WORD EXPERT PARSING¹

Steven L. Small

Department of Computer Science
University of Maryland
College Park, Maryland 20742

This paper describes an approach to conceptual analysis and understanding of natural language in which linguistic knowledge centers on individual words, and the analysis mechanisms consist of interactions among distributed procedural experts representing that knowledge. Each word expert models the process of diagnosing the intended usage of a particular word in context. The Word Expert Parser performs conceptual analysis through the interactions of the individual experts, which ask questions and exchange information in converging on a single mutually acceptable sentence meaning. The Word Expert theory is advanced as a better cognitive model of natural language understanding than the traditional rule-based approaches. The Word Expert Parser models parts of the theory, and the important issues of control and representation that arise in developing such a model form the basis of the technical discussion. An example from the prototype LISP implementation helps explain the theoretical results presented.

1. Introduction

Computational understanding of natural language requires complex interactions among a variety of distinct yet redundant mechanisms. The construction of a computer program to perform such a task begins with the development of an organizational framework which inherently incorporates certain assumptions about the nature of these processes and the environment in which they take place. Such cognitive premises affect profoundly the scope and substance of computational analysis for comprehension as found in the program.

This paper describes a theory of conceptual parsing which considers knowledge about language to be distributed across a collection of procedural experts centered on individual words. Natural language parsing with word experts entails several new hypotheses about the organization and representation of linguistic and pragmatic knowledge for computational language comprehension. The Word Expert Parser [1] demonstrates how the word expert view, coupled with certain other choices based on previous work, affect structure and process in a cognitive model of parsing.

The Word Expert Parser is a cognitive model of conceptual language analysis in which the unit of linguistic knowledge is the word and the focus of research is the set of processes underlying comprehension. The model is aimed directly at problems of word sense ambiguity and idiomatic expressions, and in greatly generalizing the notion of word sense, promotes these issues to a central place in the study of language parsing. Parsing models typically cope unsatisfactorily with the wide heterogeneity of usages of particular words. If a sentence contains a standard form of a word, it can usually be parsed; if it involves a less prevalent form which has a different part of speech, perhaps it too can be parsed. Distinguishing among the many senses of a common verb, adjective, or pronoun, for example, or correctly translating idioms are rarely possible.

At the source of this difficulty is the reliance on rule-based formalisms, whether syntactic or semantic (e.g., cases), which attempt to capture the linguistic contributions inherent in constituent chunks of sentences that consist of more than single words. A crucial assumption underlying work on the Word Expert Parser is that the fundamental unit of linguistic knowledge is the word, and that understanding its sense or role in a particular context is the central parsing process. In the parser to be described, the word expert constitutes the kernel of linguistic knowledge and its representation the elemental data structure. It is procedural in nature and executes directly as a process, cooperating with the other experts for a given sentence to arrive at a mutually acceptable sentence meaning.

Certain principles behind the parser do not follow directly from the view of word primacy but from other recent theories of parsing. The cognitive processes involved in language comprehension comprise the focus of linguistic study of the word expert approach. Parsing is viewed as an inferential process where linguistic knowledge of syntax and semantics and general pragmatic knowledge are applied in a uniform manner during

interpretation. This methodological position closely follows that of Riesbeck (see [2] and [3]) and Schank [4]. The central concern with word usage and word sense ambiguity follows similar motivations of Wilks [5]. The control structure of the Word Expert Parser results from agreement with the hypothesis of Marcus that parsing can be done deterministically and in a way in which information gained through interpretation is permanent [6]. Rieger's view of inference as intelligent selection among a number of competing plausible alternatives [7] of course forms the cornerstone of the new theory. His ideas on word sense selection for language analysis ([8] and [9]) and strategy selection for general problem solving [10] constitute a consistent cognitive perspective.

Any natural language understanding system must incorporate mechanisms to perform word sense disambiguation in the context of open-ended world knowledge. The importance of these mechanisms for word usage diagnosis derives from the ubiquity of local ambiguities, and brought about the notion that they be made the central processes of computational analysis and understanding. Consideration of almost any English content word leads to a realization of the scope of the problem -- with a little time and perhaps help from the dictionary, many distinct usages can be identified. As a simple illustration, several usages each for the words "heavy" and "ice" appear in Figure 1. Each of these seemingly benign words exhibits a rich depth of contextual use. An earlier paper contains a list of almost sixty verbal usages for the word "take" [11].

The representation of all contextual word usages in an active way that insures their utility for linguistic diagnosis led to the notion of word experts. Each word expert is a procedural entity cognizant of all possible contextual interpretations of the word it represents.² When placed in a context formed by experts for the other words in a sentence, each expert should be capable of sufficient context-probing and self-examination to determine successfully its functional or semantic role, and further, to realize the nature of that function or the precise meaning of the word. The representation and control issues involved in basing a parser on word experts are discussed below, following presentation of an example execution of the existing Word Expert Parser.

2. Model Overview

The Word Expert Parser successfully parses the sentence

"The deep philosopher throws the peach pit
into the deep pit."

through cooperation among the appropriate word experts. Initialization of the parser consists of retrieving the experts for "the", "deep", "philosopher", "throw", "s",³

²An important assumption of the word expert viewpoint is that the set of such contextual word usages is not only finite, but fairly small as well.

³The perspective of viewing language through lexical contributions to structure and meaning has naturally led to the development of word experts for common morphemes that are not words (and even, experimentally, for punctuation). Especially important is the word expert for "-ing", which aids significantly in helping to

¹The research described in this report is funded by the National Aeronautics and Space Administration under grant number NSG-7253. Their support is gratefully acknowledged.

Some word senses of "heavy"

1. An overweight person is politely called "heavy":
"He has become quite heavy."
2. Emotional music is referred to as "heavy":
"Mahler writes heavy music."
3. An intensity of precipitation is "heavy":
"A heavy snow is expected today."

Some word senses of "ice"

1. The solid state of water is called "ice":
"Ice melts at 0°C."
2. "Ice" participates in an idiomatic nominal describing a favorite delight:
"Homemade ice cream is delicious."
3. "Dry ice" is the solid state of carbon dioxide:
"Dry ice will keep that cool all day."
4. "Ice" or "iced" describes things that have been cooled (sometimes with ice):
"One iced tea to go please."
5. "Ice" also describes things made of ice:
"The ice sculptures are beautiful!"
- 6,7. "Ice hockey" is the name of a popular sport which has a rule penalizing an action called "icing":
"He iced the puck causing a face-off."
8. The term "ice box" refers to both a box containing ice used for cooling foods and a refrigerator:
"This ice box isn't plugged in!"

Figure 1: Example contextual word usages

"over", and so forth, from a disk file, and organizing them along with data repositories called word bins in a left to right order in the sentence level workspace. Note that three copies of the expert for "the" and two copies of each expert for "deep" and "pit" appear in the workspace. Since each expert executes as a process,⁴ each process instantiation in the workspace must be put into an executable state. At this point, the parse is ready to begin.

The word expert for "the" runs first, and is able to terminate immediately, creating a new concept designator (called a concept bin and participating in the concept level workspace) which will eventually hold the data about the intellectual philosopher described in the input. Next the "deep" expert runs, and since "deep" has a number of word senses,⁵ is unable to terminate (i.e., complete its discrimination task). Instead, it suspends its execution, stating the conditions upon which it should be resumed. These conditions take the form of associative trigger patterns, and are referred to as

disambiguate expressions involving gerunds or participles such as "the man eating tiger". A full discussion of this will appear in [12].

⁴Although I call them "processes" word experts are actually coroutines resembling CONNIVER's generators [13], and even more so, the stack groups of the MIT LISP Machine [14].

⁵It should be clear that the notion of "word sense" as used here encompasses what might more traditionally be described as "contextual word usage". Aspects of a word token's linguistic environment constitute its broadened "sense".

restart demons. The "deep" expert creates a restart demon to wake it up when the sense of the nominal to its right (i.e., "philosopher") becomes known. The expert for "philosopher" now runs, observes the control state of the parser, and contributes the fact that the new concept refers to a person engaged in the study of philosophy. As this expert terminates, the expert for "deep" resumes spontaneously, and, constrained by the fact that "deep" must describe an entity that can be viewed as a person, it finally terminates successfully, contributing the fact that the person is intellectual.

The "throw" expert runs next and successfully prunes away several usages of "throw" for contextual reasons. A major reason for the semantic richness of verbs such as "throw", "take", and "jump", is that in context, each interacts strongly with a number of succeeding prepositions and adverbs to form distinct meanings. The word expert approach easily handles this grouping together of words to form larger word-like entities. In the particular case of verbs, the expert for a word like "throw" simply examines its right lexical neighbor, and bases its own sense discrimination on the combination of what it expects to find there, what it actually finds there, and what this neighbor tells it (if it goes so far as to ask). No interesting particle follows "throw" in the current example, but it should be easy to conceive of the basic expert probes to discriminate the sense of "throw" when followed by "away", "up", "out", "in the towel", or other words or word groups. When no such word follows "throw", as is the case here, its expert simply waits for the existence of an entire concept to its right, to determine if it meets any of the requirements that would make the correct contextual interpretation of "throw" different from the expected "propel by moving ones arm" (e.g., "throw a party"). Before any such substantive conceptual activity takes place, however, the "s" expert runs and contributes its standard morphological information to "throw's" data bin. This execution of the "s" expert does not, of course, affect "throw's" suspended status.

The "the" expert for the second "the" in the sentence runs next, and as in the previous case, creates a new concept bin to represent the data about the nominal and description to come. The "peach" expert realizes that it could be either a noun or an adjective, and thus attempts what I call a "pairing" operation with its right neighbor. It essentially asks the expert for "pit" if the two of them form a noun-noun pair. To determine the answer, both "pit" and "peach" have access to the entire model of linguistic and pragmatic knowledge. During this time, "peach" is in a state called "attempting pairing", which is different from the "suspended" state of the "throw" expert. "Pit" answers back that it does pair up with "peach" (since "pit" is aware of its run-time context) and enters the "ready" state. "Peach" now determines its correct sense and terminates. And since only one meaningful sense for "pit" remains, the "pit" expert executes quickly, terminating with the contextually appropriate "fruit pit" sense. As it terminates, the "pit" expert closes off the concept bin in which it participates, spontaneously resuming the "throw" expert. An examination of the nature of fruit pits reveals that they are perfectly suited to propelling with ones arm, and thus, the "throw" expert terminates successfully, contributing its word sense to its event concept bin.

The "into" expert runs next, opens a concept bin (of type "setting") for the time, location, or situation about to be described, and suspends itself. On suspension, "into's" expert posts an associative restart condition that will enable its resumption when a new picture concept is opened to the right. This initial action takes place for most prepositions. In certain cases, if the end of a sentence is reached before an appropriate expected concept is opened, an expert will take alternative action. For example, one of the "in" experts restart trigger patterns consists of control state data of just this kind -- if the end of a sentence is reached, and no conceptual object for the setting created by "in" has been found, the "in" expert will resume nonetheless, and create a default concept, or perform some kind of intelligent reference determination. The sentence "The doctor is in." illustrates this point.

In the current example, the "the" expert that executes immediately after "into's" suspension creates the expected picture concept. The word expert for "deep" then runs and, as before, cannot immediately discriminate among its several senses. "Deep" thus suspends, waiting for the expert for the word to its right to help. At this point, there are two experts suspended, although the control flow remains fairly simple. Other examples exist in which a complex set of conceptual dependencies cause a number of experts to be suspended simultaneously. These situations usually resolve themselves with a cascading of expert resumptions and terminations. In our "deep pit" example, "deep" posts expectations on the central tableau of global control state knowledge, and waits for "pit" to terminate. "Pit's" expert now runs, and since this

bulletin board contains "deep"'s expectations of a person, volume, or printed matter, "pit" maps immediately onto a large hole in the ground. This in turn, causes both the resumption and termination of the "deep" expert as well as the closure of the concept bin to which they belong. At the closing of the concept bin, the "into" expert resumes, marks its concept as a location, and terminates. With all the word experts completed and all concept bins closed, the expert for "." runs and completes the parse. The concept level workspace now contains five concepts: a picture concept designating an intellectual philosopher, an event concept representing the throwing action, another picture concept describing a fruit pit which came from a peach, a setting concept representing a location, and the picture concept which describes precisely the nature of this location. Work on the mechanism to determine the schematic roles of the concepts has just begun, and is described briefly later. A program trace that shows the actions of the Word Expert Parser on the example just presented is available on request.

3. Structure of the Model

The organization of the parser centers around data repositories on two levels -- the sentence level workspace contains a word bin for each word (and sub-lexical morpheme) of the input and the concept level workspace contains a concept bin (described above) for each concept referred to in the input sentence. A third level of processing, the schema level workspace, while not yet implemented, will contain a schema for each conceptual action of the input sentence. All actions affecting the contents of these data bins are carried out by the word expert processes, one of which is associated with each word bin in the workspace. In addition to this first order information about lexical and conceptual objects, the parser contains a central tableau of control state descriptions available to any expert that can make use of self referential knowledge about its own processing or the states of processing of other model components. The availability of such control state information improves considerably both the performance and the psychological appeal of the model -- each word expert attempting to disambiguate its contextual usage knows precisely the progress of its neighbors and the state of convergence (or the lack thereof) of the entire parsing process.

Word Experts

The principal knowledge structure of the model is the word sense discrimination expert. A word expert represents the linguistic knowledge required to disambiguate the meaning of a single word in any context. Although represented computationally as coroutines, these experts differ considerably from ad hoc LISP programs and have approximately the same relation to LISP as an augmented transition network [15] grammar.⁶ Just as the graphic representation of an augmented transition network demonstrates the basic control paradigm of the ATN parsing approach, a graphic representation for word experts exists which embodies its functional framework. Each word expert derives from a branching discrimination structure called a word sense discrimination network or sense net. A sense net consists of an ordered set of questions (the nodes of the network), and for each one, the set of possible answers to that question (the branches emanating from each node). Traversal of a sense network represents the process of converging on a single contextual usage of a word. The terminal nodes of a sense net represent distinct word senses of the word modeled by the network. A sense net for the word "heavy" appears in part (a) of Figure 2. Examination of this network reveals that four senses are represented -- the three adjective usages shown in Figure 1 plus the nominal sense of "chug" as in "Joe's heavy told me to beat it."

Expert Representation

The network representation of a word expert leaves out certain computational necessities of actually using it for parsing. A word expert has two fundamental activities. (1) An expert asks questions about the lexical and conceptual data being amassed by its neighbors, the control states of various model components, and more general issues requiring common sense or knowledge of the physical world.⁷ (2) In addition, at each node an expert performs actions to affect the lexical and conceptual contents of the workspaces, the control states of itself, concept bins,

⁶An ATN without arbitrarily complex LISP computations on each arc and at each node, that is.

⁷In addition to common sense knowledge of the physical world, this could include information about the plot, characters, or focus of a children's story, or in a specialized domain such as medical diagnosis [17], could include highly domain specific knowledge.

and the parser as a whole, and the model's expectations. The current procedural representation of the word expert for "heavy" appears as part (b) of Figure 2.

Each word expert process includes three components -- a declarative header, a start node, and a body. The header provides a description of the expert's behavior for purposes of inter-expert constraint forwarding. If sense discrimination by a word expert results in the knowledge that a word to its right, either not yet executed or suspended, must map to a specific sense or conceptual category, then it should constrain it to do so, thus helping it avoid unnecessary processing or fallacious reasoning. Since word experts are represented as processes, constraining an expert consists of altering the pointer to the address at which it expects to continue execution. Through its descriptive header, an expert conditions this activity and insures that it takes place without disastrous consequences.

Each node in the body of the expert has a type designated by a letter following the node name, either Q (question), A (action), S (suspend), or T (terminal). By tracing through the question nodes (treating the others as vacuous except for their goto pointers), a sense network for each word expert process can be derived. The graphical framework of a word expert (and thus the questions it asks) represents its principal linguistic task of word sense disambiguation. Each question node has a type, shown following the Q in the node -- MC (multiple choice), C (conditional), YN (yes/no), and PI (possible/impossible). In the example expert for "heavy", node n1 represents a conditional query into the state of the entire parsing process, and node n12 a multiple choice question involving the conceptual nature of the word to "heavy"'s right in the input sentence.

Multiple choice questions typically delve into the basic relations among objects and actions in the world. For example, the question asked at node n12 of the "heavy" expert is typical:

"Is the object to my right better described as
an artistic object, a form of precipitation, or
a physical object?"

Action nodes in the "heavy" expert perform such tasks as determining the concept bin to which it contributes, and posting expectations for the word to its right. In terms of its side effects, the "heavy" expert is fairly simple. A full account of the word expert representation language will be available next year [12].

Expert Questions

The basic structure of the Word Expert Parser depends principally on the role of individual word experts in affecting (1) each other's actions and (2) the declarative result of computational analysis. Experts affect each other by posting expectations on the central bulletin board, constraining each other, changing control states of model components (most notably themselves), and augmenting data structures in the workspaces.⁸ They contribute to the conceptual and schematic result of the parse by contributing object names, descriptions, schemata, and other useful data to the concept level workspace. To determine exactly what contributions to make, i.e., the accurate ones in the particular run-time context at hand, the experts ask questions of various kinds about the processes of the model and the world at large.

Four types of questions may be asked by an expert, and whereas some queries can be made in more than one way, the several question types solicit different kinds of information. Some questions require fairly involved inference to be answered adequately, and others demand no more than simple register lookup. This variety corresponds well, in my opinion, with human processing involved in conceptual analysis. Certain contextual clues to meaning are structural; taking advantage of them requires solely knowledge of the state of the parsing process (e.g., "building a noun phrase"). Other clues subtly present themselves through more global evidence, usually having to do with linking together high order information about the specific domain at hand. In story comprehension, this involves the plot, characters, focus of attention, and general social psychology as well as common sense knowledge about the world. Understanding texts dealing with specialized subject matter requires knowledge about that particular subject, other subjects related to it, and of course, common sense. The questions asked by a word expert in arriving at the correct contextual interpretation of a word probe sources of both kinds of information, and take different forms.

⁸The blackboard of the Hearsay speech understanding system [16] is analogous to the entire workspace of the parser, including the word bins, concept bins, and bulletin board.

The Importance of Multiple Choice

Multiple choice questions comprise the central inferential component of word experts. They derive from Rieger's notion that intelligent selection among competing alternatives by relative differencing represents an important aspect of human problem solving [7]. The Word Expert Parser, unlike certain standardized tests, prohibits multiple choice questions from containing a "none of the above" choice. Thus, they demand the most "reasonable" or "consistent" choice of potentially unappealing answers. What does a child (or adult) do when faced with a sentence that seems to state an implausible proposition or reference implausible objects? He surely does his best to make sense of the sentence, no matter what it says. Depending on the context, certain intelligent and literate people create metaphorical interpretations for such sentences. The word expert approach interprets metaphor, idiom, and "normal" text with the same mechanism.

Multiple choice questions make this possible but answering them may require tremendously complex processing. A substantial knowledge representation formalism based on semantic networks, such as KRL [19], with multiple perspectives, procedural attachment, and intelligent description matching, must be used to represent in a uniform way both general world knowledge and knowledge acquired through textual interpretation. In KRL terms, a multiple choice question such as "Is the object RAIN more like ARTISTIC-OBJECT, PHYSICAL-OBJECT, or PRECIPITATION?" must be answered by appeal to the units representing the four notions involved. Clearly, RAIN can be viewed as a PHYSICAL-OBJECT; much less so as an ARTISTIC-OBJECT. However, in almost all contexts, RAIN is closest conceptually to PRECIPITATION. Thus, this should be the answer. This multiple choice mechanism has many uses in conceptual parsing and full-scale language comprehension as well as in general problem solving [20]. That any fragment of text (or other human sensual input) has some interpretation from the point of view of a particular reader constitutes a fundamental underlying idea of the word expert approach.

Expert Side Effects

Word experts take two kinds of actions -- actions explicitly intended to affect sense discrimination by other experts, and actions to augment the conceptual information that constitutes the result of a parse. Each path through a sense network represents a distinct usage of the modeled word, and at each step of the way, the word expert must update the model to reflect the state of its processing and the extent of its knowledge. The "heavy" expert of Figure 2(b) exhibits several of these actions. Nodes n2 and n3 of this word expert process represent "heavy"'s decision about the concept bin (i.e., conceptual notion) in which it participates. In the first case, it decides to contribute to the same bin as its left neighbor; in the second, it creates a new one, eventually to contain the conceptual data provided by itself and perhaps other experts to its right. At node n10, "heavy" posts its expectations regarding the word to its right on the central bulletin board. When it temporarily suspends execution at node n11, its "suspended" control state description also appears on this tableau.

Control state descriptions such as "suspended", "terminated", "attempting pairing" (see above), and "ready" are posted on this bulletin board, which contains a state designation for each expert and concept in the workspace, as well as a description of the parser state as a whole. Under restricted conditions, an expert may affect the state descriptions on this tableau. An expert that has determined its nominal role, may, for example, change the state of its concept (the one to which it contributes) to "bounded" or "closed", depending on whether or not all other experts participating in that concept have terminated. Word experts may post expectations on the bulletin board to facilitate handshaking between themselves and subsequently executing neighbors. In the example parse, the "deep" expert expects an entity that it can describe; by saying so in detail, it enables the "pit" expert to terminate successfully on first running, something it would not be able to do otherwise.

The initial execution of a word expert must accomplish certain goals of a structural nature. If the word participates in a noun-noun pair, this must be determined; in either case, the expert must determine the concept bin to which it contributes all of its descriptive data throughout the parse. This concept

⁹An exception arises when an expert creates a default concept bin to represent a conceptual notion referenced in the text, but to which no words in the text contribute. The automobile in "Joanie parked." is an example.

(a) Network representation of "heavy" expert

[word-expert heavy

```

<header
  category (PA . n1)
  sense <descriptors (LARGE-PHYSICAL-MASS . nt1)
                    (INTENSE-QUANTITY . nt3)
                    (SERIOUS-OR-EMOTIONAL . nt2)>>

<start n0>

<expert
  n0:A (REFUSE)
      (NEXT n1)
  n1:Q C parser-state t
      (open-picture . n2)
      (* . n3)
  n2:A (CONCEPT old (1w))
      (NEXT n4)
  n3:A (CONCEPT new PICTURE)
      (NEXT n4)
  n4:A (CATEGORY PA)
      (NEXT n10)
  n10:A (EXPECT (rw) view/PP ART)
       (EXPECT (rw) view/PP PRECIPITATION)
       (EXPECT (rw) view/PP PHYSOBJ)
       (NEXT n11)
  n11:S wait-for-right-word
       (RESUME (trigger 'expert-state (rw) 'terminated))
       (QUEUE first)
       (NEXT n12)
  n12:Q MC view/PP (rw)
       {art . nt2}
       {precipitation . nt3}
       {physobj . nt1}
  nt1:T PA LARGE-PHYSICAL-MASS
  nt2:T PA SERIOUS-OR-EMOTIONAL
  nt3:T PA INTENSE-AMOUNT]>
  
```

(b) Process representation of "heavy" expert

Figure 2: Word expert representation

The explicit representation of control state and structural information facilitates its use in parsing -- conditional and yes/no questions perform simple lookup operations in the PLANNER-like associative data base [18] that stores the workspace data. Questions about the plot of a story or its characters, or common sense questions requiring spatial or temporal simulations are best phrased as possible/impossible (or yes/no/maybe) questions. Sometimes during sense discrimination, the plausibility of some general fact leads to the pursuit of different information than its implausibility. Such situations occur with enough frequency to justify a special type of question to deal with them.

could either be one that already exists in the workspace or a new one created by the expert at the time of its decision. After deciding on a concept, the principal role of a (content) word expert is to discriminate among the possibly many remaining senses of the word. Note that a good deal of this disambiguation may take place during the initial phase of concept determination. After asking enough questions to discover some piece of conceptual data, this data augments what already exists in the word's concept bin, including declarative structures put there both by itself and by the other lexical participants in that concept. The parse completes when each word expert in the workspace has terminated. At this point, the concept level workspace contains a complete conceptual interpretation of the input text.

Conceptual Case Resolution

Adequate conceptual parsing of input text requires a stage missing from this discussion and constituting the current phase of research --- the attachment of each picture and setting concept (bin) to the appropriate conceptual case of an event concept. Such a mechanism can be viewed in an entirely analogous fashion to the mechanisms just described for performing local disambiguation of word senses. Rather than word experts, however, the experts on this level are conceptual in nature. The concept level thus becomes the main level of activity and a new level, call it the schema level workspace, turns into the main repository for inferred information. When a concept bin has closed, a concept expert is retrieved from a disk file, and initialized. If it is an event concept, its function is to fill its conceptual cases with settings and pictures; if it is a setting or picture, it must determine its schematic role. The activity on this level, therefore, involves higher order processing than sense discrimination, but occurs in just about the same way. The ambiguities involved in mapping known concepts into conceptual case schemata appear identical to those having to do with mapping words into concepts. Discovering that the word "pit" maps in a certain context to the notion of a "fruit pit" requires the same abilities and knowledge as realizing that "the red house" maps in some context to the notion of "a location for smoking pot and listening to records". The implementation of the mechanisms to carry out this next level of inferential disambiguation has already begun. It should be quite clear that this schematic level is by no means the end of the line -- active expert-based plot following and general text understanding fit nicely into the word expert framework and constitute its logical extension.

4. Summary and Conclusions

The Word Expert Parser is a theory of organization and control for a conceptual language analyzer. The control environment is characterized by a collection of generator-like coroutines, called word experts, which cooperatively arrive at a conceptual interpretation of an input sentence. Many forms of linguistic and non-linguistic knowledge are available to these experts in performing their task, including control state knowledge and knowledge of the world, and by eliminating all but the most persistent forms of ambiguity, the parser models human processing.

This new model of parsing claims a number of theoretical advantages: (1) Its representations of linguistic knowledge reflect the enormous redundancy in natural languages -- without this redundancy in the model, the inter-expert handshaking (seen in many forms in the example parse) would not be possible. (2) The model suggests some interesting approaches to language acquisition. Since much of a word expert's knowledge is encoded in a branching discrimination structure, adding new information about a word involves the addition of a new branch. This branch would be placed in the expert at the point where the contextual clues for disambiguating the new usage differ from those present for a known usage. (3) Idiosyncratic uses of language are easily encoded, since the word expert provides a clear way to do so. These uses are indistinguishable from other uses in their encodings in the model. (4) The parser represents a cognitively plausible model of sequential coroutine-like processing in human language understanding. The organization of linguistic knowledge around the word, rather than the rewrite rule, motivates interesting conjectures about the flow of control in a human language understander.

ACKNOWLEDGEMENTS

I would like to thank Chuck Rieger for his insights, encouragement, and general manner. Many of the ideas presented here Chuck has graciously allowed me to steal. In addition, I thank the following people for helping me with this work through their comments and suggestions: Phil Agre, Milt Grinberg, Phil London, Jim Reggia, Hanan Samet, Randy Trigg, Rich Wood, and Pamela Zave.

REFERENCES

- [1] Rieger, C. and S. Small, Word Expert Parsing, Proceedings of the 6th International Joint Conference on Artificial Intelligence, 1979.
- [2] Riesbeck, C., Computational Understanding: Analysis of Sentences and Context, AI-Memo 238, Stanford University, 1974.
- [3] Riesbeck, C. and R. Schank, Comprehension by Computer: Expectation-based Analysis of Sentences in Context, Research Report 78, Yale University, 1976.
- [4] Schank, R., Conceptual Dependency: A Theory of Natural Language Understanding, Cognitive Psychology, vol. 3, no. 4, 1972.
- [5] Wilks, Y. Making Preferences More Active, Artificial Intelligence, vol. 11, no. 3, 1978.
- [6] Marcus, M., Capturing Linguistic Generalizations in a Parser for English, Proceedings of the 2nd National Conference of the Canadian Society for Computational Studies of Intelligence, 1978.
- [7] Rieger, C., The Importance of Multiple Choice, Proceedings of the 2nd Conference on Theoretical Issues in Natural Language Processing, 1978.
- [8] Rieger, C., Viewing Parsing as Word Sense Discrimination, A Survey of Linguistic Science, Dingwall (ed.), Greylock Pub., 1977.
- [9] Rieger, C., Five Aspects of a Full Scale Story Comprehension Model, Associative Networks -- The Representation and Use of Knowledge in Computers, Findler (ed.), Academic Press, 1979.
- [10] Rieger, C., An Organization of Knowledge for Problem Solving and Language Comprehension, Artificial Intelligence, vol. 7, no. 2, 1976.
- [11] Small, S., Conceptual Language Analysis for Story Comprehension, Technical Report 663, University of Maryland, 1978.
- [12] Small, S., Word Experts for Conceptual Language Analysis, Ph.D. Thesis (forthcoming), University of Maryland, 1980.
- [13] McDermott, D. and G. Sussman, The Conniver Reference Manual, AI-Memo 259a, Massachusetts Institute of Technology, 1974.
- [14] Lisp Machine Group, LISP Machine Progress Report, AI-Memo 444, Massachusetts Institute of Technology, 1977.
- [15] Woods, W., Transition Network Grammars for Natural Language Analysis, Communications of the ACM, vol. 13, no. 10, 1970.
- [16] Erman, L. and V. Lesser, A Multi-Level Organization for Problem Solving using Many, Diverse, Cooperating Sources of Knowledge, Proceedings of the 4th International Joint Conference on Artificial Intelligence, 1975.
- [17] Reggia, J., Representing and Using Medical Knowledge for the Neurological Localization Problem (First Report of the NEUREX Project), Technical Report 693, University of Maryland, 1978.
- [18] Sussman, G., T. Winograd, and E. Charniak, Micro-Planner Reference Manual, AI-Memo 203a, Massachusetts Institute of Technology, 1971.
- [19] Bobrow, D. and T. Winograd, An Overview of KRL, A Knowledge Representation Language, Cognitive Science, vol. 1, no. 1, 1977.
- [20] London, P., Dependency Networks as a Representation for Modeling in General Problem Solvers, Technical Report 698, University of Maryland, 1978.