

Toward a Unified Theory of Lexical Error Recovery

Kurt P. Eiselt

College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280
eiselt@cc.gatech.edu

Jennifer K. Holbrook

Department of Psychology
Albion College
Albion, MI 49224
jholbrook@albion.bitnet

Abstract

The ambiguity inherent in natural language requires us to make many decisions about the meaning of what we hear or read. Yet most studies of natural language understanding have assumed that although language may be ambiguous, we always make the right choice when faced with a decision about ambiguity. Consequently, very little is said about how to recover from incorrect decisions. In this paper we look at two rare examples of investigations into recovering from erroneous decisions in resolving lexical ambiguity. After examining the corresponding theories, we find that what at first appear to be competing theories can in fact be resolved into a unified theory of lexical error recovery based upon a highly parallel architecture for language understanding.

Why error recovery?

Researchers in natural language understanding have often assumed that although spoken or written text may be ambiguous in many ways, it is not misleading. In other words, it is assumed that when the understander resolves an ambiguity, the choice that is made is always the correct one. Consequently, the question of how an understander can correct its mistakes in interpretation is seldom asked. Yet to ignore the issue of error recovery in natural language understanding is to ignore one of the most vital aspects of the human language understander. As the following example, due to Lashley (1951), illustrates, spoken or written text is often misleading, and humans frequently appear to be able to recover from mistakes in semantic interpretation quite gracefully, without any conscious awareness that a mistake has been made:

Text 1 Rapid righting with his uninjured hand saved from loss the contents of the capsized canoe.

People will often hear the second word as “writing” and realize their mistake only when they hear “the capsized canoe.” This ability to recover from errors serves us well. Ambiguity in natural language allows us to economize in our communications by eliminating much of what could be said or written and relying on the listener or reader to supply the missing knowledge. By doing so, however, we also increase the potential for misunderstanding. The ability to recover from erroneous word sense decisions when resolving ambiguities is an important compensatory mechanism that allows useful

communication to continue when the spoken or written text is misleading. Understanding the error recovery capability of the human language processor will aid in the development of more useful language understanding systems, and will lend insight into the architecture of the human language processor itself. In this paper we examine two models of lexical error recovery in text understanding, and show how the two models might be merged into a single unified model of lexical error recovery under a highly parallel architecture for language understanding.

Error recovery without text reprocessing

As Text 1 above clearly indicates, it is not necessary to reprocess misleading input in order to recover from erroneous decisions. When Text 1 is presented aurally, the listener who hears “writing” instead of “righting” will be unable to reprocess earlier text when the mistake is later revealed after hearing “the capsized canoe.” Because the memory of the verbatim text is available for only a short period of time (e.g., only until a clause boundary, according to Jarvella (1970)), some mechanism other than reprocessing must account for the listener’s ability to recover from this mistake.

A theory of lexical access and disambiguation, called *conditional retention*, accounts for a human understander’s ability to recover from an incorrect choice of word meaning without reprocessing the text (Granger, Holbrook, & Eiselt, 1984; Holbrook et al., 1988). According to this theory, all meanings of an ambiguous word are retrieved, the meaning most appropriate to the preceding context is chosen, and the other, less appropriate meanings are temporarily deactivated but retained. If later text contradicts the initially chosen meaning, the retained meanings are reconsidered in light of the updated context and a new meaning selected.

The combined deactivation and retention of the meanings not chosen accomplishes two goals for the language understander. First, it permits the processing at the lexical level to continue to make immediate decisions about the meanings of subsequent words in the context of a single, plausible interpretation of the preceding text instead of multiple interpretations of varying plausibility. If the unchosen meanings were not completely deactivated, there could be resulting confusion in making decisions about new word meanings. Second, it allows the retained decisions to be used by other processes in correcting wrong decisions made by the original process

without reprocessing the original input text, at least for a short time. Tracking retained meanings allows the error recovery to be done without maintaining separate copies of all possible interpretations of the text processed so far, thus reducing both storage and processing overhead.

One argument against conditional retention is that several cross-modal lexical priming (CMLP) experiments have shown that, within 200 msec after a meaning of an ambiguous word has been selected, the alternate meanings are as inactive as unrelated concepts that are used for comparison—they are as inactive as if they had never been activated in the first place (Onifer & Swinney, 1981; Seidenberg et al., 1982; Swinney, 1979; Tanenhaus, Leiman, & Seidenberg, 1979). This model of lexical access and disambiguation is called *active suppression*. There are two problems with using this body of data to argue in favor of the active suppression model over conditional retention. The first is that most of these studies used texts that ended with the ambiguous word (cf. Onifer & Swinney, 1981).¹ For such texts, then, meaning selection will always occur at the end of the text when there is no chance of disconfirming information. Under this condition, the conditional retention theory also holds that unselected meanings will be forgotten. But these experiments do not test what happens when the ambiguous word is embedded in the text, with more text to follow the ambiguous word. Hudson and Tanenhaus (1984) found that meanings were available longer when preceded by neutral context and embedded within the text. Holbrook (1989) has found evidence that unselected meanings are retained when the ambiguous word is embedded in a longer text. Thus, the CMLP studies that best support active suppression alone did not include materials that would test for conditional retention under appropriate conditions. The second problem is that while the CMLP studies that support active suppression do in fact demonstrate that unselected word meanings are very quickly deactivated at the end of a text, these experiments do not thoroughly test whether the meanings have been completely forgotten. They use only one type of measure (reaction time) on two closely-related tasks (lexical decision and naming).

The experiments above do not address the question of how recently-deactivated word meanings might differ from those which have been inactive for a much longer duration. This difference might appear as different degrees of sensitivity to re-activation while processing additional input, or some other quality that does not correspond to the relative activation level and is therefore immeasurable in CMLP studies. Furthermore, lexical decision tasks and lexical naming tasks are good for acquiring data about subjects' response times, which in turn correlate to the degree of facilitation of the individual word meanings. However, the conditional retention theory suggests that retention may not show up as facilitation because retained meanings have been de-activated. This suggests the need for an experimental methodology other than

¹Although Onifer and Swinney (1981) used texts that continued for several words after the ambiguous word, their materials were not as carefully controlled as has become common; specifically, meaning frequency, number of syllables of probes, and similar issues could have erased effects of meaning retention.

the lexical decision or lexical naming task.

New experimental evidence

We designed a study which uses a binary forced-choice task to test for conditional retention of unselected meanings across sentential boundaries. The forced-choice task offered the ability to detect indirectly the existence of conditional retention by studying the subjects' decisions instead of their response times. The use of texts with sentences following the ambiguous word allowed us to test whether suppression of unselected word meanings always occurs at the end of a sentence, or whether meanings can be retained when information with the potential to change the representation is at hand.

In this experiment, subjects were asked to read short texts of one to three sentences in length. Experimental texts were two sentences in length, and the others were filler texts. Filler texts were of varying length in order to ensure that subjects would not know whether there was more information to be added to the representation after their word choice. This provided the appropriate conditions under which retention is thought to occur. There were two types of experimental texts that are important to the current discussion:

- **Consistent bias surrounds ambiguous word:** The context preceding the ambiguity biases towards one meaning but does not preclude the other. The context following the ambiguity provides information requiring the original biased-for meaning.

Example: Mary realized that she had examined the wrong bat. She took it back and got one that was aluminum.

- **Conflicting bias surrounds ambiguous word:** The context preceding the ambiguity biases towards one meaning but does not preclude the other. The context following the ambiguity provides information consistent with the early context but requiring the unselected meaning.

Example: Mary realized that she had examined the wrong bat. She took it back and got one that was male.

The texts were presented on a computer monitor a few words at a time, with each group of words replacing the group before it. An information probe was displayed on the monitor at one of two points in the text: either between the ambiguous word and the disambiguating text or after the disambiguating text. The probe consisted of a pair of words, and the subject's task was to decide which of the two words was more related to the text. The choice was indicated by pressing one of two buttons, each corresponding to one of the two probe words. An example of materials presentation is shown below:

Mary realized
that she had examined
the wrong bat. She took
it back and got
CAVE PITCH
one that was
male.

The first line would appear, centered on the monitor, for 640 msec. The second line would then replace the first line for

consistent bias surrounds (n=18)	Cell 1		Cell 2		Cell 3	
	correct word	incorrect word	correct word	unrelated word	incorrect word	unrelated word
CR	75	25	100	0	100	0
AS	100	0	100	0	50	50
data	61	39	89	11	94	6

conflicting bias surrounds (n=27)	Cell 4		Cell 5		Cell 6	
	correct word	incorrect word	correct word	unrelated word	incorrect word	unrelated word
CR	25	75	100	0	100	0
AS	0	100	50	50	100	0
data	33	67	70	30	85	15

Table 1: Predictions and results at early probe point during forced choice task.

640 msec, and so on. When the two capitalized probe words appeared on the screen, the subject would press a designated key on the left side of the computer keyboard if he or she thought the word on the left was more appropriate to the text than the word on the right, or would press a designated key on the right side of the keyboard if the word on the right was thought to be more appropriate. Presentation of additional text did not continue until the subject pressed one of the two keys.

An early probe point and a late probe point were used for each text in a between-subjects design. The early probe point occurred some time after the ambiguous word but before the disambiguating information in the second sentence. The late probe point occurred at the end of the second sentence, after the disambiguating information had been presented. (For the sake of brevity, we will discuss only the results obtained at the early probe point in detail.) The probe word pairs were rotated between three types of words: a word related to the meaning of the ambiguous word that was correct at the end of the text, a word related to the incorrect meaning at the end of the text, and a word unrelated to either meaning or to the text as a whole. In the example above, the three words used were "PITCH," "CAVE," and "PAIR." "PITCH" is related to the baseball meaning of "bat," and is more appropriate at the early probe point.² "CAVE" is related to the animal meaning of "bat," and is more appropriate to the text at the late probe point. "PAIR" is unrelated to either meaning of

²The bias of the first sentence of each text was guaranteed with two groups of informants. The first group, naive judges, read and paraphrased each sentence. Those sentences which were paraphrased to reflect one meaning in 100% of the paraphrases were read by a second group of informants. These judges were informed that each sentence contained an ambiguous word. Each sentence read by this second group was followed by the two target words that were semantically related to the ambiguous word at the end of the sentence. The informants were asked to choose the word more related to the meaning of the sentence, or to indicate that both were equally related. A text was used in the experiment if the intended target word was chosen more than 80% of the time and was never judged as related to the other target word.

the ambiguous word or the text.

In this experiment, the theories of conditional retention and active suppression are set up as opposing theories, so it is useful to compare the two theories' predictions of the outcomes for the different conditions of this experiment. The predictions of the subjects' responses at the early probe point, and the data gathered at this point are summarized in Table 1. This table refers to three types of word stimuli. **Correct** word stimuli are words related to the meaning of the ambiguous word that could be integrated with context at the conclusion of the text. **Incorrect** word stimuli are words related to the meaning of the ambiguous word that could not be integrated with context at the conclusion of the text. **Unrelated** word stimuli are words that were unrelated to either meaning of the ambiguous word at any point in the text. The tables give predicted choices as percentages of the total number of responses.

The first type of text reported in Table 1, called "consistent bias surrounds," was designed so that the context which occurs before the ambiguous word is encountered biases towards one meaning of the ambiguous word. Thus, the reader will have enough information from the text on which to base a decision, and will choose the meaning which is more related to the previous context. The context which follows the ambiguous words for these texts agrees with the context that precedes the ambiguous words, so the meaning choice that was made remains correct throughout the text.

Conditional retention (CR) predicts that the unselected meaning of the ambiguous word will be retained throughout the text, and that this will cause interference in the forced-choice task between the correct and incorrect probe words (Cell 1). The simplest model of this interference is that it will be reflected in 25% of the responses (the intermediate point between a prediction of no interference, which would be reflected in 0% of the choices, and complete interference, which would be reflected in 50% of the choices). However, when the correct word is paired with the unrelated word (Cell 2), there is no reason to select the unrelated word over the correct word because it is not being retained. Therefore, the correct

word should always be chosen in this condition. When the incorrect word is paired with an unrelated word (Cell 3), the incorrect word will always be chosen by virtue of its relationship to the retained meaning of the ambiguous word.

Active suppression (AS) makes a different set of predictions. When a meaning for the ambiguous word is selected, the unselected meaning will be actively suppressed. At the point of the forced-choice task, the suppressed meaning should have no effect on the word chosen in the task. Thus, in Cell 1, active suppression predicts that the correct word will be chosen 100% of the time. In Cell 2, the subject again will always choose the correct word. In Cell 3, there is no reason to suppose that the incorrect word is chosen with more probability than the unrelated word. The conditional retention theory is supported by the results for all three conditions, while the active suppression theory is supported only by the results of Cell 2.

The other type of text reported in Table 1, called "conflicting bias surrounds," was designed so that the context which occurs before the ambiguous word is encountered biases towards one meaning of the ambiguous word. During the meaning decision process for the ambiguous word, the meaning that is a better fit with the previous context will be chosen. The context which follows the ambiguous words for these texts disagrees with the context that precedes the ambiguous words, so the meaning that is contextually appropriate at the early probe point will be inappropriate at the end of the text. Correspondingly, the probe word that is correct at the end of the text in "consistent bias" texts is incorrect at the end of "conflicting bias" texts, and the probe word that is incorrect in the former case is correct in the latter.

In Cell 4, conditional retention and active suppression both predict a significant effect of context type and target type. Although significance did not obtain, a binomial probability test supported the conditional retention theory. Conditional retention predicted this difference in Cell 5, but active suppression did not. In Cell 6, both theories predicted the finding. Conditional retention is again supported by all three conditions, whereas active suppression is only supported by one condition.

Overall, the weight of the evidence from this experiment clearly supports the conditional retention theory. The two theories predicted different results in four of the six different test conditions shown in Table 1: Cells 1, 3, 4, and 5. Analysis of the data showed that the conditional retention theory predicted the results better than the active suppression theory did in all four cells. An analysis of the data obtained with a late probe point also supported the conditional retention theory (Eiselt, 1989; Holbrook, 1989). Holbrook (1989) performed two additional related experiments which also support the conditional retention theory over active suppression. Still further support comes from another experiment which used a divided visual field methodology to find that the time course of the activation and suppression of two meanings of an ambiguous word followed significantly different paths in the different hemispheres of the brain (Burgess & Simpson, 1988). When the ambiguous word was presented only to the right visual field, and therefore to the left hemisphere of the brain, the more likely meaning was activated while the less

likely meaning was suppressed, much like active suppression. When the ambiguous word was presented to the left visual field, however, the right hemisphere kept the less likely meaning active and suppressed the more likely meaning, indicating that a form of retention was taking place.

Error recovery with text reprocessing

The evidence offered above strongly indicates that the human language processor does have a mechanism which would enable recovery from errors in lexical ambiguity resolution without text reprocessing. Yet common sense tells us that error recovery by backtracking and reprocessing is also a viable option. In fact, studies of human readers' eye fixations during the reading of misleading texts indicate that one method used by readers for recovery from lexical decision errors is simple reprocessing of the text. The following passage was used in an experiment by Carpenter and Daneman (1981, p. 137):

Text 2 The young man turned his back on the rock concert stage and looked across the resort lake. Tomorrow was the annual one-day fishing contest and fishermen would invade the place. Some of the best bass guitarists in the country would come to this spot. The usual routine of the fishing resort would be disrupted by the festivities.

Subjects in this experiment were asked to read passages such as the one above while the duration and location of their eye fixations were automatically recorded. In the example above, most readers initially interpreted the word "bass" as a kind of fish because the preceding text is biased toward this interpretation. The interpretation is contradicted, however, by the next word, "guitarists," which forces a reinterpretation of "bass" as a low-frequency musical note. Carpenter and Daneman found that most readers' eyes fixated on "bass," then moved forward to and fixated on "guitarists," then *regressed* back to "bass," moved forward to "guitarists" again and continued reading the remainder of the passage.

The reprocessing heuristic is just one of several error recovery heuristics proposed by Carpenter and Daneman. Another heuristic involves making a larger-than-normal inference encompassing both the inconsistent concept with the preceding text. This is done, they say, if the contradiction is only "mildly semantically inconsistent" and does not involve a syntactic inconsistency (Carpenter & Daneman, 1981, p. 141). Still another error recovery heuristic is to continue reading the text with the expectation that later information will resolve the inconsistency.

In addition to permitting error recovery without reprocessing, Carpenter and Daneman's theory also demonstrates the need for retention of some sort. Their proposed heuristic of checking the previous text for words that caused processing difficulties, such as ambiguous words, does not specifically address how the reader knows which word or words to reread, but Carpenter and Daneman theorize that difficulties encountered during processing may leave a memory trace which makes finding the ambiguous word much easier. The Carpenter and Daneman model follows the premise that the

activation levels of those concepts not selected for use in the interpretation either decay or are actively dampened to a base level, which would preclude the possibility that the memory trace is represented as activation, so this model strongly suggests a retention mechanism which is related to but distinct from the activation mechanism.

Integrating the models

We have seen experimental evidence for two different theories of recovery from erroneous word sense choices, one relying on reprocessing of the text, and the other relying on conditional retention. If we believe the evidence, then both approaches to error recovery are used by the human language understander, although perhaps in different situations. This apparent dichotomy gives rise to the following question: how does the language understander know the difference between these two situations? That is, how does the reader know when to backtrack and when to use retention to recover from an erroneous lexical decision?

Characterizing the differences

Consider again the two experiments discussed above. Target texts used by Carpenter and Daneman were multi-sentence passages containing an ambiguous word that had two different pronunciations and a different word sense associated with each pronunciation. The passages were constructed so that the initial context more strongly primed for one meaning of the ambiguous word. In half of the passages, information that was inconsistent with the primed meaning followed the homograph, while in the other half of the passages, the information following the ambiguous word was consistent with the primed meaning. A single passage was presented at once, thereby allowing the reader to backtrack as necessary. This construction is illustrated by Text 2, presented earlier. Of particular interest here is how closely the disambiguating information follows the ambiguity. In Text 2, the disambiguating information is the word immediately following the ambiguity. A survey of Carpenter and Daneman's target texts reveals that the disambiguating information was seldom separated from the ambiguity by more than two words, it was always in the same sentence as the ambiguity, and it was usually in the same clause as the ambiguity.

On the other hand, the target texts used in our experiment consisted of two sentences. The ambiguous word was always the last word of the first sentence, and the disambiguating information was contained in the second sentence. The text was presented to the subject in parts, so that only a few words of the entire text were available at any one time. In addition, the presentation was constructed so that there was at least one line of text presented on the monitor after the line containing the ambiguous word and before the line containing the disambiguating information. In other words, there was always at least a 640 msec delay between the reading of the ambiguity and the reading of the disambiguating information—long after active suppression of less appropriate word meanings should have taken place.

Resolving the differences

The essential difference between the two experiments, at least for the purposes of this discussion, can be summarized as follows: Carpenter and Daneman's experiment establishes a situation conducive to error recovery by reprocessing, and subjects demonstrated reprocessing behavior. Our experiment prevented subjects from backtracking, and subjects exhibited behavior consistent with conditional retention. In either case, *the subjects appeared to make the best use of whatever information was available to them to interpret the target text correctly.*

As noted above, Carpenter and Daneman propose a suite of error recovery heuristics, of which reprocessing is just one example. We propose, however, that the two different error recovery techniques described herein, reprocessing and retention, can be explained by a single mechanism. Recall that Carpenter and Daneman suggest that processing difficulties, such as word sense ambiguities, leave a memory trace to the ambiguous word which is independent of the activation levels of the word senses themselves. This is supported by their observations of subjects' eye movements which indicate that subjects did not simply reread the text; instead, the eye movements indicate that the subjects searched selectively for the source of the ambiguity.

It is not entirely clear that this memory trace or pointer to the ambiguous word is different from the retained but unchosen meanings of that ambiguous word. However, because research indicates that memory for verbatim text is limited by clause boundaries (Jarvella, 1970), and Carpenter and Daneman's experiment tested only very small separations between the ambiguous word and the following disambiguating information, we can assume that the pointer to the ambiguous word has a relatively brief life span. In contrast, conditional retention effects persist across clause boundaries, so we assume that the pointer to the ambiguous word exists independently of the retained meanings (although this remains to be explored).

Following this assumption, we propose that a reader's error recovery process works as follows:

1. Upon reading an ambiguous word, the reader selects the context-appropriate meaning, deactivates but retains the unselected meanings, and retains a pointer to the source of the ambiguity.
2. If disambiguating information which conflicts with the chosen meaning follows the ambiguity, the reader attempts to reinterpret the text by first following the memory trace back to the source of the ambiguity and making a new choice of meaning based on the additional contextual information.
3. If the memory trace no longer exists, either because of a clause boundary or because the actual text is no longer available, the reader will then attempt to reinterpret the text in light of the additional context by re-evaluating the retained meanings, although the actual word which gives rise to those different meanings is no longer known. A computational model of such a mechanism, called ATLAST, has been implemented successfully using a marker passing (or spreading activation) architecture to retrieve and evaluate

the alternate meanings in parallel. This work is described in greater detail elsewhere (Eiselt, 1987; 1989).

4. The process as described so far is essentially automatic or unconscious. However, if the retained meanings are no longer available, perhaps because of time or working memory limitations, the reader will not be able to resolve the contradiction without devoting attentional cognitive resources to the problem (i.e., conscious and unselective rereading of the text).

We are currently in the process of revising ATLAST along these lines, in order to test the plausibility of our theory of lexical error recovery. ATLAST's parallel marker passing architecture is well suited to this task, making the revisions fairly straightforward. Adapting a more traditional serial language processing architecture to this task would be far more problematic.³ In addition, we are designing an experiment to test the validity of this theory. Briefly, this experiment will involve using Carpenter and Daneman's target texts in the forced-choice paradigm we employed in our experiment described above. In this new paradigm, for example, the reader of a text like Text 2 would be unable to look back at the ambiguous word "bass" after reading "guitarists." If our theory is correct, the reader should still be able to arrive at a new correct interpretation of the text, even though the option of reprocessing the text has been eliminated, because of conditional retention.

Conclusion

There has been very little study of lexical error recovery, and the two theories of lexical error recovery herein at first glance seem to be contradictory rather than complementary. We have demonstrated in this paper, however, that it is possible to combine these two theories into a simpler unified theory of lexical error recovery. This unified theory of lexical error recovery is consistent with a highly parallel model of human language understanding.

Acknowledgments

This research was supported in part by the National Science Foundation under grant IST-85-12419.

References

- Burgess, C., & Simpson, G.B. 1988. Neuropsychology of lexical ambiguity resolution: The contribution of divided visual field studies. In S.L. Small, G.W. Cottrell, & M.K. Tanenhaus (Eds.), *Lexical ambiguity resolution: Perspectives from psycholinguistics, neuropsychology, and artificial intelligence* (pp. 411-430). San Mateo, CA: Morgan Kaufmann.
- 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.

Eiselt, K.P. 1987. A parallel-process model of on-line inference processing. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, 863-869.

Eiselt, K.P. 1989. *Inference processing and error recovery in sentence understanding* (Technical Report 89-24). Doctoral dissertation. Irvine: University of California, Department of Information and Computer Science.

Eiselt, K.P., & Granger, R.H., Jr. Process independence and concurrency in a model of sentence understanding. In U. Hahn and G. Adriaens (Eds.), *Parallel natural language processing*. Norwood, NJ: Ablex Publishing. Forthcoming.

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*, 172-178.

Holbrook, J.K. 1989. *Studies of inference retention in lexical ambiguity resolution*. Unpublished doctoral dissertation. Irvine: University of California, School of Social Sciences.

Holbrook, J.K., Eiselt, K.P., Granger, R.H., Jr., & Matthei, E.H. 1988. (Almost) never letting go: Inference retention during text understanding. In S.L. Small, G.W. Cottrell, & M.K. Tanenhaus (Eds.), *Lexical ambiguity resolution: Perspectives from psycholinguistics, neuropsychology, and artificial intelligence* (pp. 383-409). San Mateo, 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*, 188-192.

Jarvella, R.J. 1970. Effects of syntax on running memory span for connected discourse. *Psychonomic Science*, 19, 235-236.

Lashley, K.S. 1951. The problem of serial order in behavior. In L.A. Jeffress (Ed.), *Cerebral mechanisms in behavior: The Hixon symposium* (pp. 112-136). New York: Hafner.

Onifer, W., & Swinney, D.A. 1981. Accessing lexical ambiguities during sentence comprehension: Effects of frequency of meaning and contextual bias. *Memory and Cognition*, 9, 225-236.

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

Swinney, D.A. 1979. Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, 18, 645-660.

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.

³An in-depth comparison of serial and parallel architectures for natural language understanding is given by Eiselt and Granger (forthcoming).