

UC Irvine

ICS Technical Reports

Title

Parsing with parallelism : a spreading-activation model of inference processing during text understanding

Permalink

<https://escholarship.org/uc/item/5s37q93d>

Authors

Granger, Richard H.
Eiselt, Kurt P.
Holbrook, Jennifer K.

Publication Date

1984

Peer reviewed

**Parsing with Parallelism: A Spreading-Activation
Model of Inference Processing During Text Understanding**

*Richard H. Granger
Kurt P. Eiselt
Jennifer K. Holbrook*

Artificial Intelligence Project
Computer Science Department
and

Cognitive Sciences Program
Social Sciences Department

Technical Report #~~228~~ 229
September 1984

Notice: This Material
may be protected
by Copyright Law
(Title 17 U.S.C.)

This research was supported in part by the National Science Foundation under grant IST-81-20685 and by the Naval Ocean Systems Center under contracts N00123-81-C-1078 and N66001-83-C-0255.

Parsing with Parallelism:
A Spreading-Activation Model of Inference Processing
During Text Understanding

Richard H. Granger
Kurt P. Eiselt
Jennifer K. Holbrook

Artificial Intelligence Project
Computer Science Department
and
Cognitive Sciences Program
Social Sciences Department
University of California
Irvine, California 92717

ABSTRACT

The past decade of research in Natural Language Processing has universally recognized that, since natural language input is almost always ambiguous with respect to its pragmatic implications, its syntactic parse, and even its lexical analysis (i.e., choice of correct word-sense for an ambiguous word), processing natural language input requires decisions about word meanings, syntactic structure, and pragmatic inferences. The lexical, syntactic, and pragmatic levels of inferencing are not as disparate as they have often been treated in both psychological and artificial intelligence research. In fact, these three levels of analysis interact to form a joint interpretation of text.

ATLAST (A Three-level Language Analysis System) is an implemented integration of human language understanding at the lexical, the syntactic, and the pragmatic levels. For psychological validity, ATLAST is based on results of experiments with human subjects. The ATLAST model uses a new architecture which was developed to incorporate three features: spreading activation memory, two-stage syntax, and parallel processing of syntax and semantics. It is also a new framework within which to interpret and tackle unsolved problems through implementation and experimentation.

This research was supported in part by the National Science Foundation under grant IST-81-20685 and by the Naval Ocean Systems Center under contracts N00123-81-C-1078 and N66001-83-C-0255.

1.0 Introduction

The past decade of research in Natural Language Processing has universally recognized that, since natural language input is almost always ambiguous with respect to its pragmatic implications, its syntactic parse, and even its lexical analysis (i.e., choice of correct meaning, or word-sense, for an ambiguous word), processing natural language input requires decisions about word meanings, syntactic structure, and pragmatic inferences. The lexical, syntactic, and pragmatic levels of inferencing are not as disparate as they have often been treated in both psychological and artificial intelligence research. Furthermore, these three levels of analysis interact to form an interpretation of text. For example, the choice of a word-sense affects subsequent pragmatic inference decisions or syntactic structure decisions.

ATLAST (A Three-level Language Analysis System) is a computer model of how humans parse and interpret text. For psychological validity, ATLAST is based on results of experiments with human subjects. ATLAST is also an implemented integration of language understanding at the lexical, syntactic, and pragmatic levels. It uses a new architecture which consists of three processes, developed to incorporate three features: spreading activation memory, two-stage syntax, and parallel processing of syntax and semantics. Each of the processes is involved in all levels of text interpretation. The new architecture divides the abilities of the three processes in such

a way that ATLAST not only processes texts which people understand, but has difficulty with texts which cause human readers difficulty. The model employs the results of studies of many inference phenomena from several different fields of research. This approach helps solve many of the problems associated with inference decisions at all levels of processing. It is also a new framework within which to interpret and tackle unsolved problems through implementation and psychological experimentation.

2.0 Background

2.1 Our Previous Work, Briefly

The ATLAST model is a descendant of our earlier work on pragmatic ambiguity. We had worked on models which could supplant erroneous inferences with correct ones, and models which could come up with several different plausible interpretations of text events based on different pieces of world knowledge [Granger, 1980; Granger, 1981; Schulenburg, 1982; Granger, Eiselt, & Holbrook, 1983]. As we worked on these models, we observed that the different levels of inferencing have much in common. Many pragmatic inferences are triggered by individual words, which reinforced our belief in a close relationship between the lexical and pragmatic levels. For instance, consider the following examples of ambiguity at the lexical level:

[1] The CIA called in an inspector to check for bugs.
The secretaries had reported seeing roaches.

[2] The CIA called in an inspector to check for bugs.
The secretaries had reported seeing microphones.

The word "bugs" is ambiguous in both texts until the second sentence, yet the first sentence of each text implies an unambiguous reading. In text [1], the "spy" meaning of "bug" initially appears to be more appropriate than the "insect" meaning. In text [2], both sentences suggest the "spy" reading. When reading these texts, the pragmatic inferences which are made during the first sentence are based upon the lexical inferences which are originally made. The interpretation of the stories' events are thus dependent upon which meaning of the word is selected.

2.2 Lexical Access

Because of this interdependence between inference levels, theories about pragmatic inference mechanisms must include theories about lexical access processes. Lexical access, the process by which a word's meaning is extracted from its phonological or orthographic code, must include some means for selecting the most appropriate meaning for the context in which the word appears. The recent research on lexical access has led to some unexpected conclusions.

Essentially, when an ambiguous word is seen with no context (that is, alone), all meanings of the word are accessed. Then, after about a 600 msec. delay, a default meaning is selected, and

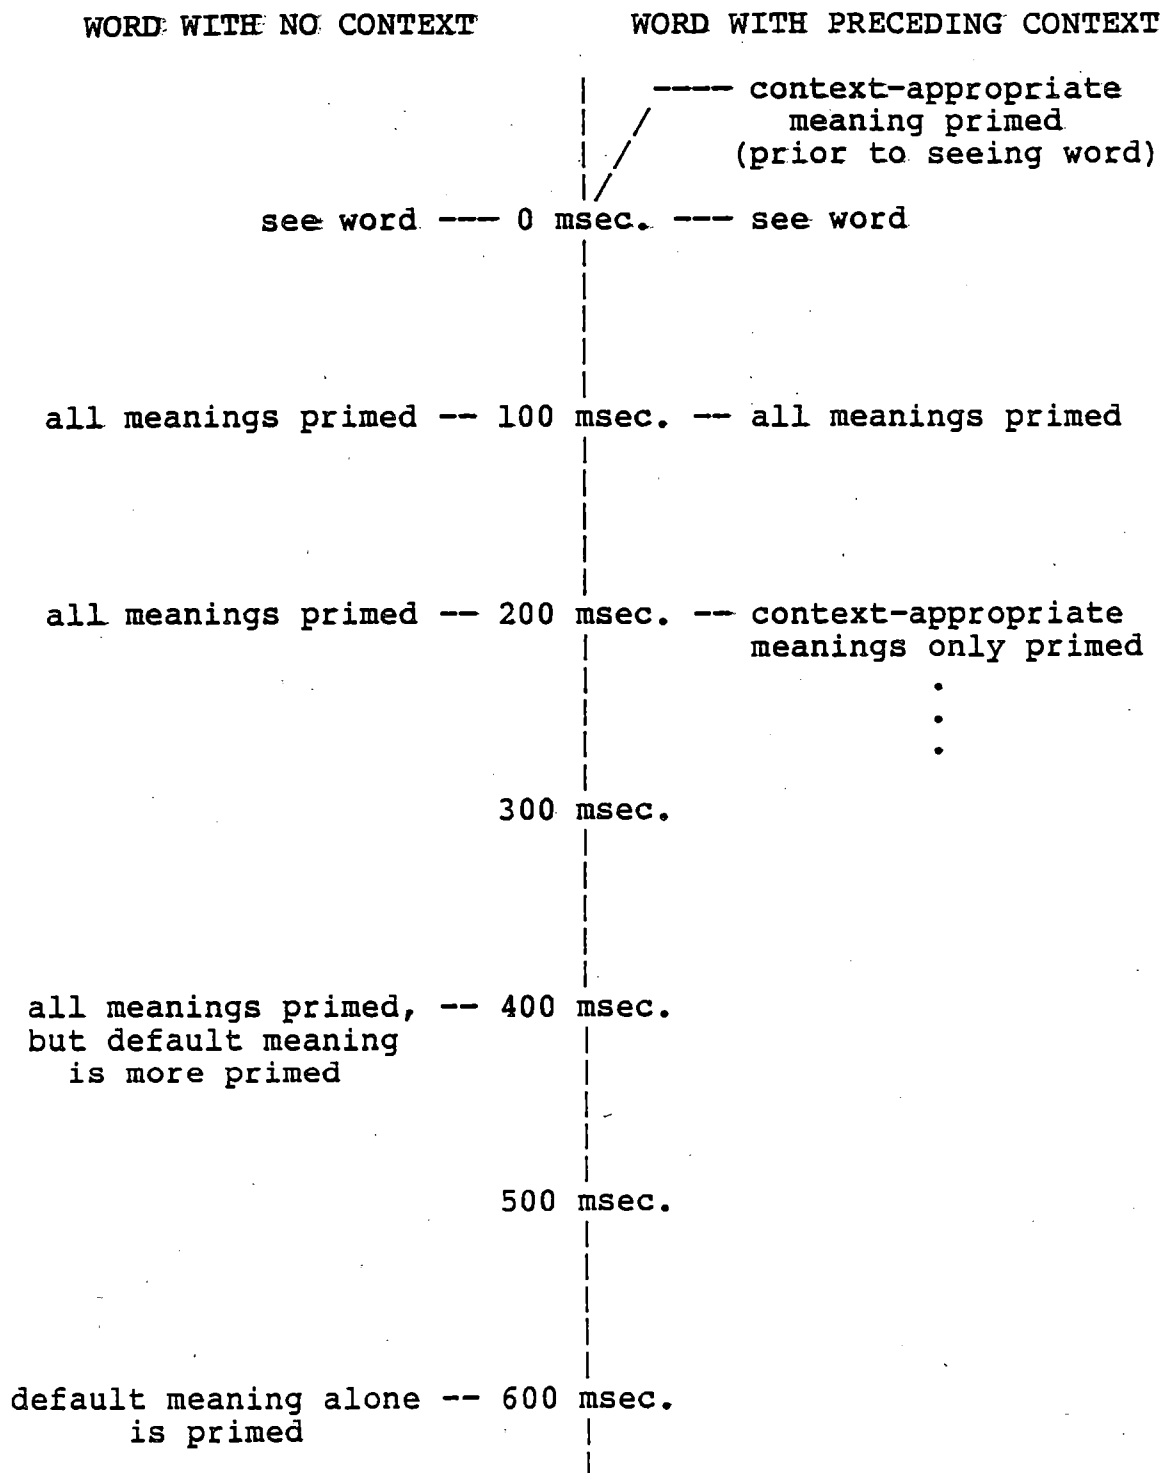


Figure 1. Lexical access timeline.

the other meanings are no longer available [Warren, 1977]. If a word is in a context (i.e., a sentence or phrase) which biases towards one of the meanings, another counter-intuitive process occurs: all meanings of an ambiguous word are accessed initially, and context is subsequently consulted to determine the most appropriate meaning (see Figure 1) [Swinney and Hakes, 1976; Tanenhaus, Seidenberg, & Leiman, 1979; Lucas, 1983]. As lexical access occurs, all meanings are primed regardless of syntactic category (e.g., "post the letter" vs. "the fence post"). This bottom-up-first, top-down-next process is used whether context is available before the ambiguous word is presented or after the ambiguous word is presented.

When an ambiguous word is presented after biasing context, it has been suggested that meanings which are inappropriate to context are actively suppressed [Tanenhaus, Seidenberg, & Leiman, 1979]. That is, they fade away much more quickly than if there had been no context at all. In other words, disambiguation would involve not only the identification of the correct meaning, but the immediate erasure of accessed but inappropriate meanings. The erasure is a special process which can only work with context.

We have proposed a modified version of the active suppression theory. We call the modified theory conditional retention [Granger, Holbrook, & Eiselt, 1984]. The conditional retention theory says that all meanings of an ambiguous word are retained until it is clear that one or more meanings are

appropriate to the whole context. Thus, if an ambiguous word appears in isolation, no meaning is inappropriate, so no meanings are suppressed. If an ambiguous word appears preceding or following a context which suggests only one meaning of the word, all other meanings will be actively suppressed. If an ambiguous word appears within context (i.e., text both precedes and follows the ambiguous word), a meaning will initially be selected which fits the context preceding the word. However, those meanings which do not fit the preceding context will not be actively suppressed until the rest of the context is available for final interpretation. Conditional retention offers an explanation as to why humans can understand texts with initially misleading contexts, as was pointed out in examples [1] and [2] above, while active suppression does not. Furthermore, experimental evidence from human subjects indicates that conditional retention provides a better explanation of human behavior than active suppression. (For a short discussion of the experimental evidence, see Appendix I.)

Because the lexical access findings indicate that all meanings are facilitated at first, with one meaning finally chosen, it seems as though all possibilities are pursued simultaneously in memory, and evaluated on the basis of a best fit with the current context. Thus, it cannot be true that correct word-senses are chosen by pursuing each possibility in turn until one fits the current context well. If this were so, we would expect that only one meaning would ever be facilitated when no context is available.

2.3 Inference as Memory Retrieval

Lexical access can be described as the retrieval of, evaluation of, and decision about specific, competing memories. The memories, in this case, are word-senses. In the same way, pragmatic inference decisions also depend upon the retrieval and evaluation of competing memories -- in this case, memories of events and event sequences. With both lexical and pragmatic memories, the evaluation consists of choosing the memory which most closely fits the current context. The choice is made through various evaluation metrics which seem to be available at both levels.

Spreading activation is a memory organization scheme which offers the ability to pursue many inference paths simultaneously, and has been employed in a number of models [e.g., Quillian, 1968; Fahlman, 1979; Charniak, 1983]. We use a spreading activation process in ATLAST to make inferences at the lexical and syntactic levels. A serious problem with spreading activation is that it can quickly lead to a combinatorial explosion of inferences if it has no inherent restrictions on which inferences will be pursued, or how far an inference will be pursued. We have addressed this problem within our system by having a separate process evaluate the inference paths which are activated, and thus controlling which inference paths will be pursued and which will be abandoned. The use of a spreading activation process for inference pursuit and another process for the evaluation of inferences has led to a new architecture for

processing and understanding text.

3.0 The New Architecture

The ATLAST model consists of three major processes: the Lexical Capsulizer, the Proposer, and the Filter. These processes run in parallel. These three processes were developed to incorporate three features: spreading activation memory, two-stage syntax, and parallel processing of syntax and semantics. Each of these features reflects a decision on how to make the model as psychologically valid as possible.

3.1 Activation and Inhibition

The first decision which affected ATLAST's architecture was the use of a spreading activation memory process, which we called Proposer. As discussed above, spreading activation allows several inference paths to be pursued simultaneously, which is apparently the way the human inference mechanism works. The Proposer has no inherent restrictions on which inference paths to follow and which to ignore.

Each path is pursued by Proposer until inhibited by the Filter, a process which runs concurrently with the Proposer. The Filter evaluates each inference path, using a set of evaluation metrics such as parsimony, cohesion, and specificity [Granger, 1980; Wilensky, 1983]. The metrics are plausibility indicators for making decisions about which inferences are to be pursued,

and for recognizing which inference paths intersect. In most cases, Filter will be able to detect and inhibit pursuit of particular inference paths as soon as they are proposed. In this way, the parallel operation of Proposer and Filter allows the concurrent pursuit of alternative inferences without suffering from the combinatorial explosion effects of pursuing too many inference paths. Though the idea of beginning pursuit on all inference paths instead of just the appropriate ones may seem both counter-intuitive and counter-productive, there are two arguments for using this approach. The first is that it would seem impossible to determine which inferences may be appropriate without first evaluating all inference possibilities. The second is that this approach is consistent with experimental studies of human behavior [Granger, Holbrook, & Eiselt, 1984].

Filter's evaluation metrics often help to disambiguate sentences. However, there are some cases in which the evaluation metrics come into conflict. One such case is with doubly embedded sentences, such as:

[3] The man the woman the child kissed met died.

Human readers are often unable to make sense of this sentence, even after several attempts. ATLAST, also, would have difficulty with such a sentence, because of the conflict between evaluation metrics. This suggests that the reason such sentences are difficult for people to understand is that human readers also have evaluation metrics which can conflict with one another.

3.2 Two-Stage Syntax

The second decision we made was to divide intra-phrasal and inter-phrasal syntactic decisions between two processes. The division allowed ATLAST to parse sentences which humans are able to parse, but also caused ATLAST to be unable to parse sentences which humans are unable to parse. One type of sentence which causes both ATLAST and human readers difficulty is the garden path sentence, such as text [3]:

[4] The horse raced past the barn fell.

A system which worked out all syntactic possibilities would have no problem understanding such a sentence; it would not make a decision when it came to the word "raced" as to whether "raced" begins a modifying clause or is the main verb. However, humans do make such a decision: they decide that "raced" is the main verb (the more common usage), but they are wrong, and they cannot parse the sentence. A parser which makes initial intra-phrasal decisions and later inter-phrasal decisions has the same problem that humans have [cf. Frazier & Fodor, 1978].

ATLAST has a Lexical Capsulizer which provides initial syntactic groupings, or "capsules", of words in a text. The Capsulizer activates much of the information immediately available about a given word, including how it can be used syntactically, phrases associated with the word, and so on. The syntactic information is accumulated by the Capsulizer as it processes words to make initial decisions about syntactic relationships within phrases (intra-phrasal syntax). Filter, on

the other hand, contains inference evaluation rules which include syntactic information as well as metrics for lexical and pragmatic inference decisions, so that Filter can make decisions about the syntactic relationships between phrases (inter-phrasal syntax).

At this point, it may seem that Filter does an unusual amount of work, and that there is no reason why another process could not be added to do inter-phrasal syntax. However, the kinds of decisions which Filter makes for guiding Proposer's search are the same kinds of decisions necessary for making inter-phrasal decisions; both tasks are simply a matter of applying evaluation metrics. In fact, many of Filter's decisions about Proposer's possible inference paths are based in part upon syntactic considerations (e.g., possible meanings of a word are limited by the syntactic category of each meaning). Filter is using Proposer's suggestions to fill missing parts of the interpretation, which include such syntactic considerations as Actor, Object, and Action. Thus, inter-phrasal syntax works better within Filter than as a separate process.

3.3 Concurrent Operation of Syntax and Semantics

The third decision which affected ATLAST's architecture was to have concurrent operation of inter-phrasal syntactic analysis (Capsulizer) and pragmatic inference generation (Proposer). To see the advantage of such parallelism, consider text [5]:

[5] The boy genius athlete was given a medal.

As an understander processes this text, it is unclear whether "boy" will be a noun or a modifier. A system might guess, but it is equally unclear which of the two "genius" will be, and which of the two "athlete" will be, until the word "was" is processed. (See [Gershman, 1977] for a thorough discussion of noun group analysis.) Yet, introspection indicates that an understander does not wait until a syntactic category is assigned to a word before beginning to build up a representation of the situation so far. Furthermore, once the sentence has been parsed through the word "was", it is still not clear whether the words preceding the word "was" make up an actor or an object (compare to "The boy genius athlete was running"), but an understander knows that the words up to "was" constitute a noun phrase. Thus, as has been pointed out by many other researchers [e.g., Charniak, 1983], syntactic decisions need not be made before semantic possibilities are explored.

3.4 The Processes and Their Functions

We have introduced ATLAST's three major processes: the Lexical Capsulizer, the Proposer, and the Filter (see Figure 2). Each of the decisions discussed above contributed to the design of the three processes and their functions. In keeping with the two-stage syntax, the Lexical Capsulizer provides initial capsules of words in a text using intra-phrasal syntactic rules. The Capsulizer recognizes a word by checking a letter string

against entries in the lexicon, and noticing a match. Capsulizer activates all the syntactic information about a given word. The syntactic information is accumulated by Capsulizer as it processes words to make intra-phrasal syntactic decisions. These decisions are made available to Filter (via the capsules), which uses the information to perform such tasks as differentiating between actors and objects.

The Proposer, which is the spreading activation mechanism, can be thought of as an emergent property of memory organization. When a match is found between a letter string from the text and the lexicon, Proposer triggers the alternate meanings of the word and pursues all possible inference paths from the associations with each meaning simultaneously. The inference paths lead to associated higher-level memory organization packets (MOPs) [Kolodner, 1984; Schank, 1982]. Each path is pursued until that path is inhibited by Filter.

The Filter, which runs concurrently with the Capsulizer and the Proposer, inhibits apparently unfruitful searches by the Proposer without expunging them, and allows the Proposer to pursue promising inference paths. The Filter applies evaluation metrics; the metrics are plausibility indicators for making decisions about which inferences are to be pursued, and for recognizing when inference paths intersect. When no more text is available, Filter "expunges", or rejects, all currently inhibited inferences.

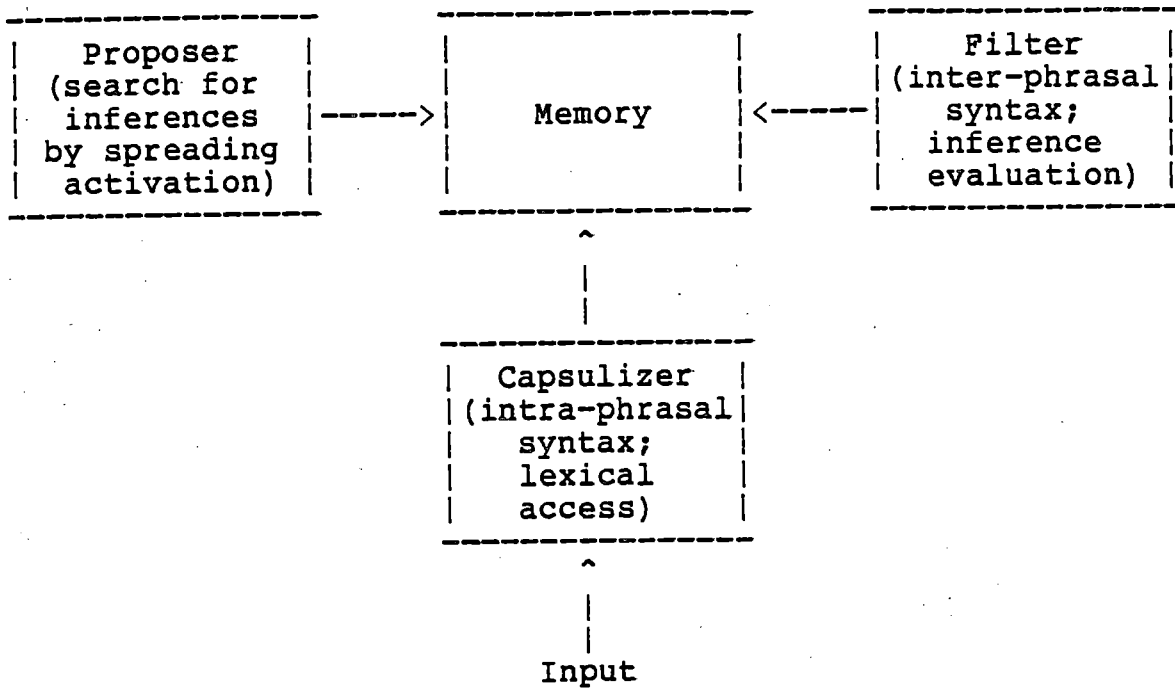


Figure 2. The organization of ATLAST's major components.

The Filter also maintains the various connections between the episodes which make up the alternative interpretations. The Filter makes inter-phrasal syntactic decisions using appropriate evaluation metrics. These metrics include rules about filling in slots in the representation, such as the Actor and Object slots. They also include rules which have to do with agreement of tense, number, and gender, as well as keeping track of referents across phrases, understanding when a phrase is modifying another phrase, and so forth.

Proposer, Filter, and Capsulizer all run simultaneously, although they may or may not be working on the same information at the same moment. For example, Filter cannot evaluate an inference path until Proposer begins to pursue it. Proposer cannot begin pursuing inference paths until Capsulizer finds a match between a letter string and a word in the lexicon. However, the concurrent operation of the three processes allows quick evaluation and inhibition of inferences, easy maintenance of alternative interpretations of text (and thus, easy supplanting of incorrect interpretations [Granger, 1980]), and fast, correct parsing of texts with which human readers have no trouble. In addition, the split syntax means that ATLAST has difficulty parsing the same types of texts as humans do.

Other models of language comprehension have tried to integrate some of the levels of inference behavior. There are models which integrate the syntactic and pragmatic levels [e.g., Dyer, 1982; Lebowitz, 1980], as well as models which integrate

lexical access and syntactic parsing [e.g., Small, Cottrell, & Shastri, 1982]. ATLAST is an implemented integration of language understanding on the lexical, syntactic, and pragmatic levels [cf. Charniak, 1983].

4.0 ATLAST: The Program

What follows is actual annotated run-time output from the first ATLAST prototype program. This example illustrates primarily how ATLAST disambiguates between two possible meanings of the word "bugs" in the text, "The CIA checked for bugs." In the interest of brevity and clarity, we use a very short text and just enough of a knowledge base to process this example. ATLAST is written in UCI-LISP on a Decsystem-20, so the parallelism which is so important to the theory is necessarily simulated in its implementation.

Processing begins

Input text is: (THE CIA CHECKED FOR BUGS *PERIOD*)

Capsulizer:

Retrieving lexical entry: THE
 No MOPs will be activated from lexical entry
 Begin sentence
 Begin noun phrase

Proposer:

No activity

Filter:

No activity

The first word, "the", is processed by ATLAST. Though Capsulizer recognizes that this marks the beginning of a noun phrase, there are no relevant structures in memory to be activated. Thus, Proposer and Filter are idle at this time.

Capsulizer:
Retrieving lexical entry: CIA

Proposer:
Initializing CENTRAL-INTELLIGENCE-AGENCY
Spreading from CENTRAL-INTELLIGENCE-AGENCY
Activating SPY-AGENCY

Filter:
No activity

In this cycle, the memory structure CENTRAL-INTELLIGENCE-AGENCY is activated as a result of reading "CIA". Proposer then begins to search along the links leading from CENTRAL-INTELLIGENCE-AGENCY for related memory structures, thus activating the more general SPY-AGENCY.

Capsulizer:
Retrieving lexical entry: CHECKED
Sending capsule
End noun phrase
Begin verb phrase

Proposer:
Initializing SEARCH
Spreading from SEARCH
Activating REMOVE
Spreading from SPY-AGENCY
Activating GET-OTHERS-SECRETS
Activating PRESERVE-OWN-SECRETS
Activating GENERIC-EMPLOYER

Filter:
ACTOR slot filled by CENTRAL-INTELLIGENCE-AGENCY

The next word, "checked", terminates the noun phrase and begins a verb phrase. Capsulizer sends a "capsule" consisting of the word-senses initially activated by the noun phrase (i.e., CENTRAL-INTELLIGENCE-AGENCY) to Filter. Filter, looking for an actor for this sentence, fills the slot with this noun-phrase capsule.

Capsulizer:
Retrieving lexical entry: FOR
No MOPs will be activated from lexical entry
Sending capsule
Begin prepositional phrase

Proposer:
Spreading from REMOVE
Activating REMOVE-OTHERS-LISTENING-DEVICE
Activating REMOVE-HEALTH-HAZARD
Spreading from GENERIC-EMPLOYER

Activating PRESERVE-HEALTHY-ENVIRONMENT
 Activating MANAGEMENT
 Spreading from PRESERVE-OWN-SECRETS
 Found connections at REMOVE-OTHERS-LISTENING-DEVICE
 Path from CENTRAL-INTELLIGENCE-AGENCY to SEARCH
 No MOPS activated from PRESERVE-OWN-SECRETS
 Spreading from GET-OTHERS-SECRETS
 Activating PLANT-OWN-LISTENING-DEVICE

Filter:

New path discovered: IPATH0
 Path from CENTRAL-INTELLIGENCE-AGENCY to SEARCH
 CENTRAL-INTELLIGENCE-AGENCY is a special case of SPY-AGENCY
 SPY-AGENCY has the goal PRESERVE-OWN-SECRETS
 PRESERVE-OWN-SECRETS has the plan REMOVE-OTHERS-LISTENING-DEVICE
 REMOVE-OTHERS-LISTENING-DEVICE is a special case of REMOVE
 REMOVE has the precondition SEARCH
 ACTION slot filled by SEARCH

The preposition "for" does not activate any new memory structures, but it does begin a modifying prepositional phrase. Capsulizer sends the verb component of the verb phrase (SEARCH) to Filter, which then assigns the capsule to the action slot.

Proposer, looking for intersections among the "wavefronts" of spreading activation, finds a connection, or inference path (IPATH0), between CENTRAL-INTELLIGENCE-AGENCY and SEARCH, and notifies Filter. Filter knows of only one inference path at this time, so there is no basis for comparison and evaluation of inference paths yet.

Capsulizer:

Retrieving lexical entry: BUGS

Proposer:

Initializing INSECT
 Initializing MICROPHONE
 Spreading from INSECT
 Found connections at REMOVE-HEALTH-HAZARD
 Path from INSECT to SEARCH
 No MOPS activated from INSECT
 Spreading from MICROPHONE
 Found connections at PLANT-OWN-LISTENING-DEVICE
 Path from MICROPHONE to CENTRAL-INTELLIGENCE-AGENCY
 Found connections at REMOVE-OTHERS-LISTENING-DEVICE
 Path from MICROPHONE to CENTRAL-INTELLIGENCE-AGENCY
 Path from MICROPHONE to SEARCH
 No MOPS activated from MICROPHONE
 Spreading from REMOVE-HEALTH-HAZARD
 Found connections at PRESERVE-HEALTHY-ENVIRONMENT
 Path from SEARCH to CENTRAL-INTELLIGENCE-AGENCY
 Found connections at INSECT
 Path from SEARCH to INSECT

No MOPS activated from REMOVE-HEALTH-HAZARD
 Spreading from REMOVE-OTHERS-LISTENING-DEVICE
 Found connections at PRESERVE-OWN-SECRETS
 Path from SEARCH to CENTRAL-INTELLIGENCE-AGENCY
 Found connections at MICROPHONE
 Path from SEARCH to MICROPHONE
 No MOPS activated from REMOVE-OTHERS-LISTENING-DEVICE
 Spreading from PLANT-OWN-LISTENING-DEVICE
 Found connections at MICROPHONE
 Path from CENTRAL-INTELLIGENCE-AGENCY to SEARCH
 Path from CENTRAL-INTELLIGENCE-AGENCY to MICROPHONE
 No MOPS activated from PLANT-OWN-LISTENING-DEVICE
 Spreading from MANAGEMENT
 No MOPS activated from MANAGEMENT
 Spreading from PRESERVE-HEALTHY-ENVIRONMENT
 Found connections at REMOVE-HEALTH-HAZARD
 Path from CENTRAL-INTELLIGENCE-AGENCY to INSECT
 Path from CENTRAL-INTELLIGENCE-AGENCY to SEARCH
 No MOPS activated from PRESERVE-HEALTHY-ENVIRONMENT

Filter:

New path discovered: IPATH1

Path from CENTRAL-INTELLIGENCE-AGENCY to SEARCH
 CENTRAL-INTELLIGENCE-AGENCY is a special case of SPY-AGENCY
 SPY-AGENCY can be viewed as GENERIC-EMPLOYER
 GENERIC-EMPLOYER has the goal PRESERVE-HEALTHY-ENVIRONMENT
 PRESERVE-HEALTHY-ENVIRONMENT has the plan REMOVE-HEALTH-HAZARD
 REMOVE-HEALTH-HAZARD is a special case of REMOVE
 REMOVE has the precondition SEARCH

New path discovered: IPATH2

Path from CENTRAL-INTELLIGENCE-AGENCY to INSECT
 CENTRAL-INTELLIGENCE-AGENCY is a special case of SPY-AGENCY
 SPY-AGENCY can be viewed as GENERIC-EMPLOYER
 GENERIC-EMPLOYER has the goal PRESERVE-HEALTHY-ENVIRONMENT
 PRESERVE-HEALTHY-ENVIRONMENT has the plan REMOVE-HEALTH-HAZARD
 REMOVE-HEALTH-HAZARD has the role-filler INSECT

New path discovered: IPATH3

Path from CENTRAL-INTELLIGENCE-AGENCY to MICROPHONE
 CENTRAL-INTELLIGENCE-AGENCY is a special case of SPY-AGENCY
 SPY-AGENCY has the goal GET-OTHERS-SECRETS
 GET-OTHERS-SECRETS has the plan PLANT-OWN-LISTENING-DEVICE
 PLANT-OWN-LISTENING-DEVICE has the role-filler MICROPHONE

New path discovered: IPATH4

Path from CENTRAL-INTELLIGENCE-AGENCY to SEARCH
 CENTRAL-INTELLIGENCE-AGENCY is a special case of SPY-AGENCY
 SPY-AGENCY has the goal GET-OTHERS-SECRETS
 GET-OTHERS-SECRETS has the plan PLANT-OWN-LISTENING-DEVICE
 PLANT-OWN-LISTENING-DEVICE has the role-filler MICROPHONE
 MICROPHONE is a role-filler of REMOVE-OTHERS-LISTENING-DEVICE
 REMOVE-OTHERS-LISTENING-DEVICE is a special case of REMOVE
 REMOVE has the precondition SEARCH

New path discovered: IPATH5

Path from SEARCH to MICROPHONE
 SEARCH is a precondition of REMOVE
 REMOVE has the special case REMOVE-OTHERS-LISTENING-DEVICE

REMOVE-OTHERS-LISTENING-DEVICE has the role-filler MICROPHONE
 New path discovered: IPATH6
 Path from SEARCH to INSECT
 SEARCH is a precondition of REMOVE
 REMOVE has the special case REMOVE-HEALTH-HAZARD
 REMOVE-HEALTH-HAZARD has the role-filler INSECT
 New path discovered: IPATH7
 Path from MICROPHONE to CENTRAL-INTELLIGENCE-AGENCY
 MICROPHONE is a role-filler of REMOVE-OTHERS-LISTENING-DEVICE
 REMOVE-OTHERS-LISTENING-DEVICE is a plan of PRESERVE-OWN-SECRETS
 PRESERVE-OWN-SECRETS is a goal of SPY-AGENCY
 SPY-AGENCY has the special case CENTRAL-INTELLIGENCE-AGENCY
 Parsimony metric -- IPATH7 explains more input than IPATH3
 Specificity metric -- IPATH4 more specific than IPATH1
 Parsimony metric -- IPATH0 shorter than IPATH4

Capsulizer reads the ambiguous word "bugs", which results in the activation of two word-senses: INSECT and MICROPHONE. Proposer's search has uncovered several new inference paths. When two different inference paths connect the same two word-senses, Filter applies inference evaluation metrics to the two paths to determine which of the two provides the better explanation of the input text. The rejected paths are "de-activated", or ignored, until later text results in activating that path again.

Capsulizer:

Retrieving lexical entry: *PERIOD*
 No MOPs will be activated from lexical entry
 Sending capsule
 End prepositional phrase
 End verb phrase
 End sentence

Proposer:

No activity

Filter:

OBJECT has competing slot fillers: INSECT vs. MICROPHONE
 Specificity metric -- IPATH7 more specific than IPATH2
 Parsimony metric -- IPATH5 explains more input than IPATH6
 Word-sense ambiguity resolution: MICROPHONE vs. INSECT
 All paths through INSECT have been de-activated
 The ambiguity is resolved -- MICROPHONE selected
 OBJECT slot filled by MICROPHONE

Capsulizer encounters the end of the text and sends to Filter a capsule containing the word-senses activated by the prepositional phrase. Filter determines that the capsule contains the object of the action SEARCH, and that this object is ambiguous. Filter attempts to resolve this ambiguity by applying the inference evaluation metrics to the remaining active inference paths. Because MICROPHONE and INSECT are now known to

be competing word-senses, Filter treats IPATH7 and IPATH2 as competing inference paths. That is, although IPATH7 connects MICROPHONE to CENTRAL-INTELLIGENCE-AGENCY and IPATH2 connects INSECT to CENTRAL-INTELLIGENCE-AGENCY, the two different paths are evaluated as if they connected the same two word-senses because INSECT and MICROPHONE were activated by the same lexical entry ("bugs"). For this same reason, IPATH5 is evaluated against IPATH6. This evaluation results in the two remaining inference paths containing INSECT to be de-activated, so Filter resolves the ambiguity in favor of MICROPHONE. Below is the active memory structure after all processing has ended, followed by the pointers into the structure.

Processing completed

Active memory structure:

Path from MICROPHONE to CENTRAL-INTELLIGENCE-AGENCY

MICROPHONE is a role-filler of REMOVE-OTHERS-LISTENING-DEVICE
 REMOVE-OTHERS-LISTENING-DEVICE is a plan of PRESERVE-OWN-SECRETS
 PRESERVE-OWN-SECRETS is a goal of SPY-AGENCY
 SPY-AGENCY has the special case CENTRAL-INTELLIGENCE-AGENCY

Path from SEARCH to MICROPHONE

SEARCH is a precondition of REMOVE
 REMOVE has the special case REMOVE-OTHERS-LISTENING-DEVICE
 REMOVE-OTHERS-LISTENING-DEVICE has the role-filler MICROPHONE

Path from CENTRAL-INTELLIGENCE-AGENCY to SEARCH

CENTRAL-INTELLIGENCE-AGENCY is a special case of SPY-AGENCY
 SPY-AGENCY has the goal PRESERVE-OWN-SECRETS
 PRESERVE-OWN-SECRETS has the plan REMOVE-OTHERS-LISTENING-DEVICE
 REMOVE-OTHERS-LISTENING-DEVICE is a special case of REMOVE
 REMOVE has the precondition SEARCH

Pointers to memory structure:

Actor: CENTRAL-INTELLIGENCE-AGENCY

Action: SEARCH

Object: MICROPHONE

5.0 Summary and Conclusions

As a model of the behavior of human inference processes during text understanding, ATLAST is quite different from those which have been proposed to date [e.g., Lebowitz, 1980; Dyer, 1982; Riesbeck, 1982]. The features which distinguish the ATLAST model from others can be summarized as follows:

1. ATLAST unifies inference processing at three distinct levels: the lexical, syntactic, and pragmatic levels.

2. The separation of intra-phrasal and inter-phrasal syntactic analysis enables ATLAST to process texts which humans understand and to make the same mistakes a human understander makes.

3. The use of a spreading-activation memory model allows ATLAST to pursue competing inference paths simultaneously until syntactic or semantic information suggests otherwise. Previous models of inference decision processes either left a loose end or chose a default inference when faced with an ambiguity [Granger, 1980; Granger, 1981; Wilensky, 1983; DeJong, 1979; Lebowitz, 1980; Dyer, 1982].

4. The concurrent operation of ATLAST's Capsulizer, Proposer, and Filter permits pragmatic interpretations to be evaluated independently of syntactic decisions. This parallel organization also allows immediate evaluation and inhibition of competing inference paths, thus minimizing combinatorial explosion effects.

5. ATLAST conforms to the results of controlled experiments on human subjects.

5.1 Future work

The ATLAST framework has been applied only to relatively short texts. We will be applying ATLAST to longer texts as well, to look at such factors as distance between inference points, which we believe will also affect inference processes, especially at the pragmatic level. We will also be applying ATLAST to different types of text to discover further rules for inference processing.

The model makes several predictions about what kind of behavior to expect from human readers. Because ATLAST is meant to be a model of human behavior, and not simply a program which

can read texts, ATLAST has difficulty parsing certain kinds of text which human readers also find very difficult to parse. The causes of ATLAST's difficulties are predictions of the causes of human readers' difficulties. We are currently designing and running several experiments on human subjects which allow us to test the predictions. For example, we are designing more lexical access experiments to decide how disambiguation occurs with longer texts, different experimental methodologies, and more specific predictions which will allow us to divide more specifically the roles of Proposer, Capsulizer, and Filter in disambiguation. We are also designing experiments to test our predictions about why garden path sentences are so difficult for human readers. Still another set of proposed experiments has to do with devising and testing various evaluation metrics, and testing what happens when the evaluation metrics conflict with one another.

ATLAST is a model of language understanding which employs the results of studies of many inference phenomena from several different fields of research. This approach helps solve many of the problems associated with inference decisions at all levels of processing. It is also a new framework within which to interpret and tackle unsolved problems through implementation and psychological experimentation.

6.0 References

- Charniak, E. Passing markers: A theory of contextual influence in language comprehension. Cognitive Science, 7, 171-190, 1983.
- DeJong, G.F. Skimming stories in real time: An experiment in integrated understanding. Ph.D. Thesis. Research Report #158. Department of Computer Science. Yale University, New Haven, CT, 1979.
- Dyer, M.G. In-depth understanding: A computer model of integrated parsing for narrative comprehension. Ph.D. Thesis. Research Report #219. Department of Computer Science. Yale University, New Haven, CT, 1982.
- Fahlman, S.E. NETL: A system for representing and using real-world knowledge. Cambridge, MA: MIT Press, 1979.
- Frazier, L., & Fodor, J.D. The sausage machine: A new two-stage parsing model. Cognition, 6, 291-325, 1978.
- Gershman, A.V. Conceptual analysis of noun groups in English. Proceedings of the 5th International Joint Conference on Artificial Intelligence, Cambridge, MA, 1977.
- Granger, R.H. When expectation fails: Towards a self-correcting inference system. Technical Report #162. Department of Information and Computer Science. University of California, Irvine, CA, 1980.
- Granger, R.H. Directing and Re-directing Inference Pursuit: Extra-textual Influences on Text Interpretation. Technical Report #171. Department of Information and Computer Science. University of California, Irvine, CA, 1981.
- Granger, R.H., Eiselt, K.P., & Holbrook, J.K. STRATEGIST: A program that models strategy-driven and content-driven inference behavior. Proceedings of the National Conference of Artificial Intelligence, Washington, D.C., 1983.
- Granger, R.H., Holbrook, J.K., & Eiselt, K.P. Interaction effects between word-level and text-level inferences: On-line processing of ambiguous words in context. Proceedings of the Sixth Annual Conference of the Cognitive Science Society, Boulder, CO, 1984.
- Hudson, S., & Tanenhaus, M. Ambiguity resolution in the absence of contextual bias. The Sixth Annual Conference of the Cognitive Science Society, Boulder, CO, 1984.
- Kolodner, J.L. Retrieval and organizational strategies in memory: A computer model. Hillsdale, NJ: Erlbaum, 1984.
- Lebowitz, M. Generalization and memory in an integrated understanding system. Ph.D. Thesis. Research Report #186.

Department of Computer Science. Yale University, New Haven, CT, 1980.

Lucas, M. Lexical access during sentence comprehension: Frequency and context effects. Proceedings of the Fifth Annual Conference of the Cognitive Science Society, Rochester, NY, 1983.

Quillian, M.R. Semantic memory. In M. Minsky (ed.), Semantic Information Processing. Cambridge, MA: MIT Press, 1968.

Riesbeck, C.K. Realistic language comprehension. In W.G. Lehnert & M.H. Ringle (eds.), Strategies for Natural Language Processing. Hillsdale, NJ: Erlbaum, 1982.

Schank, R.C. Dynamic Memory: A Theory of Reminding and Learning in Computers and People. New York: Cambridge University Press, 1982.

Schulenburg, D.A. Generating alternate interpretations: Directing inference during story understanding. Unpublished Master's Thesis. University of California, Irvine, 1982.

Small, S., Cottrell, G., & Shastri, L. Toward connectionist parsing. Proceedings of the National Conference on Artificial Intelligence, Pittsburgh, PA, 1982.

Swinney, D.A., & Hakes, D.T. Effects of prior context upon lexical access during sentence comprehension. Journal of Verbal Learning and Verbal Behavior, 15, 681-689, 1976.

Tanenhaus, M., Leiman, J., & Seidenberg, M. Evidence for multiple stages in processing of ambiguous words in syntactic contexts. Journal of Verbal Learning and Verbal Behavior, 18, 427-440, 1979.

Warren, R.E. Time and the spread of activation in memory. Journal of Experimental Psychology: Human Learning and Memory, 3 (4), 458-466, 1977.

Wilensky, R. Planning and Understanding. Reading, MA: Addison-Wesley, 1983.

7.0 Appendix I: Conditional Retention Experiment

One of the experiments we ran had subjects read two-sentence texts like texts [1] and [2] above. The last word in the first sentence of each of the texts was ambiguous (e.g., "bugs"). The second sentence was always written in such a way that only the last word would disambiguate the text. For example, in text [1], the disambiguating word is "microphone". In text [2], it is "roaches". Immediately after the subject read a text, they saw a pair of words. Each word was related to one of the meanings of the ambiguous word (e.g., "spy" and "ant"). The subject was to circle the word most closely related to the story's events. They did this as quickly as possible, because they thought they were being timed on their decision. After a subject had circled one of the words, they answered several questions related to the story events.

Active suppression would predict that stories in which the two sentences were biased toward different interpretations (conflicting context stories, such as text [2]), understanding would be difficult, because the initially inappropriate meaning, which is necessary to understand the story, is suppressed after the first sentence. If understanding is difficult, then story comprehension questions should have a high error rate. Furthermore, active suppression would predict that when the subject had to choose between the two words, the choice would be easy, because the inappropriate meanings would have been suppressed, so that the correct meaning would be the only one

primed. (If the subjects were not making their choice as quickly as possible, the error rate should in fact increase because the story interpretation difficulty should interfere.)

The conditional retention theory makes opposite predictions. Understanding the conflicting context stories should be easy, because both meanings are available. Thus, answering the comprehension questions should be easy. However, the error rate for the word choice task should be high because both meanings are still available. The results of this experiment are in Tables 1 and 2.

In the word choice task, the difference in error rate between the two-context condition and the other conditions was highly significant. However, the difference in error rate between the other conditions is not significant. This is exactly what conditional retention would predict. In the question-answering task, the difference in error rate between all four conditions was insignificant. This, too, is what conditional retention would predict. The results of this experiment agree with the conditional retention theory. We are also testing the conditional retention theory using several other methodologies. Experiments done by Hudson and Tanenhaus [1984] using a different methodology, confirm some parts of our results. Their experiments, however, were not designed to test for conditional retention, so they did not control for some of the important variables which would be necessary for a full confirmation.

Type of Text*	Distractor items (lexically unambiguous)	Neutral/ Bias (1 context)	Bias/ Same bias (1 context)	Bias/ Different Bias (2 contexts)
Error Rate	0%	7%	7%	54%

Table 1. Results of the word choice experiment.

Type of Text*	Distractor items (lexically unambiguous)	Neutral/ Bias (1 context)	Bias/ Same bias (1 context)	Bias/ Different Bias (2 contexts)
Error Rate	3.75%	6.25%	2.2%	4.9%

Table 2. Results of the question answering experiment.

*Described as the context type of the first sentence followed by the context type of the second sentence. A context could be neutral with respect to selecting the meaning of the ambiguous word, or biased toward one or the other meaning of the ambiguous word.