrnym of 'cat* in W7, because one of the definitions of 'cat* is "to *take* in through the mouth..." Thus, a plausible meaningful connection has been established between 'fork" and "eat" using an INSTRUMENT link; this constitutes a strong positive indication in favor of the likelihood of the interpretation "INSTRUMENT(eat.fork)". The connection found can be viewed as forming a path in a semantic network, as illustrated in Figure 5:

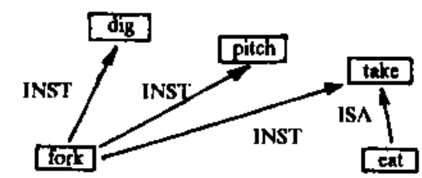

Figure 5. A semantic path connecting "cat" and "fork" in W7

It must be emphasized that this network is only an implicit structure illustrating the connections found by the system, and has not been provided as hand-coded knowledge. The main characteristic of our approach is that connections between concepts are established by tracing word relationships in the natural language entries of a standard dictionary.

No stronger indication will be found for the other possible construct "(WITH <FISH> <FORK>)," or for the other possible meanings of the preposition (in particular, there is no PARTOF pattern in the definitions of "fork" that would allow us to connect "fork" to "fish"). Thus, "INSTRUMENT(cat,fork)" will be the preferred interpretation of this example.

Let us now turn our attention to sentence (2). When evaluating the construct "(WITH <FISH> <DONE>)", the system will find the following PARTOF pattern in the definition of bone:

bone: 1. *One of the hard parts of the skeleton of a vertebrate.*

This yields two possible PARTOF indications: "skeleton" and "vertebrate." The system tries to relate each of these terms to "fish." Since "fish" is a direct hyponym of "vertebrate" in W7, we find an easy connection (shown figure 6), which gives a strong indication in favor of the interpretation "PARTOF(fish,bone)".

Again, as no shorter path will be found for any other possible interpretation, this will be the preferred interpretation.

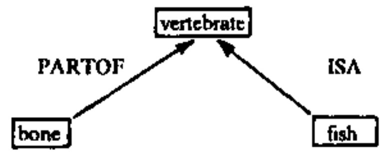

Figure 6. A semantic path connecting "fish" and "bone" in W7

IV. USING INFERENCES

In many cases, the dictionary entry for the prepositional complement does not provide the necessary information. Specific inference rules can then be tried. These rules depend on the nature of the connection that the program is trying to establish, and they will usually direct the examination of other dictionary entries.

Thus, in sentence (4), the system will not be able to find directly a useful INSTRUMENT pattern in the definitions of the prepositional complement "bus". But we can make the reasonable inference that a term can have the same instrumental uses as its hypemyms. The system will go up one level in the hierarchy, and try to find INSTRUMENT patterns for the direct hypernyms of "bus". One of these hypernyms in W7 is "vehicle", which offers the following INSTRUMENT pattern:

vehicle: 3. *a means of carrying or transporting something...*

This gives us two instrument indications: "carry" and "transport." The taxonomic information offers different ways to connect "carry" and "go", as shown in Figure 7.

The value of the likelihood factor returned by a rule is related to the length of the connection path found. Thus the instrumental interpretation will receive a weaker likelihood factor for (4) than for (1). However, once again, no stronger preference can be obtained for the other interpretations, and "INSTRUMENT(go,bus)" will be preferred.

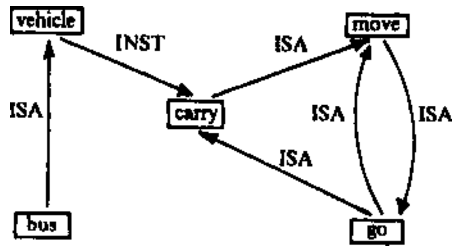

Figure 7. Semantic paths between 'bus* and "go* in W7

In sentence (3), neither "finger" nor its direct hypernyms will offer any useful INSTRUMENT pattern in their definitions. However, "finger" is defined as *"one of the five terminating members of a hand,* a "hand" is characterized as *a grasping organ* and "grasp is related to "eat" via "take." By making tie inference that if something can play an instrumental role, then a part of this something can play the same role, the system is able to establish the following connection:

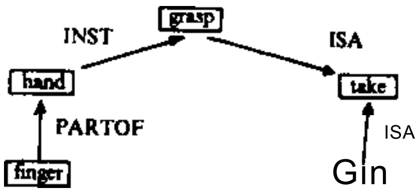

Figure 8. A semantic path connecting "finger" and "eat" in W7

This inference is at best approximate, and the connection is thus regarded as weaker than the ones in Figures 5 and 7. However, no other possibility provides a stronger connection, and *INSTRUMENT(eat,finger)* will be the preferred interpretation.

Example sentence (5) illustrates another kind of inference using a PARTOF link. The criterion for evaluating the TIME interpretation of "by" is to try to establish a connection between the complement "noon" and one of the terms "time" or "period". No direct taxonomic relation can be found. However "noon" is defined as "the *middle of the day",* and "day* is a direct hyponym of "time" in W7. By using the inference that a part of a time can also be a time, the system establishes a strong connection between "noon" and "time," which will lead to the preference of the interpretation "TIME(finish,noon)":

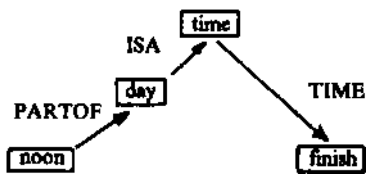

Figure 9. A semantic path connecting "finish" and "noon" in W7

V. IMPLEMENTATION OF THE DICTIONARY SEMANTIC EXPERT

A standard dictionary is plagued by incompleteness and inconsistency; moreover, the information that it carries, being expressed in natural language form, is complicated. Therefore our rules are designed as *heuristics,* which can only provide approximate results. The cumulative effect of many heuristics, and not the perfection of each one taken separately, has to do the job.

Heuristics can make use of syntactic and semantic information for determining likelihood factors (LF), as can be seen in the following English description of the rule handling the PARTOF meaning of the preposition "with" (as, for example, in the construction "a fish with bones"):

H2- checking for a PARTOF relation between a head and a "with'' complement:

1. if the head is not a noun, the relation doesn't hold (LF - -i);
2. if some PARTOF pattern exists in the dictionary definition of the complement, and if this pattern points to a defining term that can be linked with the head, then the PAR I OF relation probably holds (LF - 0.7);
3. else assume that there is more chance that the relation doesn't hold (LF - -0.3).

Each rule appears as a list of clauses, with each clause containing a condition and an action. The conditions are evaluated in sequence until one of them is satisfied. The corresponding action (indicated in parentheses in the rule description) defines the resulting likelihood factor (LP) The first clause of rule H2 is a syntactic test expressing the fact that we want only to consider a possible PARTOF meaning for 'with* when the head is a noun. The second clause directs the search for a semantic connection using a PARTOF link, as illustrated in section 3 The third clause provides a default value to be used when all other tests in the rule have failed.

Different preposition meanings may require different types of inferences in order to establish meaningful connections and may also have different syntactic cues Therefore the system includes at least one rule for each meaning of each preposition it knows about Several rules may apply to the same preposition meaning, their resulting likelihood factors will then be combined, as in MYCIN, by the formula "LF = 1 - (1- LF1)*(1-LF2)" where LF1 and LF 2 arc the factors to be combined

As an example, we give below another rule used when checking for PARTOF connections This rule uses a syntactic cue. It expresses our strong intuition that the presence in the prepositional phrase of a possessive which disagrees in person with the possible head noun precludes a PARTOF relation (as, for example, in the construction "a fish with my fingers')

H3- for checking for a PARTOF relationship between a head and a "with" complement:

1. if
   a. the head is a noun, and
   b. the complement is qualified by a possessive, and
   c. this possessive doesn't agree in person with the head,

   then the PARTOF relation probably doesn't hold (LF= -0.9),
2. else this heuristic doesn't give any indication (LF = 0).

H3 will be applied in parallel with 112. Most of the time, the result will be the neutral value zero, which does not affect the results of other rules. When the syntactic cue is present, the rule will decrease strongly the resulting likelihood factor for the PARTOF interpretation. Generally speaking, the rule structure offers a choice between sequentially and parallelism. Tests which should be performed sequentially will be included as successive clauses in a rule. Tests which should reinforce each other will be implemented as separate rules.

The condition of a clause may define subgoals, which will be evaluated by applying other rules. Thus the attempt to use taxonomic or synonymic relations to connect two terms (such as "fish" and "vertebrate" in sentence (2) is defined as a separate subgoal, which will be invoked by the second clause of 112 (this can be seen in the formal implementation of the rule, given in Figure 10 below). Different rules may be applied to this subgoal, one checking for a direct taxonomic link, another for a synonymic link, a third if the two terms have a common parent. The evaluation of the subgoal will also yield a likelihood factor which, if positive, will be multiplied by the one expressed in the action part of the rule, thereby decreasing it.

The choice of likelihood factors rests mainly on intuition. Some choices are easy; some inferences, for example, are obviously weaker than others. In other cases the values have to be adjusted by trial and error, by processing many examples. It is interesting to note that, as our corpus of examples increases, the likelihood factors are converging toward apparently stable values.

The resulting likelihood factor of an interpretation must also be inversely tied to the length of the semantic path corresponding to that interpretation. This is basically accomplished by designing the rules so that a given rule will search only for one link in a path, the other links being identified by subgoals called by that rule. The longer the path, the longer the reasoning chain, and the lower the final likelihood factor.

The rules are written in the PLNI,P language (Heidorn 1975), which we shall not describe here. We shall only give below, as an illustration, the formal version of rule 112:

```
H2(CONTROLER*PTR,COMPLEMENT*PTR,
   <SEGTYPE(CONTROLER).NE.'NOUN',
      <-LF<'PARTOF',]"-1.0"»,
   <SUBG0AL2<'PREF'...CONTROLER...
                ENTRIES<COMPLEMENT,'PARTOF?>...
                SEGTYPE(CONTROLER)...2>,
      <-LF<'PARTOF\]"0.8"»,
   <-LF<'PARTOF']"-0.3">)
```

Figure 10. PLNLP formulation of rule H2

Our prototype system includes currently twenty rules and is able to handle the prepositions "with", "by", "after* and "in." It has been tested on about 50 examples so far.

VI. LEARNING USEFUL FACTS

Re-parsing all the definitions of a word every time that word is encountered will soon lead in practice to unsurmountablc efficiency problems. Human beings do not look up repetitively the same word in the dictionary Eventually, they will learn something about its meaning

We have included in our system a simple but helpful form of learning which consists of remembering useful information extracted from dictionary definitions. Each time the pattern-matching mechanism identifies semantic relations, they are stored in a "memory file" where they will remain available for processing further input. If the same word occurs again later on, the system will check to see if the remembered relations are useful before attempting to parse the definitions. Figure 11 below shows some of the information remembered for sentences (1) and (2).

```
(PUT 'FORK  'PARTOF      'NONE)
(PUT 'FORK  'INSTRUMENT  '(TAKE PITCH DIG))
(PUT 'BONE  'PARTOF      '(SKELETON VERTEBRATE))
(PUT 'BONE  'INSTRUMENT  '(STIFFEN))
```

Figure 11 Information learned from the definitions of "fork" and "bone"

The remembering mechanism offers two significant advantages. First, it can speed up significantly the disambiguation process. Second, it provides an interesting tool for finding out what information is useful and how to store it. Obviously, however, there is still room for considerable improvement in this learning process. In particular, it does not attempt to infer new information, induce new concepts, or even to structure the information in more useful ways. In the long term, this line of research could lead to the automatic building of formal knowledge bases from natural language sources.

VII. CONCLUSIONS

We have shown that it is possible to consult automatically a standard machine-readable dictionary as a knowledge base, in order to resolve ambiguities in the computational parsing of natural language text. Only prepositional phrase attachment ambiguities have been studied so far, but the techniques described here should be extendible to ambiguities that involve all other classes of phrases and clauses. Furthermore, there is no need to stop with dictionaries: encyclopedias and other online written works are available, valuable sources of information.

All types of information present and available in dictionary entries should eventually be used to help in the disambiguation process: examples, synonyms, and usage notes, as well as the definitions themselves. Other important areas for future development include the development of new heuristics; the adjustment and improvement of the likelihood factors; the analysis of relationships that are signaled by prepositions, conjunctions, and other function words; and also the addition of information on phrasal verbs (verb-particle pairs, verb-preposition pairs), and an investigation of how these will help in the parsing process.

This line of research highlights several principles of computer cognition:

- Knowledge is represented in natural language.
- Natural language text can be accessed by a syntactic parser and a reasoning module (in this case, heuristic reasoning).
- Heuristics can be used to help the program reach new levels of understanding.

The work should benefit all areas of natural language processing -- from text analysis to machine translation to data base query -- by opening up the prospect of a genuinely broad-coverage approach to meaning and understanding.

ACKNOWLEDGMENTS

Thanks are due to George Heidorn for the many useful discussions that we had with him during this research and for helpful comments on a previous draft of this paper.

REFERENCES

1. Binot, Jean-Louis. 1985. *SABA: vers un systeme portable danalyse du francais ecrit.* Ph.D. dissertation, University of Liege, Liege, Belgium.
2. Dahlgren, K. and J. McDowell. 1986. "Using commonsense knowledge to disambiguate prepositional phrase modifiers" *Proceedings of AAAI-86,* Philadelphia, Pennsylvania, August 1986.
3. Heidorn, George E. 1975. "Augmented phrase structure grammars" in Nash-Webber and Schank, eds., *Theoretical Issues in Natural Language Processing.* Association for Computational Linguistics.
4. Jensen, Karen, forthcoming. "PEG 1986: A broad-coverage computational syntax of English." IBM Research Report.

   Johnson, David. 1984. "Design of a Robust, Portable Natural I.anguage Interface Grammar." IBM Research Report RC 10867, December 5, 1984.
5. Kayser, Daniel. 1982. "An experiment in knowledge representation", ECAI82, pp. 119-124.
6. Lcnat, Doug, Mayank Prakash, and Mary Shepherd. 1986. "CYC: Using Common Sense Knowledge to Overcome Brittleness and Knowledge Acquisition Bottlenecks." *The AI Magazine,* Winter 1986, pp.65-85.

7. Lytinen, Steven L. 1986. "Dynamically combining syntax and semantics in natural language processing." *Proceedings of AAAI-86,* Philadelphia, Pennsylvania, August 1986.
8. Markowitz, Judith, Thomas Ahlswede, and Martha Evens. 1986. "Semantically significant patterns in dictionary definitions." *Proceedings of the 24th Annual Meeting of the ACL,* Columbia University, June 1986.
9. Sager, Naomi. 1981. *Natural language Information Processing: A Computer Grammar of English and Its Applications.* Addison-Wesley Publishing Co., Inc.
10. Schubert, Lenhart K. 1986. "Are there preference trade-offs in attachment decisions?" *Proceedings of AAAI-86,* Philadelphia, Pennsylvania, August 1986.
11. Shortliffe, EH. 1976. *Computer-based medical consultation: MYCIN.* Artificial Intelligence Series. Elsevier.
12. Sowa, J.F. and E.C. Way, Implementing a semantic interpreter using conceptual graphs, *IBM Journal of Research and Development* 30:1, 1986.
13. Winograd, T. and Flores, F. 1986. *"Understanding computers and cognition"*, Ablex Publishing Co.
14. Zadrozny, Wlodzimierz. forthcoming. "Intended models, circumscription and commonsense reasoning." IBM Research Report.