Recovering from Erroneous Inferences

Kurt P. Eiselt*
Irvine Computational Intelligence Project
Department of Information and Computer Science
University of California
Irvine, California 92717

Technical Report 87-06
February 1987

Abstract

Many models of natural language understanding make inference decisions as they process a text, but few models address the issue of how to correct their interpretation when later text reveals that earlier inference decisions are wrong. This paper describes how ATLAST, a marker-passing model of text understanding, approaches this problem. The keys to ATLAST's error recovery capability are a means for remembering the choices it could have made but didn't, and a means for initiating the re-evaluation of those previously rejected choices at the appropriate times. This paper also discusses some of the arguments for and against the psychological validity of a theory of inference retention in human text understanding.

*This research was supported in part by the National Science Foundation under grants IST-81-20685 and IST-85-12419 and by the Naval Ocean Systems Center under contracts N00123-81-C-1078 and N66001-83-C-0255.

1 Introduction

As we read, we make unconscious decisions about the meaning of ambiguous words, sentences, or passages based on incomplete information. Often those decisions are wrong and we must revise our understanding of the text. For example, consider the following simple story:

```
Text 1: Fred asked Wilma to marry him.
        Wilma began to cry.
```

Interpreting this text requires that a causal relationship between Fred's proposal and Wilma's tears be inferred. One such possible relationship is that Wilma was happy about Fred's proposal and was crying "tears of joy." Another equally likely inference is that Wilma was crying because she was saddened or upset by the proposal.[1]

Now consider this variation of Text 1:

```
Text 2: Fred asked Wilma to marry him.
        Wilma began to cry.
        She was saddened by the proposal.
```

Assuming that after processing the first two sentences of Text 2, the text understander has inferred that Wilma is happy, how does the understander resolve that inference with the contradictory third sentence?

One solution is to postpone making inferences for as long as possible so that potential conflicts are resolved before any decisions are made. However, this solution becomes less viable as texts increase in length. A better solution is to make inferences as the opportunities arise, then revise initial inferences if later text shows them to be incorrect. This paper describes how one model of text understanding, ATLAST, simplifies the error recovery process by remembering the alternative inferences it could have made but did not, and reconsidering those alternatives when subsequent text suggests they might now be correct.

Most models of natural language understanding fail to address the problem of recovery

[1] Experimental evidence indicates that either interpretation is equally likely when this text is presented to human subjects (Granger, Eiselt, & Holbrook, 1983; Granger & Holbrook, 1983).

from erroneous inferences, but there have been exceptions. Granger's ARTHUR (1980) was able to supplant incorrect inferences by using a map of pointers to all of the inferences generated during the processing of a text, whether or not they appeared in the final representation. O'Rorke (1983) proposed a design for a story understander called RESUND that used non-monotonic dependencies to correct false assumptions. Norvig's FAUSTUS (1983) temporarily stored rejected inferences using a process similar to the retention process discussed in this paper. FAUSTUS represented inferences as frames, and rejected frames were stored in a separate data base in case later text forced revision of earlier decisions. While previous papers on ATLAST have discussed how it makes decisions in the face of ambiguity (Eiselt, 1985; Granger, Eiselt, & Holbrook, 1986), this is the first paper to describe how ATLAST revises the interpretation when its decisions are contradicted by later text.

2 How it works

ATLAST's ability to revise its original interpretation in the face of new information depends upon two features: the ability to remember inference paths that it originally decided should not be part of its interpretation of the input text, and a mechanism for recognizing when these rejected paths should be reconsidered. Explaining how this works will be facilitated by a brief summary of the relevant aspects of ATLAST's architecture and the assumptions under which it runs.

2.1 Architecture and assumptions

ATLAST uses marker-passing to search a relational network for paths which connect meanings of open-class words from the input text. A single path is a chain of nodes, representing objects or events, connected by links, corresponding to relationships between the nodes. Any nodes in a path which are not explicitly mentioned in the text are events or objects which are inferred; therefore, these paths are called *inference paths.* A set of inference paths which joins all of the words in the text into a connected graph represents one possible interpretation of the text. In this respect ATLAST resembles a number of other models of text understanding that utilize marker-passing or spreading activation (e.g., Charniak, 1983; Cottrell, 1984; Hirst, 1984; Quillian, 1969; Riesbeck & Martin, 1986; Waltz & Pollack, 1984). The paths which form the current interpretation are called *active paths.*

For any given text, however, there may be a great number of possible interpretations, many of which are nonsensical. The problem then is determining which of the possible interpretations provides the best explanation of the text. ATLAST deals with this problem by applying inference evaluation metrics. These metrics are used to compare two competing inference paths and select the more appropriate one. Two inference paths compete when they connect the same two nodes in the relational network via different combinations of links and nodes. The path that fits better with the current interpretation is activated (i.e., it becomes part of the interpretation). The other path is de-activated but not discarded. Instead, that path is *retained* in order to facilitate error recovery as described in Subsection 2.2. The choice of one inference path over another is made as soon as ATLAST discovers that the two paths compete; ATLAST does not postpone inference decisions. As the marker-passing search mechanism finds more paths, ATLAST constructs an interpretation consisting of those paths which survive the evaluation process. When the marker-passing and evaluation processes end, the surviving active paths make up the final interpretation of the text.

Before they are even considered by the evaluation process, however, paths proposed by the marker-passing search process must pass minimum acceptability requirements. For example, ATLAST will ignore paths which contain cycles. Another example is that paths which connect components of different sentences (which ATLAST naively assumes to represent different states or events) must contain links denoting a causal relationship between the sentences (cf. Schank & Abelson, 1977). These constraints, along with a limit on the spread of marker passing, serve to limit combinatorial growth of the number of paths that could be discovered and evaluated.

In theory, the search for inference paths and their evaluation take place simultaneously. In practice, though, ATLAST simulates this concurrency by alternating between marker-passing and path evaluation. During each of these cycles, a new word is read from the input, its meanings are recalled and marked, and all markers in the network are passed a fixed distance. Any path discovered in this way is then examined to see if it competes with an active path in the interpretation as it stands at that time. If so, the evaluation metrics are applied and a choice between the two competing paths is made.

2.2 Retention and re-evaluation

The two keys to error recovery in ATLAST are the retention of previously rejected inference paths and the ability to re-evaluate possibly relevant retained paths at the appropriate times. Without a mechanism for knowing when and how to re-evaluate the retained paths, the retention feature alone provides no benefit.

There are two ways in which the re-evaluation of a retained path can be initiated. The first is through direct rediscovery of the retained path by the search process. Because the passing of markers begins in different places at different times during the processing of text, the same inference path may be discovered (or more appropriately, rediscovered) more than once. If a rediscovered path is not currently part of ATLAST's interpretation of the text (i.e., the path has been discovered earlier, rejected by the evaluation metrics, but retained), that path is re-evaluated against the competing path which is part of the interpretation. This rediscovery process initiates reconsideration of some of the retained paths, but it is not dependent upon retention because these paths would be reconsidered even if they had not been retained.

Some retained paths, though, will not be rediscovered, but the inferences made from later text may change the interpretation in such a way that these paths now should be included. ATLAST uses a method of "piggy-backing" the re-evaluation of these paths onto the evaluation of paths which are directly discovered or rediscovered by the search process. If a (re)discovered path is evaluated against a competing path in the current interpretation, any subpaths or superpaths of the (re)discovered path are also evaluated against the current interpretation. In this way, ATLAST attempts to limit re-evaluation to those paths that are currently relevant. (While constraining reconsideration to just those paths that completely contain or are completely contained by the (re)discovered path works for the example of Section 3, it has proven to be too restrictive for another text. In that case, one retained path which should have been part of the final representation was neither directly nor indirectly chosen for re-evaluation. Future work with ATLAST will include relaxing this constraint to see if its error recovery ability can be improved while still avoiding the re-evaluation of every retained path on every cycle.) Without the ability to force re-evaluation of paths rejected early in processing but not rediscovered later, ATLAST's final interpretation probably will be incorrect. Indirectly initiating the re-evaluation of previously rejected inference paths is essential to ATLAST's error recovery capability and is dependent upon inference retention.

2.3 The three steps to error recovery

Another way to view error recovery is as a three-step process (Norvig, 1983). The three steps are (1) recognizing that an error has occurred, (2) locating the source of the error, and (3) correcting the error. From this perspective, ATLAST's operation can be described as follows:

- *Recognizing that an error has occurred:* Each inference path has only two endpoints, and for any two given endpoints there will be at most one active inference path between them. When a new path is discovered (or an old path is rediscovered), the set of currently active paths is searched for a path which shares the same endpoints. If such a path exists, it is possible that the currently active path was incorrectly included in the representation of the text.

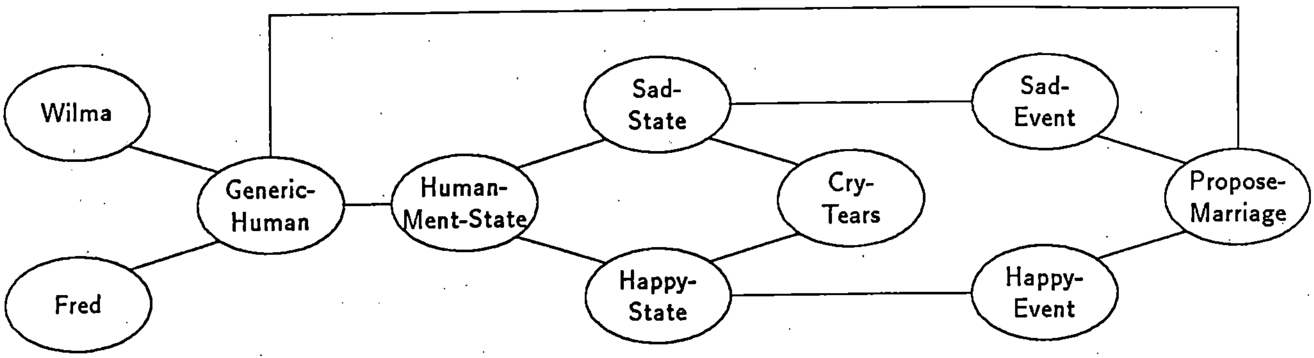

Figure 1: The organization of nodes in the memory structure for Text 3.

- *Locating the source of the error:* This step is effectively subsumed by the previous one. If competition between inference paths has been detected, the competition will always be between a path which is currently part of the representation and one which is not. If an error has in fact occurred, the source of the error will be the currently active path.
- *Correcting the error:* The competing paths are evaluated in the context of the current interpretation of the text (minus the active path being evaluated). The path which is more appropriate to the current interpretation is added to the interpretation, while the less appropriate path is added to the set of retained paths. If the interpretation changed as a result, then an error has been detected, located, and corrected.

3 ATLAST in action

An example of ATLAST processing a simple but potentially misleading text will illustrate the program's capacity for error recovery. This section describes the operation of ATLAST as it arrives at an interpretation for a simplified version of Text 2:

```
Text 3: Fred proposed.
        Wilma cried.
        Wilma was sad.
```

Although this is a simplified version of the original text (because ATLAST's syntactic abilities are limited), the relevant inference decisions should be the same for both texts. In the following example, many of the steps are left out for the sake of brevity (e.g., simple syntactic analysis and thematic role binding). The corresponding memory structure is shown in Figure 1.

As ATLAST reads the first sentence from left to right, it finds a path from `proposed` to `Fred`. At this point, there is no candidate interpretation for the text, thus no competing inference paths, so this path becomes the first member of the set of active paths:

```
New path discovered: path0
  Path from PROPOSE-MARRIAGE to FRED
    PROPOSE-MARRIAGE has the role-filler GENERIC-HUMAN
    GENERIC-HUMAN has the instance FRED
    Activating: path0
```

While processing the second sentence of the text, ATLAST finds a path denoting a causal relationship between `proposed` and `cried`. This path represents the inference that the crying results from a state of happiness which in turn results from the proposal of marriage. This path is added to the set of active paths:

```
New path discovered: path4
  Path from CRY-TEARS to PROPOSE-MARRIAGE
    CRY-TEARS is a result of HAPPY-STATE
    HAPPY-STATE is a result of HAPPY-EVENT
    HAPPY-EVENT has the instance PROPOSE-MARRIAGE
    Activating: path4
```

Next, ATLAST discovers a path which provides an alternate interpretation to that offered by the previous path. During this example, ATLAST was instructed to give preference to older paths over newer paths when no other evaluation metric was able to make a decision.[2] Thus, the newer path is not added to the set of active paths:

```
New path discovered: path5
  Path from CRY-TEARS to PROPOSE-MARRIAGE
    CRY-TEARS is a result of SAD-STATE
    SAD-STATE is a result of SAD-EVENT
    SAD-EVENT has the instance PROPOSE-MARRIAGE
Perseverer metric -- path4 older than path5
    De-activating: path5
```

[2]This tendency to prefer older inferences over newer ones results from the work on differences in human inference decision behavior noted in the previous footnote. The theory that was proposed to explain the differences suggests that some subjects prefer older inferences when faced with a choice between competing inferences, while other subjects prefer newer inferences. The people who prefer older inferences are called "perseverers" while those who prefer newer inferences are called "recencies." ATLAST is capable of modeling either kind of behavior by changing one of its evaluation metrics; it recovers from erroneous inferences in either mode.

ATLAST now finds a path which connects `cried` to `Wilma` and adds it to the set of active paths:

```
New path discovered: path9
  Path from CRY-TEARS to WILMA
    CRY-TEARS is a result of SAD-STATE
    SAD-STATE is an instance of HUMAN-MENT-STATE
    HUMAN-MENT-STATE is an attribute of GENERIC-HUMAN
    GENERIC-HUMAN has the instance WILMA
    Activating: path9
```

The interpretation now contains three paths: path 0, path 4, and path 9. There is a semantic contradiction among the active paths at this time in that path 9 is an inference that Wilma cried because she was sad while path 4 says that the tears were shed due to a state of happiness induced by the marriage proposal. ATLAST does not notice the contradiction because the two paths are not competing paths. This is the best interpretation based on the paths discovered so far. During the same cycle, a competing path from `Wilma` to `cried` is found. This new path, path 11, shares more nodes with other active paths than does its competing path, path 9; this is one criteria which ATLAST employs to decide which path explains more of the input. In this case, path 11 explains more input so it is added to the set of active paths and path 9 is moved to the set of retained paths:

```
New path discovered: path11
  Path from CRY-TEARS to WILMA
    CRY-TEARS is a result of HAPPY-STATE
    HAPPY-STATE is an instance of HUMAN-MENT-STATE
    HUMAN-MENT-STATE is an attribute of GENERIC-HUMAN
    GENERIC-HUMAN has the instance WILMA
More-reinforcement metric -- path11 has more shared nodes than path9
    De-activating: path9
    Activating: path11
```

As the final sentence is processed, ATLAST discovers a path connecting `proposed` to `sad`. This path is added to the set of active paths. In addition, this new path has three superpaths among the set of retained paths, and these paths are re-evaluated. One of these superpaths, path 5, is now preferred over the active path 4 because it is reinforced by path 15 (i.e., it contains the active path 15 as a subpath). Path 4 is moved from the set of active paths to the retained paths, and path 5 is moved from the retained paths to the active paths:

```
New path discovered: path15
  Path from SAD-STATE to PROPOSE-MARRIAGE
    SAD-STATE is a result of SAD-EVENT
    SAD-EVENT has the instance PROPOSE-MARRIAGE
  Also reconsidering: (path10 path5 path2)
    Activating: path15
Shorter-path metric -- path4 shorter than path10
    De-activating: path10
More-reinforcement metric -- path5 has more shared nodes than path4
    De-activating: path4
    Activating: path5
Shorter-path metric -- path0 shorter than path2
    De-activating: path2
```

The previous step demonstrates the need for inference path retention. Path 5 has been found directly several times prior to this point. Each time, the evaluation metrics have determined that path 4 fit better with the context. Now that path 15 is part of that context, path 5 is determined to be more appropriate than path 4. Had path 5 not been retained after being rejected earlier, it could not have been reconsidered at this time, nor would it ever have been reconsidered because the search process will not find path 5 again. If path 5 had not been retained, path 4 would incorrectly end up in the final representation of the story. In fact, this is what happens when ATLAST's retention capability is disabled while processing Text 3.

Continuing with the example, ATLAST finds a new path from cried to sad and adds it to the active paths:

```
New path discovered: path18
  Path from SAD-STATE to CRY-TEARS
    SAD-STATE has the result CRY-TEARS
  Also reconsidering: (path16 path13)
    Activating: path18
Shorter-path metric -- path15 shorter than path16
    De-activating: path16
```

```
Shorter-path metric -- path11 shorter than path13
    De-activating: path13
```

ATLAST then discovers the last new path to be added to the set of active paths. This path connects Wilma and sad. This new path also forces the reconsideration of one retained superpath, path 9, which is now preferred over its old competitor, path 11, because path 9 now shares more nodes with other active paths than does path 11. Path 9 is returned to the set of active paths and path 11 becomes a retained path, again illustrating the usefulness of inference retention:

```
New path discovered: path20
  Path from SAD-STATE to WILMA
    SAD-STATE is an instance of HUMAN-MENT-STATE
    HUMAN-MENT-STATE is an attribute of GENERIC-HUMAN
    GENERIC-HUMAN has the instance WILMA
  Also reconsidering: (path9)
    Activating: path20
More-reinforcement metric -- path9 has more shared nodes than path11
    De-activating: path11
    Activating: path9
```

The marker-passing mechanism will uncover nine more new paths to be considered, and rediscover many others, which will in turn force the re-evaluation of a number of retained subpaths and superpaths of those paths. However, none of these paths will be incorporated into the final interpretation of the text, which follows:

```
Active memory structure:
  Path from CRY-TEARS to WILMA
    CRY-TEARS is a result of SAD-STATE
    SAD-STATE is an instance of HUMAN-MENT-STATE
    HUMAN-MENT-STATE is an attribute of GENERIC-HUMAN
    GENERIC-HUMAN has the instance WILMA
  Path from SAD-STATE to WILMA
    SAD-STATE is an instance of HUMAN-MENT-STATE
    HUMAN-MENT-STATE is an attribute of GENERIC-HUMAN
    GENERIC-HUMAN has the instance WILMA
  Path from SAD-STATE to CRY-TEARS
    SAD-STATE has the result CRY-TEARS
  Path from CRY-TEARS to PROPOSE-MARRIAGE
```

```
    CRY-TEARS is a result of SAD-STATE
    SAD-STATE is a result of SAD-EVENT
    SAD-EVENT has the instance PROPOSE-MARRIAGE
  Path from SAD-STATE to PROPOSE-MARRIAGE
    SAD-STATE is a result of SAD-EVENT
    SAD-EVENT has the instance PROPOSE-MARRIAGE
  Path from PROPOSE-MARRIAGE to FRED
    PROPOSE-MARRIAGE has the role-filler GENERIC-HUMAN
    GENERIC-HUMAN has the instance FRED
```

The evaluation metrics used in this example, with the exception of the perseverer metric, all represent attempts to quantify different aspects of the principle of parsimony: explaining the greatest amount of input text with the least amount of representation. Variations of this principle have been employed by a number of diverse models of language understanding, including those of Crain and Steedman (1985), Granger (1980), Kay (1983), McDermott (1974), Quillian (1969), and Wilks (1978).

4 Retention issues

The principle of retaining rejected inference paths is inspired by experimental work which has led to a theory of lexical disambiguation called *conditional retention* (Granger, Holbrook, & Eiselt, 1984). According to this theory, lexical disambiguation is an automatic process in which all meanings of an ambiguous word are retrieved, the meaning most appropriate to the preceding context is chosen, and the other meanings are temporarily retained. In the case where the ambiguous word appears within a short text, the meanings are retained until the end of the text. Should later text contradict the initially chosen meaning, the retained meanings for that word are reconsidered in light of the updated context, and a new meaning is selected without repeating the lexical retrieval process. The theory of conditional retention thus offers an explanation of how readers can recover from an incorrect choice of word meaning without reprocessing the text. Because the choice of a word meaning will affect the inferences which are made during the understanding of a text, the theory of conditional retention has implications for making inference decisions at levels above the lexical level. Following this assumption, ATLAST uses the inference retention mechanism described in Sections 2 and 3 to recover from both incorrect lexical inferences (i.e., choices of word meaning) as well as erroneous pragmatic inferences.

However, the theory of conditional retention is by no means widely accepted, and the criticisms of conditional retention should be taken into consideration when evaluating ATLAST's utility as a cognitive model. One argument against conditional retention is a large body of experimental evidence which shows that, almost immediately after a meaning of an ambiguous word has been selected, the alternate meanings seem as if they had never been recalled (e.g., Seidenberg, Tanenhaus, Leiman, & Bienkowski, 1982; Tanenhaus, Leiman, & Seidenberg, 1979). This has been interpreted by some as proof that retention does not occur. On the other hand, these experiments were not specifically designed to look for evidence of retention. Also, as shown by Holbrook, Eiselt, Granger, and Matthei (to appear), the results of some experiments (e.g., Hudson & Tanenhaus, 1984) can be interpreted in such a way as to support the theory of conditional retention, though not conclusively. The one experiment to date which was designed to look for retention (Granger, Holbrook, & Eiselt, 1984) also yielded inconclusive results.

A frequent and deserved criticism of the conditional retention theory is that it offers no concrete answer to the question of how long alternate choices are retained; it says only that the choices are retained until the end of the text if the text is short. The experiment described by Granger, Holbrook, and Eiselt (1984) did not address this issue, but recent tests with ATLAST may point the way to an answer. In these tests, ATLAST was modified so that a path was given a time stamp indicating the cycle during which it was added to the set of retained paths. In addition, a limit was established on the number of cycles that a path could be retained without being reconsidered. Thus, a path retained on the first cycle, for example, was deleted from the set of retained paths if that path had not been re-evaluated by the end of the third cycle. If the path had been re-evaluated in that time but had not been added to the current interpretation, it was given an updated time stamp and returned to the set of retained paths.

While the example text of Section 3 was run without time stamps or limits, subsequent tests on this text, which consists of two two-word sentences and one three-word sentence, revealed that a minimum limit on retention of two cycles was sufficient to enable ATLAST to arrive at the same representation as that found in the example of Section 3. In a test with another text consisting of a five-word sentence followed by a four-word sentence, the minimum limit was four cycles. Obviously, there are too few data points to draw any conclusions, and the units of measure (e.g., cycles and words) are almost too coarse to be useful, but there is the slightest hint of a correlation between the amount of time that inferences are retained and the length of sentences in the text. Further work with ATLAST along these lines may lead to predictions that the duration of retention is tied to structural cues such as clause boundaries, or that the duration is independent of syntax and is controlled by the number of paths being retained, just to name two possibilities. In any case, if interesting predictions do arise from work with the model, it may be possible to test these predictions in the laboratory with human subjects.

Another problem with the conditional retention theory is that it assumes human readers recover from errors without backtracking and rereading the text. However, as Carpenter and Daneman (1981) demonstrate through studies of eye fixations of human subjects while reading, there are texts and conditions which cause a reader to backtrack when a semantic inconsistency is discovered in an ambiguous text. Carpenter and Daneman's error recovery heuristics include checking previous words that caused processing difficulty. The problem then becomes one of how to find the words that caused the inconsistency. Carpenter and Daneman state that their data allow for a model which utilizes memory of previous decisions, though this is not the only interpretation they offer (p. 152). Thus, while ATLAST differs in many ways from the model of Carpenter and Daneman, especially in regard to the issue of reprocessing the input text, the latter model at least recognizes the plausibility of the principle of retention in explaining a reader's ability to recover from incorrect inferences made while reading misleading text.

5 Conclusion

The principle of retaining rejected inference paths within the larger framework of a relational network provides a simple but effective mechanism for recovering from erroneous inferences during text understanding, but only if there is a way to locate and re-evaluate the retained paths at the appropriate times.

From a practical perspective, the principle of inference retention could be incorporated into new or existing text understanding systems in order to enable them to correct erroneous decisions. From a cognitive modeling perspective, however, the jury is still out on the issue of inference retention. While a model like ATLAST demonstrates the plausibility of the theory, only psycholinguistic experiments designed specifically to test for retention will be able to confirm or deny the validity of the theory.

6 References

Carpenter, P.A., & Daneman, M. (1981). Lexical retrieval and error recovery in reading: A model based on eye fixations. *Journal of Verbal Learning and Verbal Behavior, 20,* 137-160.

Charniak, E. (1983). Passing markers: A theory of contextual influence in language comprehension. *Cognitive Science, 7,* 171-190.

Cottrell, G.W. (1984). A model of lexical access of ambiguous words. *Proceedings of the National Conference on Artificial Intelligence*, Austin, TX.

Crain, S., & Steedman, M. (1985). On not being led up the garden path: The use of context by the psychological syntax processor. In D.R. Dowty, L. Karttunen, & A.M. Zwicky (Eds.), *Natural language parsing: Psychological, computational, and theoretical perspectives.* Cambridge, England: Cambridge University Press.

Eiselt, K.P. (1985). A parallel-process model of on-line inference processing. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, CA.

Granger, R.H. (1980). When expectation fails: Towards a self-correcting inference system. *Proceedings of the First Annual National Conference on Artificial Intelligence*, Stanford, CA.

Granger, R.H., Eiselt, K.P., & Holbrook, J.K. (1983). STRATEGIST: A program that models strategy-driven and content-driven inference behavior. *Proceedings of the National Conference on Artificial Intelligence*, Washington, DC

Granger, R.H., Eiselt, K.P., & Holbrook, J.K. (1986). Parsing with parallelism: A spreading-activation model of inference processing during text understanding. In J.L. Kolodner & C.K. Riesbeck (Eds.), *Experience, memory, and reasoning.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Granger, R.H., & Holbrook, J.K. (1983). Perseverers, recencies, and deferrers: New experimental evidence for multiple inference strategies in understanding. *Proceedings of the Fifth Annual Conference of the Cognitive Science Society*, Rochester, NY.

Granger, R.H., Holbrook, J.K., & Eiselt, K.P. (1984). 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.

Hirst, G. (1984). Jumping to conclusions: Psychological reality and unreality in a word disambiguation program. *Proceedings of the Sixth Annual Conference of the Cognitive Science Society*, Boulder, CO.

Holbrook, J.K., Eiselt, K.P., Granger, R.H., & Matthei, E.H. (to appear). (Almost) never letting go: Inference retention during text understanding. In S.L. Small, G.W. Cottrell, and M.K. Tanenhaus (Eds.), *Lexical ambiguity resolution in the comprehension of human language.* Los Altos, CA: Morgan Kaufmann.

Hudson, S.B., & Tanenhaus, M.K. (1984). Ambiguity resolution in the absence of contextual bias. *Proceedings of the Sixth Annual Conference of the Cognitive Science Society*, Boulder, CO.

Kay, P. (1983). *Three properties of the ideal reader.* Berkeley Cognitive Science Report No. 7. Cognitive Science Program, Institute of Cognitive Studies, University of California, Berkeley, CA.

McDermott, D. (1974). *Assimilation of new information by a natural language-understanding system.* Technical report no. 291. Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA. Cited in L. Birnbaum (1985). Lexical ambiguity as a touchstone for theories of language analysis. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, CA.

Norvig, P. (1983). Six problems for story understanders. *Proceedings of the National Conference on Artificial Intelligence*, Washington, DC

O'Rorke, P. (1983). Reasons for beliefs in understanding: Applications of non-monotonic dependencies to story processing. *Proceedings of the National Conference on Artificial Intelligence*, Washington, DC

Quillian, M.R. (1969). The teachable language comprehender: A simulation program and theory of language. *Communications of the ACM*, *12*(8), 459-476.

Riesbeck, C.K., & Martin, C.E. (1986). Direct memory access parsing. In J.L. Kolodner & C.K. Riesbeck (Eds.), *Experience, memory, and reasoning.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Seidenberg, M.S., Tanenhaus, M.K., Leiman, J.M., & Bienkowski, M. (1982). Automatic access of the meanings of ambiguous words in context: Some limitations of knowledge-based processing. *Cognitive Psychology*, *14*, 489-537.

Schank, R.C. & Abelson, R.P. (1977). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures.* Hillsdale, NJ: Lawrence Erlbaum Associates.

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

Waltz, D.L., & Pollack, J.B. (1985). Massively parallel parsing: A strongly interactive model of natural language interpretation. *Cognitive Science*, *9*(1), 51-74.

Wilks, Y. (1978). Making preferences more active. *Artificial Intelligence*, *11*, 197-223.