Using Commonsense Knowledge to Disambiguate Prepositional Phrase Modifiers

Dr. K. Dahlgren

IBM Los Angeles Scientific Center
11601 Wilshire Blvd.
Los Angeles, CA 90025

J. McDowell

University of Southern California
Los Angeles, CA 90089

Abstract

This paper describes a method using commonsense knowledge for discarding spurious syntactic ambiguities introduced by post-verbal prepositional phrase attachment during parsing. A completely naive parser will generate three parses for sentences of the form NP-V-Det-N-PP. The prepositions alone are insufficiently precise in meaning to guide selection among competing parses. The method is imbedded in the Kind Types System (KT) which employs commonsense knowledge of concepts, including prototype and inherent features (generic information) and ontological classifications. The generic information is drawn from published psycholinguistic studies on how average people typically view the world. This method is employed in preference strategies which appeal to the meaning of the preposition combined with information about the verbs and nouns associated with it drawn from the text and from the generic and ontological databases. These determine which syntactic structures generated by a semantically naive parser are commonsensically plausible. The method was successful in 93% of cases tested.

1. Semantically Implausible Syntactic Ambiguities

A problem for text understanding systems is that syntactic rules alone produce numerous ambiguities, many of which are not semantically possible (or likely) interpretations. Consider sentence (1), for which any standard parser would produce three distinct syntactic structures. Figure 1 is a syntactic tree showing the parse for (1) in which the key belongs to the lock. The *with*-phrase is a constituent of the noun phrase headed by *lock* (NP constituency). Figure 2 displays the parse in which the *with*-phrase is a constituent of the verb phrase (VP constituency). Figure 3 shows the parse in which the *with*-phrase modifies the sentence (S-modification), so that the event of buying the lock takes place with the key. Only one of these syntactic possibilities is semantically possible for (1), namely the one in which the prepositional phrase is a complement of the NP whose head is *lock*. Similarly, only VP constituency is semantically possible for (2), and only S-modification for (3).

1) John bought the lock with the key.
2) John bought the lock with five dollars.
3) John bought the lock in the afternoon.
4) John took the key to the lock.

Clearly, the semantically impossible syntactic ambiguities generated for (1)-(3) are spurious. On the other hand, some syntactic ambiguities correspond to possible semantic ambiguities. In sentence (4), both the VP constituency and NP constituency parses are semantically possible. It is easy to imagine a situation in which John physically carries the key over to the lock. However, in this case the preferred reading maps to NP constituency because the head of the *to*-phrase is typically 'a part of' or 'used for' the head of the direct object NP. A text understanding system that can guess NP constituency in this case is not only practical and workable, it is also superior to one which chooses randomly. The commonsense disambiguation method to be described in this paper assigns constituency for prepositional phrases according to commonsense preference, and the only ambiguities which remain

Figure 1 NP Constituency

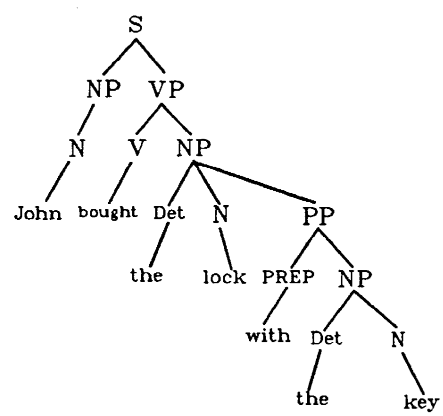

Figure 2 VP Constituency

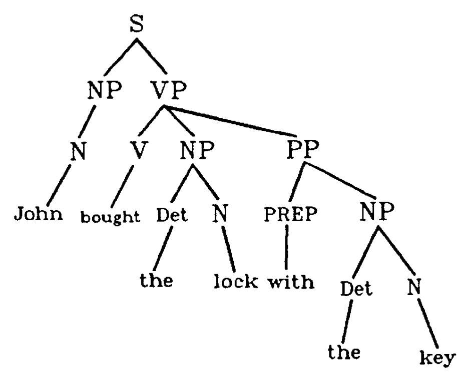

Figure 3 S – Modification

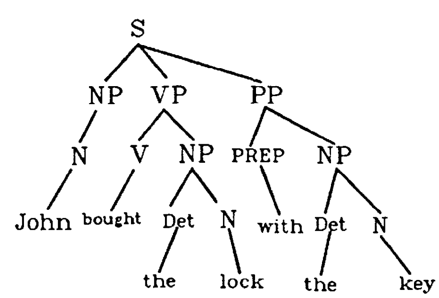

after the preference strategy has been invoked are the semantically and commonsensically possible ambiguities, such as those in (5).

5) John saw the man in the park.
John built the houses by the sea.

(1)-(4) are all of the form NP-V-NP-PP. The same considerations apply when there are multiple PPs. Consider sentence (6), which has eight parses, of which only one is semantically possible (*to the pasture* is a VP constituent, and *in the afternoon* modifies the S).

6) The boy took the cow to the pasture in the afternoon.

2. Using Commonsense Knowledge to Disambiguate

One solution to the problem of spurious ambiguities is semantically-driven parsing, which forces you to give up the speed and parsimony of autonomous parsing (Arens, 1981, Wilks, 1985). Another solution is to ask the user to disambiguate, as in Tomita (1985). This works well in the database querying environment, but not for text understanding, where human intervention is not feasible. A third possibility, to be described here, is to use the knowledge already needed to understand the text to eliminate spurious parses. Exemplifying the method with sentences (1) - (3), in (1), the *with*-phrase must be a NP constituent because locks typically have keys. English speakers know this, and that is why (1) is unambiguous. In (2), the *with*-phrase is a VP constituent because *buy* is a verb of exchange. In (3), the fact that *afternoon* is a temporal noun forces the interpretation in which the PP modifies the S because only events have a temporal argument.

The method described here uses commonsense knowledge associated with concepts to choose among possible parses for a sentence with a prepositional phrase to the right of the verb. (Prepositional phrases in the subject of the sentence are not ambiguous in the same way, because PPs after the head of the subject noun must be NP constituents. We assume that preposed PPs as in "In the spring, we go dancing" show up at the end of the S after parsing.) A text understanding system can eliminate spurious parses by employing preference strategies in the spirit of Warren (1981). The disambiguation method is independent of the type of parser and grammar. After each parse is generated, and before semantic interpretation, word-level commonsense knowledge is employed to decide whether the parse is preferred. The knowledge used here derives from empirical psycholinguistic studies of prototypes associated with common nouns and with verbs (Rosch et al, 1976, Graesser, 1985, Ashcraft, 1976, Dahlgren, 1985). In order to make the system seem to understand the text, this knowledge is needed anyway (Hayes 1985, Hobbs et al, 1985, Lenat et al, 1986). We found that in limited text understanding, where the grammaticality of the text can be assumed, the level of detail called for in Waltz (1981), Gawron (1982), and Herskovits (1985) is not necessary.

This paper will first outline the text understanding system and its use of commonsense knowledge, then describe how this knowledge is used in preference strategies for prepositional ambiguities, and finally list the preference strategy rules and discuss their implementation.

3. Commonsense Knowledge in the Kind Types System (KT)

The Kind Types (KT) system (Dahlgren and Mcdowell, 1986) is written in VM/PROLOG. KT reads geography text using an IBM-internal parser and logic translator (Stabler & Tarnawsky), and partially understands it because of commonsense knowledge in its Ontology and Generic Lexicon. It can answer queries both of what the text says directly, and of information the typical speaker of English would infer from the text because of commonsense knowledge. KT's representations are based upon a theory which identifies lexical meaning with commonsense knowledge of concepts, so that there is no difference in form of representation between word meanings and encyclopedic knowledge. This theoretical basis is shared with other approaches such as KL-ONE (Brachman and Schmolze, 1985), which differ in employing representations which are richer than first order logic.

In addition to a standard syntactic lexicon, KT employs commonsense knowledge encoded in PROLOG axioms in several databases. The Ontological Database encodes basic distinctions such as REAL vs ABSTRACT. The Generic Lexicon lists prototypical and inherent features of nouns and verbs. The Feature Typing Database classifies features as colors, etc. The Kind Types Database lists constraints on feature types associated with kind types such as persons, artifacts, etc.

The Ontological Database reflects top-level category cuts in commonsense knowledge. It is an original cross-classification scheme, taking into account previous work (Keil, 1979, Tenenbaum, 1985). The ontology is represented as PROLOG axioms, and every common noun and verb which KT recognizes is attached at one of its leaves in a rule like (7). Using ontological information, the KT system deduces taxonomically inherited information about the sorts mentioned in the text, as described below.

The Generic Lexicon contains prototypical and inherent features of common nouns which are taken from empirical studies of prototypes. The entry for chicken is shown in (8). The first argument is a list of prototype features, and the second a list of inherent features.

```
(7)  bird(X)  ← chicken(X).

(8)  chicken(  { white.scratch(X).farm.
                 meat.eggs.roost(X)  } ,
               { lay(X,Y) & egg(Y).handleable.
                 haspart(legs,2).cluck(X)  } ).
```

The features with variables are logic translations of verb phrases. The Feature Typing database types every feature which occurs in the Generic Lexicon as a color, size, behavior, function, and so on. Thirty-three feature types account for all of the features.

The text in (9) is representative of the range of vocabulary and syntax KT works with.

9) *John is a miner who lives in a small town. John digs for coal. John raises a chicken which lays brown eggs.*

The text is translated into PROLOG axioms to form a Textual Database. Queries to the system are translated into PROLOG goals. Queries are answered by a problem solver which consults the Textual, Ontological, Generic and Typing databases. For a higher-level question such as "What color is the coal?", KT finds the generic information for coal (by Prolog predicate matching). Then it compares all the colors in the Feature Typing Database to the list of generic features and finds that "black" is an inherent feature of coal. For a lower-level question, such as "Is the coal black?", only the generic information is accessed. The system identifies higher-level questions by looking at a list of higher-level predicates (the types). The answers to queries can come from any of these sources, as illustrated in (10).

(10) **Textual Database**

Is John a miner?
-- Yes
Where does John live?
-- In a small town

Ontological Database

What is a miner?
-- A miner is a role, sentient, concrete, social, individual, real, and an entity.
What is digging?
-- A motion, goal oriented, temporal, individual, and real.

Generic Database

Is the miner rugged?
--Probably so.
Does the chicken lay eggs?
--Inherently so.
Does John use a shovel?
--Probably so.

Typing and Generic Databases

What color is the coal?
--Typically black.
What size is the chicken?
--Inherently handleable.

4. Commonsense Knowledge used in the Preference Strategy

Our preference strategy assigns PP constituency according to information from several sources: syntactic information about the preposition (PREP) and the verb (V); information from the Ontological Database about the V, direct object (DO), and object of the preposition (POBJ); and information from the Generic Database about the V, DO and POBJ. Table 1 lists the preference rules preposition by preposition.

4.1 Ontological Class of Object of the Preposition

While it is possible for the POBJ to belong to any ontological class, membership in a small set of such classes restricts the PP-assignment possibilities for many prepositions. For example,

If the POBJ is temporal, the PP modifies the S.

in the morning, for six days

If the Prep is *by*, the POBJ is sentient and the DO is propositional, the PP modifies the NP.

the book by Chomsky

If the Prep is *at* or *in* and the POBJ is abstract, the PP modifies the S.

at once, in detail

If the Prep is *on*, and the POBJ is propositional, the PP modifies the NP.

the book on love

Sometimes it is necessary to consider not only the ontological class of the object of the preposition, but also the ontological class of the DO. For example, in *the report by the committee* it is necessary to know not only that *report* is PROPOSITIONAL, but that *committee* is SENTIENT. Other ontological classes which play a role in the preference strategies are PLACE, EMOTION, ROUTE, MEASURE, RELATION, and DIRECTION. A global rule assigns locative and directional PPs to the VP as in (11), though later specific prepositional rules may assign them as as S modifiers if the global rule fails as in (12).

11) John put the book in the living room.
12) John read the book in the living room.

PLACEs can be social places (*factory*, *hospital*) or natural places (*valley, mountain*). EMOTIONS enter into PPs in such phrases as *under duress, in fear, from hatred, with courage,* etc. ROUTES are terms like *way, road, path.* MEASUREs appear in PPs with *to* (*to a degree, to a point*). DIRECTIONS figure prominently in physical descriptions (*to the North, on the South*). PPs headed by *with* and *without* are NP constituents if the DO is a RELATION (*connection with, contact with*).

4.2 Ontological Class of The Direct Object

Much less crucial is information about the ontological class of the DO. As described above, the fact that a DO is PROPOSITIONAL is important only in the case of two classes of prepositional objects and then only for certain prepositions. If the DO is a MEASURE, then the PP is an NP constituent (*enough food for their needs, much about the world*).

4.3 Ontological Class of Verb

The nature of the verb itself can sometimes induce a preference for a PP assignment without reference to any other information. For instance, the verbs *be* and *stand* and other intransitive STATIVES like them automatically take any PP as a VP constituent (*be on time, stand in the rain*). For this reason the global rule for statives must precede the global rule assigning all temporal phrases as S modifiers. Mental verbs force VP constituency for PPs headed by *of, for,* and *about* (Gawron, 1982), as illustrated by (13) as opposed to (14).

13) John thought of his sweetheart while he waited.
14) John repaired the roof of the house.

Verbs of exchange, such as *buy*, typically take three arguments, the object exchanged, the source, and the goal. In addition, the medium of exchange and the benefactee of the exchange can be expressed. With such verbs, PPs headed by *for, from,* and *with,* are VP constituents.

4.4 Generic Information

So far just the ontological classification of the verb and the NPs to its right have been considered. The Generic Database was encoded originally in order to describe prototypical and inherent features of objects for purposes of understanding the meaning of text. It was found to be useful as well in addressing the problem of PP attachment. The generic relations between nouns illustrated in this section are encoded in the Generic Database for nouns. In the case of *with* and *on* it is sufficient that the POBJ is mentioned in the prototype description for the DO (*car with a wheel, birds with feathers, hair on the head*).

A locative relationship can also be part of generic information. A stove is typically in the kitchen, a book on a shelf, etc. The typical MATERIAL from which an object is made is included in the generic information for physical objects. A house is made from wood, a window from glass. We call SOLVERS the nouns which are typically associated with other nouns by means of *to* (*the key to the lock, the answer to the problem*). SIZE is a generic feature and is encoded in terms of 13 reality-related size ranges which have mnemonic names such as microscopic, person-sized, etc. An object must be of the order of at least two sizes larger than the subject of a sentence in order for it to be a suitable location for an action by the subject. That is, a PP headed by a locative preposition will be a S modifier if the prepositional object is a suitable location for the action of the subject as in (15) and otherwise it will be a VP or NP constituent as in (16).

15) John mixed the salad in the kitchen.
16) John mixed the salad in the salad bowl.

The relation between a verb and the generic INSTRUMENT with which the verbal action is carried out is encoded in the Generic Database for verbs along with other information, such as selectional restrictions on verb arguments. In (17), because *knife* is a typical INSTRUMENT for the verb *cut*, the PP is assigned to the VP, but it is assigned to the S in (18).

17) John cut the loaf with a knife.
18) John cut the loaf with glee.

4.5 Syntax

Certain syntactic constructions can also force PP interpretation. There is a large class of intransitive verbs which are ill-formed unless accompanied by certain prepositions (*depend on, look for, make up, get along with, revolve about, cooperate with, turn to, divide into, provide for*). These are conventional co-occurrence requirements, and they force the PP to be interpreted as a VP constituent. Adjectives + PP require attachment to the phrase (XP) containing the adjective (*suitable for a child, useful for parsing*). On the other hand, the comparative construction forces NP constituency (*the largest book in the library, uglier than the man at the desk*).

We showed above that the ontological class of the POBJ can determine correct PP constituency. Generics and names are syntactic classes and also guide PP parsing.

19) John read the book on Peter the Great.
20) John read the book on dogs.

For each preposition there is a general rule which takes effect when all the specific rules fail. This is a kind of Elsewhere Condition for the syntax of prepositions. Generalizing the results of the study, locative PPs are NP constituents; directional PPs are VP constituents; and time/manner/place PPs are S modifiers. However, whether or not a PP falls into one of these classes is a function of the prepositional head in combination with the verb and the prepositional object and not of the preposition alone.

5. Success Rate of the Preference Strategy

The PP rules were developed intuitively by considering the interaction of each preposition with one, two and three-place verbs. Then the PP rules were hand-tested on 4500 words of schoolbook geography texts, the original corpus upon which the lexicon and ontology in the KT system were built. The PP-attachment rules were developed independently of these texts, but the success rate was a surprising 100%. Then, as a check against these results, the rules were hand-tested on a second group of three short sample texts. These were (1) a 481-word excerpt from a novel, (2) a 415-word excerpt from a work on the history of science, and (3) a 409-word excerpt from a technical article. We assumed parser translation of the texts into strings of the form NP-V-NP-PP for submission to our rules. We also ignored passive *by*-phrases because the parser recognizes them as distinct from ordinary PPs. On the latter three texts the rules were tested as if the vocabulary was resident in our system. The overall success rate for the second group is 93%. The failures in these tests are of two types. The first type of failure was in idiomatic phrases, most of which have the function of asides or sentence qualifiers (*at all, in effect, in every case, under my eyes, in particular, according to*). We do not view this as a defect in our system since any system must be able eventually to deal with idiomatic phrases. The second type of failure was outright failure of the rules. If we ignore the failures due to idiomatic phrases then the average success rate for the second group is much higher, 98%. One reason why the success rate is so high is the high occurrence rate of *of*-phrases. These constitute 32% of the second group. In every case we have seen so far, they attach to the NP immediately to the left.

6. Implementation

Preference rules for thirteen English prepositions are listed in Table 1. First, seven global rules are attempted. If none of these rules applies, the procedure relevant to the preposition is called upon. Although there is no single, general algorithm for assigning constituency for prepositions, three points compensate for this lack of generality: the set of prepositions in English is a closed and small set, some rules are used for several prepositions, and for each preposition, the list of rules is short (usually three). The phrase structures which are input to the rules are: VP(V-DO-Prep-POBJ), VP(V-Prep-POBJ), VP(V-comparative) VP(V-Adj-Prep-POBJ) and VP(V-Prep-NP).

The seven global rules are listed below. Lexical (V + Prep) means that the relationship between the verb and the preposition is lexical, as described in Section 4.5. Stative (V) means the verb is stative, measure(DO) that the direct object is a measure. Adj and comparative mean that such a construction occurs in the sentence.

1. lexical(V+Prep) --> vp_attach(PP)
2. stative(V)--> vp_attach(PP)
3. time(POBJ) --> s_attach(PP)
4. xp(...Adj-PP...) --> xp_attach
5. measure(DO) --> np_attach(PP)
6. motion(V) & DO & endofclause --> vp_attach(PP)
7. comparative --> np_attach

To illustrate the application of the rules, consider the rule for *of*. applied to the sentence "John buys the book of poems". The global rules are tried, and they fail. Then the first *of*-rule consults the Ontological Database to see whether the verb is mental. This fails, so the solution is NP constituency.

The *with*-rule illustrates more complex reasoning. Consider the sentence "Sam bought the car with the wheel." The first *with*-rule consults the entry for *car* in the Generic Lexicon, looking for mention of *wheel* there, and finds it, as cars inherently have wheels. The rule succeeds and the PP is assigned NP constituency. In contrast, consider the sentence "Sam fixed the car with a wrench". The global rules fail, and the first *with*-rule tests whether a generic relationship exists between the DO and the POBJ, in the Generic Lexicon, and whether the DO is a relation in the Ontological Database, and fails. The next *with*-rule checks whether *wrench* is a typical instrument of the verb *fix* in the Generic Lexicon. This succeeds, so the PP is assigned VP constituency. Finally, consider the sentence "Sam fixed the car with Mary." No generic relation can be found between *car* and Mary or *fix* and Mary, so the elsewhere rule applies, and the PP modifies the S.

These generic relationships exist for a number of prepositions but are not mentioned in the rules because they are subsumed by the elsewhere condition. For example, such relationships exist for uses of *for*, as in *the wheel for the car* and *the cap for the jar*, but since the rules are written so that NP-attachment is the elsewhere rule, this kind of relationship does not show up directly.

In the *in*-rules, notice that first the generic relation of location (a place the DO is typically found) is checked for in the Generic Lexicon. If that fails, and VP constituency fails, a check is carried out in the Ontological Database for whether the POBJ is a place. This order captures the difference between (21) and (22). Our system chooses S modification for (23), but it is actually ambiguous.

21) John saw the horse in the barn.
22) John walked the horse in the city.
23) John saw the horse in the city.

The rules work for constructions which have no DO. There are several types of these. One is the type where the verb must always cooccur with a certain preposition (*depend on*). These are covered by the first global rule, which checks for such constraints in the Syntactic Lexicon. Another type is STATIVE verbs, as in "John lives in the house", which are covered by the second global rule. Notice that intransitive constructions are excluded from the sixth global rule which assigns VP constituency for sentences such as "John put the book on the table". This means that the rules will prefer S-attachment in some cases where sentences are commonsensically ambiguous, as in (24).

24) John ran at the woman.
John ran by the park.

Conclusion

The preference strategy presented here can be applied to the output of any type of parser, and the commonsense knowledge can be represented in any language desired. The content of the knowledge derives from available empirical studies. Thus the method is broadly applicable. The method interfaces autonomous syntactic and semantic components of a natural language understanding system, discarding implausible syntactic trees before they are fully interpreted semantically. It is also possible to apply the preference strategy during the parse, by first generating all the possible places to attach a PP, and looking ahead to parse the object of the PP. At this point all of the information needed by the preference strategy is available, and the rules can be applied, thus eliminating the expense of generating parses only to discard them later.

Table 1

of-rules
mental(V) --> vp_attach(PP)
Elsewhere --> np_attach(PP)

on-rules
(location(DO,POBJ) OR generic(POBJ) OR name(POBJ)
OR (propositional(DO) & abstract(POBJ)) -> np_attach(PP)
Elsewhere --> s_attach(PP)

for-rules
place(POBJ) OR sentient(POBJ) OR mental(V)
OR exchange(V) --> vp_attach(PP)
distance(POBJ) --> s_attach(PP)
Elsewhere --> np_attach(PP)

at-rules
abstract(POBJ) OR place(POBJ) --> s_attach(PP)
Elsewhere --> np_attach(PP)

in-rules
abstract(POBJ) OR emotion(POBJ) OR place(POBJ)
--> s_attach(PP)
Elsewhere --> np_attach(PP)

by-rules
location(DO,POBJ) --> np_attach(PP)
propositional(DO) & sentient(POBJ) --> np_attach(PP)
Elsewhere --> s_attach(PP)

under-rules
twosizeslarger(POBJ,SUBJ) OR propositional(POBJ)
--> s_attach(PP)
Elsewhere --> np_attach(PP)

about-rules
mental(V) OR motion(V) --> vp_attach(PP)
Elsewhere --> np_attach(PP)

to-rules
solver(DO, POBJ) OR route(DO) --> np_attach
geometric(V) & direction(POBJ) --> s_attach(PP)
place(DO) & direction(POBJ) --> np_attach(PP)
measurement(POBJ) --> s_attach
Elsewhere --> vp_attach(PP)

with and **without**-rules
partof(DO,POBJ) OR relation(DO) --> np_attach(PP)
instrument(POBJ,V) --> vp_attach(PP)
Elsewhere --> s_attach(PP)

from-rules
material(POBJ) OR emotion(POBJ) --> s_attach(PP)
exchange(V) --> vp_attach(PP)
Elsewhere --> np_attach(PP)

through-rules
Elsewhere --> s_attach(PP)

ACKNOWLEDGEMENTS

We gratefully acknowledge the suggestions of Edward P. Stabler and the support of Juan Rivero.

REFERENCES

(1) Arens, Y. Using language and context in the analysis of text. *IJCAI*, Vancouver, (1981).
(2) Ashcraft, M.H. Property norms for typical and atypical items from 17 categories *Memory and Cognition* 6:3 (1976) 227-232.
(3) Dahlgren, K. The cognitive structure of social categories. *Cognitive Science* 9 (1985) 379-398.
(4) Dahlgren, K. and J. McDowell. Kind types in knowledge representation. *COLING86*, Bonn (1986).
(5) Gawron, J.M. Lexical representations and the semantics of complementation. Dissertation. University of California, Berkeley, (1983).
(6) Graesser, A. and L. Clark. *Structure andProcedures of Implicit Knowledge* . Ablex (1985).
(7) Hayes, P.J. The second naive physics manifesto. *Formal Theories of the Commonsense World.* Ed. J.R. Hobbs and R.C. Moore. Ablex (1985).
(8) Herskovits, A. Semantics and pragmatics of locative expressions. *Cognitive Science*9 (1985) 341-378.
(9) Hobbs, Blenko, Croft, Hager, Kautz, Kube and Shoham. *Commonsense Summer: Final Report.* Stanford, California: CSLI Report (1985).
(10) Keil, F. C. *Semantic and Conceptual Development* . Harvard U Press (1979).
(11) Lenat, D., M. Prakash, and M. Shephard. CCY: Using common sense knowledge to overcome brittleness and knowledge acquisition bottlenecks. *AI Magazine* 4 (1986) 65-85.
(12) McCord, M. The lexical base for semantic interpretation in a prolog parser. Workshop on the Lexicon, Parsing and Semantic Interpretation, CUNY Graduate Center (1985).
(13) Rosch, Mervis, Gray, Johnson, D.M. and Boyes-Braem. Basic objects in natural categories. *Cognitive Psychology* 8 (1976) 382-439.
(14) Stabler, E.P., Jr., and G.O. Tarnawsky. NL/Prolog---A prolog-based natural language facility. To appear.
(15) Tenenbaum, J.D. Taxonomic reasoning. *Proc. IJCAI*, Los Angeles (1985).
(16) Tomita, M. An efficient context-free parsing algorithm for natural languages. *Proc. IJCAI*, Los Angeles (1985).
(17) Waltz, D. L. Toward a detailed model of processing for language describing the physical world. *Proc. IJCAI*, Los Angeles (1981).
(18) Warren, D.H.D. Efficient processing of interactive relational database queries expressed in logic. *7th Intl.Conf. on Very Large Databases*, Cannes, France (1981).
(19) Wilks, Y., X. Huang and D. Fass. Syntax, preference and right attachment.*IJCAI*, Los Angeles (1985).