VIEWING Word EXPERT PARSING AS LINGUISTIC THEORY*

Small

Department of Computer Science
Mathematical Sciences Building
University of Rochester
Rochester, New York 14627

Groupe d'Intelligence Artificiellc
Departernent d'Informatique
University Paris VIII - Vincennes
93200 Saint Denis FRANCE

ABSTRACT

The *Word Expert Parser* is a computer program that analyzes fragments of natural language text in order to extract their meaning in context. The construction of the program has led to the development of a linguistic theory based on notions orthogonal to those traditionally found at the heart of such theories. Word Expert Parsing explains the understanding of textual fragments containing highly idiosyncratic elements, such as idioms, collocations, cliches, and colligations, as well as lexical sequences that contain interesting structural phenomena. The theory perceives the individual word of language as the organizing unit for linguistic knowledge, and views understanding as consisting of *lexical Interactions* among procedural word experts. This paper describes four classes of lexical interaction required to explain the understanding of sentences in context, *Idiosyncratic Interaction, linguistic Interaction, discourse interaction,* and *logical interaction.* The paper purposely avoids programming details in order to focus on Word Expert Parsing as linguistic theory.

1. INTRODUCTION

The *Word Expert Parser* (UEP) is a computer program that analyzes fragments of natural language text in order to extract their meaning in context. The system has been developed with particular attention paid to the wide variety of different meaning roles of words when appearing in combination with other words. The theoretical position advanced by UEP about the nature of Individual words is that words have no meaning per se, but rather, that fragments of lexical items mean something through their Interrelationships. Furthermore, the character of lexical relations runs the gamut from the simple direct knowledge that some word sequence represents some remembered concept, to the more analytical knowledge that particular kinds of lexical sequences often represent certain classes of conceptual notions.

• The initial research motivating the ideas described In this report was supported by grant NSG-7253 from the National Aeronautics and Space Administration to the University of Maryland. During the writing of this paper, the author was supported by a Fulbright Lectureship in Artificial Intelligence at the University de Paris, and computer facilities were provided by the *Institut de Recherche et Coordination AcoustIqua/Musique.* The support of NASA, the Fulbright-Hays program, and IRCAM are gratefully acknowledged.

The evolution of this perspective started with the observation that the understanding of a particular fragment of text depends fundamentally on the disambiguation of the individual words composing it. Knowing the contextual meanings of the words is tantamount to understanding the meaning of the overall fragment. Another way of saying the same thing [Rieger, 1977] is that language interpretation can be ultimately viewed as a process of word sense discrimination. Unfortunately, this perspective does not eliminate the classic problems of deciding the nature of a distinct word sense, the difference between different usages, word senses, and idioms, and so forth. The solution to these problems comes in realizing that the process of understanding the meaning of words in context does not require reference to those notions at all. The design of a parsing procedure based on determining the meaning roles of individual words in context has led to the Drthogonal linguistic notions that are the subject of this paper.

The organization of UEP is founded on the belief that the grouping together of words to form meaningful sequences is an active process which succeeds only because of highly idiosyncratic application of lexical knowledge. That is, we fragment text and understand the meaning of the pieces because we know how the particular words involved interact with each other.

Sometimes sequences of two or more words interact together to such an extent that they seem to behave as a single lexical item. Linguists have labelled such sequences *Idioms.* The notion to which this definition gives rise, however, causes several problems for linguistic theory. First of all, rarely does such a sequence hold together so tightly that it can be truly treated theoretically as a single lexical item. Secondly, rarely does such a sequence have a unique meaning. More often than not, the meaning of an idiomatic expression must be determined by disambiguation. The sequence must be analyzed in context and be treated by comprehension processes as being either (a) a cohesive whole with idiosyncratic meaning, or (b) a sequence having meaning through less specific language knowledge. There is no a *priori* way of knowing the meaning of the sequence to be the one or the other.

The notion of *idiom* falls at one end of a spectrum, an idealized end that I claim does not exist. Lexical seauences can be more *or less Idiomatic,* in the sense that the process interactions constituting the understanding of them includes greater or fewer idiosyncratic Interactions. The UEP way of looking at the most idiomatic sequences is that the special interactions among the participating words take priority over any other potential interactions involving those words. The disambiguation of idiomatic expressions, i.e., the understanding of the sequences as either idioms or non-idioms (to use the popular distinction), generally requires other process interactions besides the strictly word-specific ones. The understanding of an idiom thus differs insignificantly, from the perspective of UEP theory, from comprehension of any other kind (according to whatever classification scheme) of lexical sequence.

The notion that all fragments of language are more or less idiomatic, while radical in some linguistic quarters, has been previously suggested. In his introductory textbook, Aspects of Language, Dwight Bolinger asks "whether everything we say may be in some degree idiomatic — that is, whether there are affinities among words that continue to reflect the attachments the words had when we learned them, within larger groups" [Bolinger, 1975J. After working within what he calls "the prevailing reductionism", Bolinger began to suggest a positive answer to his pedagogical question, choosing to take "an idiomatic rather than an analytical view" [Bolinger, 1979] of language. The contribution of artificial intelligence in general, and of *Word Expert Parsing* in particular, is to develop theory from this informal view. The notion of process, and of process interaction, allows us to begin to do Just that.

2. THE UORD EXPERT PARSER

The UEP computer system maintains linguistic knowledge across a community of word-based structures called *word experts,* which represent the process of determining the contextual meaning and role of the individual words. A word expert must not be thought of as a representation for the various meanings, roles, and contributions of a word in context, but rather as a declarative representation (a network) of the process (which we shall call *disambiguation)* of determining these things. Certainly, it is the meaning contributions of individual lexical items that we wish to determine. Word experts are both data and process; they can be augmented, examined, and manipulated as data, yet parsing takes place through their interpretation as program by an expert evaluator, similar to the EVAL of Lisp.

The distributed parsing scheme of UEP works as follows. The UEP reader examines a word of text and retrieves its word expert from memory. The word expert starts executing, trying to determine the meaning role of its word in context, i.e., interacting with other word experts and with higher-order system processes to acquire the appropriate contextual knowledge to make the correct inferences. Finally, all the word experts for a particular fragment of text come to mutual agreement on the meaning of the fragment, and the local distributed process terminates. Local, in the sense that as long as there remains input text, the overall parsing process continues, while the disambiguation of individual lexical sequences making up the larger text completes.

Interaction, between individuals in the world, or between distributed processes in a computer program, requires both (a) giving information and (b) receiving information. In UEP, the experts exchange two kinds of information, called concept *structures* and *control signals.* Concept structures represent human concepts, such as "a book", "going fishing", "the box of candy 1 gave Joanie for Valentine's day in 1981", "some blue physical object", and the like. Control signals represent processing clues, such as "expect a word that can begin a lexical sequence that can describe concept structure X", "send me the concept structure representing the agent of concept structure Y or a signal saying you cannot", "wait a second and you will be sent a concept structure that will help you", and similar things. The representation and use of concepts and signals are described fully in [Small, 1980].

3, LEXICAL INTERACTIONS

I use the term *lexical interaction* to denote the sending and receiving of control signals and concept structures by word experts in UEP, This includes interactions between individual experts, as well as those between a word expert and another kind of model process (e.g., a mechanism inferring the goals of a dialogue participant). This paper discusses lexical interactions by presenting four classes of required interaction, and then arguing for the necessity and giving examples of each. The categorization is by the kind of knowledge exchanged in the communication, and includes the following.

Idiosyncratic Interaction

Linguistic Interaction

Discourse Interaction

Logical Interaction

The least general class of lexical interactions are considered i*diosyncratic* since they are word-specific and arise through simple recall memory. This type of interaction permits the understanding of idiomatic fragments. General knowledge about the syntax and semantics of some natural language gives rise to *linguistic* interactions, and are of course crucial to the understanding of lexical sequences not previously seen. Sometimes words interact with processes that monitor the development of an entire text (or parts thereof), or the goals of participants in discussion. These *discourse* interactions are often necessary for the meaningful cohesion of lexical fragments. Lastly, but certainly not least important are the *logical* interactions between words and the most general cognitive processes. Perceptions about the world, beliefs, inference-making skills, rote memory, and so forth, are basic to language understanding.

The classification of word fragments into categories such as "idiom", "collocation", "colligation", "noun phrase", "complement", and the like, does not make sense in UEP theory. Rather, individual words are viewed as having certain kinds and sequences of interactions with their neighbors to form meaningful pieces of text. Fragments often described as "Idioms" are those that are understood principally through idiosyncratic lexical interactions. A non-idiomatic structure, diagnosed as a "noun phrase", is one that involves mostly linguistic interactions to understand. A so-called "noun-noun pair" can be thought of as a lexical sequence comprehended with the help of logical interactions, with recourse to common sense memory and skills.

3.1 Idiosyncratic interaction

Since the emphasis of the UEP research effort is to construct a computer program to understand natural language, we are not qualified to make claims about just how much of human parsing involves idiosyncratic lexical interaction. Suffice it to say that no theory of language analysis can do without such a notion, and that we have come across many examples, both in our own work and in the linguistic literature, where it applies. Furthermore, our conception of idiosyncratic interaction (in conjunction with the other interaction types), provides explanatory adequacy for many linguistic phenomena, observed in diverse camps.

IdiOMS

The comprehension of idioms requires both idiosyncratic lexical interactions and context-probing disambiguation (i.e., interactions between lexical processes and higher-order memory processes). As an example, examine the following sentence, analyzed easily by the prototype WEP system described in [Small, 1980].

(1) "The fellow throws in the towel."

Linguistic description of this sentence could take several different routes. The verb in the sentence could be seen to be either one of (a) *throw,* (b) *throw in,* or (c) throw *in the towel,* depending on the theoretical perspective of the linguistic approach. While UEP does not itself (explicitly) make use of the notion of a verb, if forced to explain its behavior in these terms, I would say that UEP would consider the verb to be either (a), (b), or (c), depending on the context.

The UEP computer system would behave as follows. The word expert for *throw* would send a message to the *in* expert, consisting of both a control sifinal and a concept structure. The signal would in effect tell *1n* that throw would like to pair .up with it to form a cohesive fragment. The accompanying concept structure would tell *in* what kind of lexical sequence (or conceptual mapping of one) throw would expect to find to the right of *in* for it to have confidence in such a pairing. Of course, the *in* expert has independent control of its own interactions, and could decide (based on the nature of the input) to reject any suggestions it receives. In the above example, *in* would wait for a signal from the expert on its right (the the expert) indicating the formation of a cohesive fragment designating a towel, and then act on this Information.

The concept structure sent by throw indicates that if the words to the right of *in* are the towe, then the words throw and *in* should pair up. However, the *in* expert has the prerogative to override this suggestion, and if query to the processes modelling the focus of discourse attention (see the later section on discourse interactions) determined that some large towel was the location of some active event, would do so, transmitting a signal to *throw* rejecting its advice. Ordinarily, this would not happen, and the *in* expert would signal acceptance to *throw.* The *throw* expert would then query the discourse processes monitoring the activity context to determine if (a) someone is actually throwing a towel into some volume, or (b) someone is conceding defeat to someone in a game. When this was determined, the understanding of the fragment (in this case a sentence) would be complete.

Collocation and Colligation

What is the nature of the advice given by one word expert to another, as from the *throw* expert to the *1n* expert above, and what is its basis? These suggestions, transmitted as control signals and concept structures in WEP, are based on purely idiosyncratic criteria. While the idiomatic meaning of certain lexical sequences cannot be predicted from the parts, certain others have meanings that differ slightly from what might be inferred. These sequences are sometimes called *collocations,* and illustrate the idiosyncracy of lexical pairings and thus the basis for certain interactions in UEP analysis.

In his recent paper, Bolinger [1979] quotes from T. F. Mitchell [1971] about the word *work* and its meaning in various fragments of lexical items: "Men — specifically cement workers — work *in* cement works; others of different occupation work *on* works of art; others again, or both, perform good works. Not only are good works *performed* but cement works are boilt and works of art *produced"* (italics in original). Bolinger goes on to reflect why "builders do not *produce* a building or authors *invent* a novel". From the vantage point of the distributed lexical actors of UEP, these particularities form the basis of the inter-expert communications we have been calling *idiosyncratic interactions.*

Let us take Mitchell's example, the ambiguous word work, and see how the word experts of UEP can interact to determine its meaning in context. It should be easy to see from the discussion so far how UEP would handle the following fragment.

(2) "The Ensemble Intercontemporain performs a work."

Without knowing anything about the various entities mentioned in the sentence, it is clear that an ensemble of musicians played a musical piece. Lexical interactions between the word expert for *perform* and that for *work* make clear the meaning of the fragment. The *perform* expert would signal to *work* its expectation of a lexical sequence denoting a conceptual entity that can be "performed", such as "a service", "a musical piece", or "a series of actions". It would then be up to *work* to use this (non-binding) advice to contribute to the meaning of the overall fragment in which it participates.

3.2 Linguistic Interaction

The syntax and semantics of natural language comprise the stuff of most linguistic theory, and their use is the cornerstone of computer parsing programs. However, UEP forces us to perceive the nature and use of this knowledge from an untraditional vantage point, that of the individual word and its active processing to form meaningful fragments of text with its neighbors. From this perspective, the syntax and semantics of a language is that body of knowledge that helps us infer enough about the meaning contributions of new words to understand lexical sequences completely different from ones that we have already seen. After processing some lexical sequence, this linguistic knowledge provides us with certain expectations about upcoming lexical items and the nature of their interactions, thus helping us to fragment those items into pieces and to infer the meaning of the pieces from general information about their component parts.

Noun Phrases

How can the purely lexical UEP system require no notion of high order structural phenomena, yet still be able to account for them? The following example (provided by Yorick wilks) illustrates the lexical interactions required to analyze an interesting fragment of text.

(3) "Joanie washes the colorful dishes up."

The difficulty with this fragment is in determining that the word *dishes* contributes to the meaning of the fragment through interactions with the two words to its left, but that the word *up* contributes by association with the word *washes,* which precedes *up* by many intervening words. The reason that I am avoiding the use of traditional linguistic jargon for describing this phenomenon is the following belief: An understanding of UEP requires the viewing of language interpretation from the vantage point of the individual word and its interactions. An important way to achieve this is to describe the analysis process with reference to the very notions (not the traditional ones) around which it is organized.

In the analysis of the example fragment, UEP would find the referent of *Joanie* and then proceed as follows. The *wash* expert would begin executing, trying to determine its own meaning role in some lexical fragment, and at the same time, trying to provide information to other lexical agents to permit them to do the same. The meaning of *wash* in context depends on a number of factors, including the nature of the words succeeding it, and their own actions in determining their meaning and role contributions in context. The *wash* expert must thus prepare for a number of contingencies, or different things that could happen in the text, and then, wait to see if any of them actually occur. If the word up appears to the right of *wash,* for example, the words could choose to pair up into a meaningful fragment (as in *throw in* above). Under certain conditions, the word *up* could appear later on in the text, and still pair up with *wash* (as must occur for correct interpretation of the example sentence).

What are the contextual conditions that would permit this? One of the contingencies that *wash* anticipates is the grouping of the words to its right into a meaningful fragment of their own (i.e., a *concept structure).* The *wash* expert knows that (a) the nature of this concept structure may be important for its own sense disambiguation, and (b) that the word immediately following the meaningful fragment could pair up with it. In the Jargon of UEP, one of the experts in such a meaningful fragment *reports* a concept structure. Since the *up* expert does not reply to *dishes* with an acceptable message to continue the ongoing concept building activity, the *dishes* expert reports the structure. It is this report that triggers some new processing by the *wash* expert, namely the examination of the next word (i.e., *up).*

The rest of the analysis takes place predictably. The *wash* expert interacts with *up* as if *up* occurred to its immediate right in the text. The pairing up of the two words results from mutual accord, and the *wash* expert creates a concept structure to represent the meaning of the washing up of dishes. Next *wash* organizes the conceptual object, *the* co7orful *dishes,* into the overall meaning of the sentence, and again waits for things to happen. This time, the word expert for the period at the end of the sentence executes, and transmits an appropriate message. The *wash* expert again executes, cleans up its business, and reports the concept structure representing the meaning, in context, of the entire fragment (sentence).

Passives and Relative Clauses

Sentences in the passive voice and those containing relative clauses are similar in being complex structural phenomena in natural language, and often suggestive of sentence-level rules as linguistic explanation. Furthermore, the understanding of such constructions by the distributed word-based approach of wEP may be far from evident, especially considering my claim that no explicit notions of structure are referenced by the computer system or used in the theory. Interpretation of textual fragments containing complex syntactic structures takes place through complex patterns of lexical interactions among the appropriate word experts. The words that normally cue a reader about the presence of such structural relations in a fragment are the ones in wEP that coordinate the process of understanding them.

The analysis of a passive sentence involves linguistic interactions among the word experts for the suffix *en,* the word *by,* and the other words composing it. The following sentence has been parsed by the existing UEP system, and discussed at length in [Small, 1980].

(4) "The case was thrown out by federal court."

The *en* expert begins executing before *throw,* and the normal attempts by the *throw* expert to coordinate the analysis of the fragment in which it participates are intercepted by an. The actions of *en* allow *throw* to pair up with out. as outlined above for *throw in* and *wash up,* but its lexical interactions to determine the nature of the object being "thrown out", and the agent doing the "throwing" are all intercepted by the *en* expert, which provides *throw* with the correct replies to its queries. Please refer to [Small, 1980] and [Small, 1981] for a fuller discussion.

Relative clauses beginning with the word *who* are analyzed by UEP through the interactions among the *who* word expert and the experts for the other words in the clause and the larger fragment containing it. The following sentence is an example of such a fragment.

(5) "The man who throws the game likes to lose."

The *who* expert in this sentence has the responsibility for interacting with the word expert for *likes* to inform *Tikes* about *the man* doing the "liking". Ordinarily, this expert would expect to find a meaningful lexical sequence to its left representing the needed concept. However, the particular structure of the fragment means that *who* must be at the other end of the relevant linguistic interactions of *likes,* rather than the expert for the word to its immediate left, which would normally perform the needed service.

The *UEP* interpretation of the example fragment proceeds as follows. The word experts for *the* and *man* agree to form a meaningful sequence and construct a concept structure to represent its meaning. The *who* expert begins executing, gets hold of this concept, and waits for the *throw* expert to start exploring the nature of the lexical sequence on its left. In addition, the *who* expert anticipates that another word expert further down the line (in the example, the expert for *7 7/ces)* will also seek out information about the sequence to its left, in exactly the way that throw does. The *who* expert, like every word expert in UEP, plans a strategy to interact with the experts involved in both its prior context and its subsequent context, cooperatively to interpret fragments of text.

The *throw* expert begins executing and investigates the nature of the lexical sequence to its left. The *who* expert provides the appropriate Information, i.e., the concept structure representing *the men,* and *throw* begins to disambiguate its meaning in context. The experts for a and *game* mutually agree on their local meaning, and through linguistic and idiosyncratic interactions with *throw* help it determine its meaning as the "throwing of a contest". The *likes* expert starts executing, and its messages in search of the person doing the "liking" are intercepted by the *who* expert, which has been on the lookout for such interactions since the beginning. Since the *who* expert knows the unique name of the concept structure representing *the man,* it sends this concept to likes, which proceeds normally, knowing nothing of the structural complexities preceeding it.

The word experts for both throw and for *likes* can be expected to explore the underlying meaning of the lexical sequences preceding them. Note the way that UEP applies this linguistic knowledge to the interpretation of fragments of natural language text containing these words. Rather than saying that throw and *Ukes* act as finite verbs in certain contexts (which are described in some relational representational scheme, such as grammar rules or logic), we say instead that these words carry on linguistic interactions with the active processes modelling the other words making up the (local linguistic) context to arrive at a mutually acceptable characterization of their individual contributions to textual meaning. The advantage of this perspective comes from the fact that linguistic interactions constitute but a portion of all possible lexical interactions that represent in UEP the process of understanding.

3.3 Discourse Interactions

While it is clear that certain lexical sequences cannot be understood solely through recourse to syntax and semantics, namely those fragments for which idiosyncratic interactions are required (i.e., specific remembered contexts), why do we need other kinds of general knowledge? We have already seen examples suggesting the answer to this question. In trying to understand the meaning of *throw 1n the towel,* the relevant word experts must find out some things about the person performing the described action, before knowing what action he is in effect performing.

If the discourse describes some sort of competition between two people (or teams), for example, *throw in the towel* could indicate a concession of defeat by one of them. The following fragment illustrates such a contextual situation.

(6) "Rick and Joanle play chess.
Rick throws in the towel."

On the other hand, if the discourse has recently made reference to a place where one might dispose of a towel, *throw 1n the towel* might be signifying the putting of some towel in that place. The following example illustrates this case.

(7) "Joanie drops a penny in the pit.
Rick throws in the towel."

I am not claiming that knowledge of the discourse context is sufficient to disambiguate the meanings of the example sentence, but rather, that such knowledge is required to understand it.

The discourse interactions required to interpret the above example take place (a) between the *throw* expert and a higher order process modelling the *activity context,* and (b) between the *1n* expert an a process modelling the discourse focus of *attention.** There are two aspects to the processing of the activity mechanism, the unsolicited sending of control signals to indicate the anticipation of certain actions in the text and concept structures to represent them, and the more data-directed interactions with word experts (and other understanding processes) to determine the nature of the actions that actually do occur. The throw expert must carry on activity *context Interactions* to determine if the discourse could be seen as discussing some competitive activity. If so, the "concession of defeat" interpretation of the example sentence is plausible. The *in* expert carries on *rocus or attention interactions* to find out if some location has recently been described in the text *in which* something might be thrown.

While the UEP system has been directed toward the understanding of fragments of text occurring in textual discourse, the issues arising in the interpretation of dialogue are very similar. The difference between the two tasks involves the nature of discourse interactions. In interpreting fragments of dialogue from the vantage point of one of the participants, word experts must interact with model processes monitoring the goals of the other participant. The following example (provided by James Allen) illustrates the question.

(8) "When is the Windsor train?"

In trying to understand this question from the perspective of the person at the information desk of a train station, the question could be directed at eliciting either of two pieces of information [Allen, 1978], i.e., the time of the next arrival from Windsor, or the time of the next departure to Windsor.

By saying that *the Windsor train* is a "noun-noun pair", we get nowhere in trying to understand it. In UEP, the word experts for *Windsor* and train would interact locally and determine the range of possible interpretations for the fragment. In the case of textual discourse, the *train* expert would carry on discourse interactions with the activity process to find out if discussion of some particular train were anticipated in the text. In the case of dialogue, these interactions would occur between *train* and an *Intention mechanism,* which might determine that the speaker in the dialogue is concerned with the trains coming from Windsor, and not with the trains leaving for Windsor. If the processes modelling the activity context or the speaker intentions cannot provide help to the train expert, the word experts for the sequence would construct a concept structure to represent the disjunct of the two possibilities, but continue to await the information that would decide between them.

* The term *activity context* describes a notion similar to the *scripts* of Schank and Abelson [1975] and to the *rrames* of Charniak [1977]. The notion of *rocus or attention* has been taken directly from the work of Grosz [1977].

3.4 Logical Interactions

The understanding of fragments of natural language text *by a particular individual* (or computer program) often requires knowledge of the beliefs of that individual. How can this be true given the fact that, even when the discourse context has not been made to render a particular fragment unambiguous, the majority of readers still interpret it the same way? The answer lies in the common experience brought to understanding by readers from the same culture. This notion carries over to the common experiences of people in all sorts of sub-cultures as well, such as scientific communities, religious groups, age and class groups, and so forth. There are pieces of text that would be understood in common by members of these groups and not by people outside these groups, and other fragments that would be understood in common by almost everyone.

A psychological experiment at Stanford University [Smith, Lance, and Shoben, 1974], described in the paper by Bolinger [1979] that I have made such extensive use of here, tested the meaning of various words in terms of class membership. Subjects were asked to judge the truth or falsity of certain sentences, such as *A robin 1s a bird* and *A chicken 1s a bird.* The fact that it generally takes longer to judge the latter true than the former led the researchers to conclude that there are degrees of birdiness. Clearly, different people with different conceptions of notions like birdiness would understand the same sentence in different ways. Furthermore, the data of these psychologists, when interpreted in terms of language understanding, suggests that a lexical sequence could be viewed as representing, *to a greater or lesser degree,* some particular conceptual notion.

It is my belief that the understanding of many sequences of words takes place through a process of comparing new sequences to already interpreted ones. Since the linguistic experience of each individual is different, language comprehension must necessarily take place in a person-specific manner. Examine the following sentence for an example of what 1 mean.

(9) "Joanie Caucus throws a seminar."

Most people reading this sentence will understand the notion of *throwing a seminar* by comparison with the other kinds of things they have (linguistically) seen *thrown.* When someone *throws* a party, for example, he is organizing a social event. When someone *throws a chess match,* he is losing the match on purpose. My claim is that by comparing *throwing a seminar* with other fragments involving *throwing,* people come up with the intended interpretation of the fragment.

Multiple Choice Perspectives

This notion of comparison with already known concepts is a fundamental method of interpretation in the Word Expert Parser. Individual word experts interact with a memory of real-world knowledge to determine whether certain conceptual notions can be perceived as other ones. The paradigm for these interactions is based on multiple choice — of all the fragments of text that have already been understood (the finite choice), which most closely resembles the one now being examined? For the example sentence above, the question in UEP would be put forth by the *throw* expert and would be some variation on the following (depending on the knowledge of *throw* stored in the *throw* word expert): "Is a seminar better viewed as a party, a tantrum, a chess game, a legal case, or a baseball?".*

This type of logical interaction (interaction between a word expert and a process modelling beliefs about the world) has a fundamental role in UEP language understanding. The queries like the example are called *(multiple choice) view interactions.* It is important to note that these multiple choice queries often take place without including a "none of the above" option. The understanding process must be directed toward the goal of providing some interpretation on each portion of input text. This means that a reader must do the best that he can to understand it, whether or not he has sufficient linguistic and cultural experience to come to the correct (intended) interpretation. Sometimes we make mistakes.

Plausibility of Propositions

The multiple choice view interactions allow word experts to use general knowledge about conceptual items in the world to understand fragments of natural language text. Another kind of general knowledge concerns structures more complex than single items, namely relations among several such conceptual items. When we perceive such a relation as being convertible into some sort of truth value, we call it a *proposition.* In UEP, propositions may have truth values along a wide range, from something we might call *completely disbelieved* to the opposite extreme we could call *completely believed.* I say that propositions in UEP work like this not because of the existence of a computer program to operate on them, but rather, because such a program must exist for language understanding within the UEP framework.

The reason involves one class of logical interactions, those between individual word experts and a process maintaining beliefs about the world. A word expert may interact with this belief modelling process to determine the *relative plausibility* of two propositions. Consider the following sentences, both of which have been successfully interpreted by the existing UEP system (with the user acting as the belief modelling process).

(10) "The man eating tiger growls."

(11) "The man eating spaghetti growls."

The difference between these two fragments from the perspective of UEP involves the relative plausibility of *tigers that eat men* and *men who eat tigers* in the first case, and of *spaghetti that eat men* and *men who eat spaghetti* in the second. Of course, in certain contexts, the problem is resolved through discourse interactions; the activity context or focus of attention could make clear the appropriate meaning of the sentences without any need at all for more general knowledge. Clearly, however, we can understand these fragments perfectly well without any guiding discourse context.

• Note the important effect such logical interactions have on motivating general inference processes in understanding. The augmentation of a real-world knowledge base depends on making appropriate inferences, and UEP motivates these by forcing the system to find relationships between previously unrelated notions. For example, by asking to relate *seminar* and *party,* the parser instigates the construction of a new concept to represent "an organized activity of people'.

The understanding of these fragments is coordinated in WEP by the word expert for the affix 1ng. The 1ng expert interacts linguistically with the experts for the words around it, helping them form meaningful sequences, and carries on logical interactions with the belief modelling process to determine the relative plausibility of the two propositions possibly signified by the larger sequence. In the first case above, the ing expert begins executing after the and man have already started constructing a concept structure to represent the meaning of the man. It awaits the report of this concept structure, as well as the one to be reported by the tiger word expert. Furthermore, 1ng carries on linguistic interactions with eat to arrive cooperatively at a concept structure representing its meaning. The ing expert then has a plausibility interaction with the belief modeller, and coordinates the remainder of the understanding process based on this important knowledge.

4. SUMMARY

Word Expert Parsing is a linguistic theory based on a lexical organization of linguistic knowledge represented procedurally in word experts. The comprehension of fragments of natural language text is viewed as a process of word interactions, where active lexical agents cooperate to form meaningful sequences of interrelated lexical items. Lexical interactions are of four types, idiosyncratic, linguistic, discourse, and logical. Idiosyncratic interactions allow UEP to explain the understanding of idiomatic (more or less idiomatic) lexical sequences, by comparing new sequences with explicitly remembered ones (called prerabs by Bolinger [1979]). Linguistic interactions enable the use of syntactic and semantic generalizations to interpret fragments, and discourse interactions provide word experts with knowledge of discourse activities and foci of attention. Logical interactions allow word experts to use knowledge about the real-world, especially about the multiple perspectives of individual conceptual objects within it and the relative plausibility of propositions about it.

S. ACKNOWLEDGEMENTS

The theory presented in this paper was originally conceived in wonderful cooperation with Chuck Rieger over the past several years. Some of the sensible perspectives in the paper have benefitted from much appreciated written and spoken suggestions and criticisms of Yorick Wilks, Dick Hudson, Dwight Bolinger, Pat Hayes, and James Allen. The Groupe d'Intelligence Artificielle of the University Paris VIII - Vincennes has provided an excellent environment for research. Thanks to Patrick Greussay, Harald Wertz, Daniel Goosens, Annette Cattenat, and Gerard Paul. Many extra thanks to Patrick for all the personal and rofessional help he has given in making my year at incennes both valuable and enjoyable.

6. REFERENCES

Allen, James (1978), Recognizing Intention in Dialogue, Technical Report, Department of Computer Science, University of Toronto.

Bolinger, Dwight (1975), Aspects of Language, Harcourt Brace Jovanovich.

Bolinger, Dwight (1979), Meaning and Memory, in Experience Forms, Haydn (ed.), Mouton.

Charniak, Eugene (1977), Ms. Malaprop, A Language Comprehension Program, Fifth International Joint Conference on Artificial Intelligence.

Chomsky, Noam (1965), Aspects of the Theory of Syntax, MIT Press.

Dik, Simon C. (1978), Functional Grammar, North Holland.

Fillmore, Charles J. (1968), The Case for Case, in Universals in Linguistic Theory, Each and Harms (eds.), Holt.

Grosz, Barbara (1977), The Representation and Use of Focus in Dialogue Understanding, Technical Note #151, Stanford Research Institute.

Hudson, Richard A. (1979), Pan-Lexicalism, Working Paper, Department of Phonetics and Lingustics, University College London.

Jackendoff, Ray (1972), Semantic Interpretation in Generative Grammar, MIT Press.

Jackendoff, Ray (1976), Toward an Explanatory Semantic Representation, Linguistic Inquiry, v. 7, n. 1.

Kaplan, Ronald M., and Joan W. Bresnan (1980), Lexical-Functional Grammar: A Formal System for Grammatical Representation, Occasional Taper #13, MIT Center for Cognitive Science.

Marcus, Mitchell P. (1979), An Overview or a Theory of Syntactic Recognition for Natural Language, AI Memo #531, MIT Artificial Inteligence Laboratory.

Miller, George A. (1978), Semantic Relations among Words, in Linguistic Theory and Psychological Reality, Halle, Bresnan, and Miller (eds.), MIT Press.

Mitchell, T. F. (1971), Linguistic "goings on"; collocations and other lexical matters arising on the syntactic record, Archivum Linguisticum, v. 2.

Rieger, Chuck (1977), Viewing Parsing as yord Sense Discrimination, in A Survey of Linguistic Science, Dingwall (ed.). Greylock Publishers.

Riesbeck, Christopher K., and Roger C. Schank (1976), Comprehension by Computer: Expectation-Based Analysis of Sentences in Context, Research Report #78, Department of Computer Science, Yale University.

Schank, Roger C., and Robert P. Abelson (1975), Scripts, Plans, and Knowledge, Fourth International Joint Conference on Artificial Intelligence.

Small, Steven (1980), Word Expert Parsing: A Theory of Distributed Word-Based Natural Language Understanding, Technical Report #954, Department of Computer Science, University of Maryland.

Small, Steven (1981), Toward a Cognitive Mechanics for Distributed Modelling, Technical Report, Department of Computer Science, University of Rochester (to appear).

Smith, Edward E., Lance J. Rips, and Edward J. Shoben (1974), Semantic Memory and Psychological Semantics, in The Psychology of Learning and Motivation, v. 8, Bower (ed.), Academic Press.

Wilks, Yorick (1980), Some Thoughts on Procedural Semantics, Cognitive Studies Centre Report #1, University of Essex.