

Interaction Effects between Word-Level and Text-Level Inferences: On-line Processing of Ambiguous Words in Context

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

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

1. Introduction

Ambiguous interpretations that arise during text understanding are triggered by meanings of words in context. Our recent research into on-line processes of text understanding has examined how readers choose between two equally plausible interpretations of a complete text [Granger & Holbrook, 1983; Granger, Eiselt, & Holbrook, 1983]. Other researchers have focused on how readers resolve ambiguity of individual words in context [e.g., Swinney, 1979; Tanenhaus, Leiman, & Seidenberg, 1979; Lucas, 1983]. The problem addressed in this paper focuses on the overlap between these two lines of research: in particular, how the on-line process of selecting from among ambiguous word-senses contributes to, and is itself affected by, the process of selecting from among alternative, equally plausible interpretations of the overall text.

Although initial context in a text may suggest a word-sense for an ambiguous word in context, the word's ambiguity often persists in reading. This can be illustrated by using a text in which the initial context is misleading. For example, compare the following three texts:

- [1] The CIA called in an inspector to check for bugs. Some of the secretaries had reported seeing roaches.
- [2] The management called in an inspector to check for bugs. Some of the secretaries had reported seeing roaches.
- [3] The management called in an inspector to check for bugs. They knew their rivals would stop at nothing to get their trade secrets.

The word "bugs" is ambiguous in all three texts until the second sentence, yet the first sentence of each text suggests a particular reading. In the first text, the "microphone" meaning of bug initially appears to be more appropriate than the "insect" meaning. Yet the second sentence of text [1] makes it clear that the "insect" meaning is correct. In text [2], both sentences suggest the "insect" reading. In text [3], the first sentence suggests the "insect" reading, while the second sentence forces "insect" to be dropped in favor of "microphone".

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.

Thus, even if a meaning of a word seems inappropriate initially, there must be some way to retrieve that meaning as more context becomes available. We call the process by which initially inappropriate meanings are reactivated, **cued reactivation**. There are two logically competing hypotheses as to what happens to the unselected word-senses once context has ruled them out. The first, proposed by Tanenhaus, et al. [1979], is known as **active suppression**. With active suppression, a meaning in accord with initial context is selected. The other meanings are no longer primed, that is, no longer available for quick recall. In fact, the priming effect for initially inappropriate meanings is lost much more quickly when the word is in a biasing context than when the word is not in a biasing context. Thus, a word and all its meanings would have to be re-examined if initial context was misleading. Consider again texts [1] and [2]:

- [1] The CIA called in an inspector to check for bugs. Some of the secretaries had reported seeing roaches.
- [2] The management called in an inspector to check for bugs. Some of the secretaries had reported seeing roaches.

If active suppression is correct, then with text [1], the context of "CIA" will choose "microphone" as the correct word-sense of "bugs". Once "microphone" is chosen, other meanings, including "insect", will be suppressed. After the word "roaches" is read, the reader will not be able to reconcile "microphone" with "roaches", so all meanings will have to be recalled, and a meaning chosen on the basis of the further context. If instead, the reader saw text [2], the "insect" meaning of "bugs" would be selected initially. Therefore, when the word "roaches" is read, there would be no conflict in meaning, and the several word-senses of bugs would not have to be re-activated. It should be possible to measure the difference in the processes needed to understand these two texts as a difference in reading times.

We propose a second theory, which we call **conditional retention**. With conditional retention, the initially inappropriate meanings of a word are not actively suppressed if further text is available. In other words, all meanings retain their priming when it is possible that further context could cause a re-interpretation of earlier context and therefore earlier word-sense choices. Thus, for example, when text [1] is read, the "microphone" meaning of "bugs" would be selected as appropriate to context. If the text ended after the first sentence, other meanings would be suppressed. However, if the text continued after the first sentence, the other meanings, including the "insect" meaning, would still be primed. Thus, when the word "roaches" is read, the connection between "roaches" and "insect" would be available, and there would be no need to re-process "bugs". Processing text [2] would be no different: the "insect" meaning would be initially selected, and the match between "insect" and "roaches" discovered in the same way. Therefore, there should be no significant difference in reading times if conditional retention is correct.

This paper describes two experiments we ran which were designed to decide between the active suppression and the conditional retention theories. The results from these experiments are presented as evidence for a new theory of on-line word-sense disambiguation during text processing. This new theory is one part of a new model of how all inference decisions are made during text comprehension. The theory we present here incorporates experimental results of other researchers [Swinney, 1979; Tanenhaus, et al., 1979; Lucas, 1983] with our own to create a more comprehensive theory of the interaction between word-level and text-level inference phenomena than has previously been possible.

2. Background

In building an earlier model of inference behavior [Granger, Eiselt, & Holbrook, 1983], we, as did other researchers, made implicit assumptions by pre-parsing the input to our model. We essentially started in the middle of the whole inference process. We have found that many of the inferences which could be made about a single statement were often triggered by a single word, and since most words are ambiguous to some degree, our model now had to include a method of resolving lexical ambiguity in a way consistent with our results from experiments on human subjects. As a result, we have incorporated some ideas about resolving lexical ambiguity which are suggested by several recent experiments in lexical access, the low-level processes that retrieve and use words and word-senses during reading [Swinney, 1979; Warren, 1977; Tanenhaus, et al., 1979; Lucas, 1983].

2.1. The Lexical Access Findings

Lexical access involves the translation of a word's phonological or orthographic code into its underlying meaning. Of course, many words are ambiguous to some degree, so some process is needed to disambiguate these words, selecting the most appropriate meaning for the context. Two such processes are possible: either context suggests the correct meaning before an ambiguous word is processed, so that the inappropriate meaning is never accessed, or all meanings of an ambiguous word are accessed initially, and context is subsequently consulted to determine the most appropriate meaning.

To decide between these two theories, experiments would utilize the effects of word-sense priming. Essentially, what has been found with priming is that when an ambiguous word is presented to a subject, words which are related in meaning to any of the ambiguous word's senses are more quickly recognized than words which are unrelated to the ambiguous word's word-senses. However, when there is no context present, one meaning of the ambiguous word is chosen as a default, and the other meanings which are not chosen lose their priming; that is, they are no longer recognized more quickly than words which are not related to any of the meanings. If the first theory is correct, so that the context narrows the search down to only the meanings which are in accord with it, and the inappropriate meanings are thus never accessed, then the inappropriate meanings would not be primed when the word is read. However, if context is consulted only after all meanings have been accessed, then all meanings should be primed when the word is read. Once a meaning is chosen, the other meanings would presumably lose their priming in the same way that non-default meanings of words presented with no context lose their priming.

Swinney [1979] and Tanenhaus, et al. [1979] found that all meanings of a word are initially accessed, and context is then consulted to determine which word-sense is most appropriate. Lucas [1983] extended their results in two ways. First, she showed that before an ambiguous word is seen, context itself primes the appropriate word-sense. Second, after the word is seen, all word senses are accessed regardless of context. At 100 msec. after the ambiguous word is shown, the default meaning is "more" primed (slightly more quickly recognized) than non-default meanings, whether the default is appropriate to context or not. Within the next 100 msec., though, the context selects the most appropriate meaning; thus, context becomes fully active again as the disambiguation process continues. At 200 msec. after the word is seen, only the context-appropriate meaning of the word is still primed.

When an ambiguous word is presented with context, the inappropriate meanings remain active

for between 100 and 200 msec., as noted above. However, Warren [1977] indicates that when an ambiguous word is presented without context, all meanings are available for much longer than 200 msec. Therefore, Tanenhaus, et al. [1979] have suggested that when context is available for disambiguation, inappropriate meanings of a word are actively suppressed; in other words, disambiguation involves not only the identification of the correct meaning, but the immediate erasure of primed but inappropriate meanings. Tanenhaus, et al. did not test this theory experimentally.

3. The Dilemma

As was shown with texts [1-3] above, initial context may suggest a misleading interpretation which is corrected by the later text. While active suppression explains the difference between the priming of meanings when a word is not in context versus when the word is in the context of a sentence, it may not explain what happens when there are two conflicting contexts, as in the texts above. It may be that when a word is presented with more context than simply the sentence in which it appears, active suppression does not occur. Readers may instead keep both meanings of the word primed, while waiting for confirming context. In contrast to the active suppression hypothesis, we will call this the conditional retention hypothesis; with this hypothesis, although initial context selects a meaning, the other meanings will be retained on the condition that there is further text which might suggest a second meaning, and suppressed on the condition that no more context is available.

Conditional retention is a possible solution to the questions raised when a theory based on active suppression is adopted: if the meaning of the word which is selected initially turns out to be incorrect, as in texts [1] and [3] above, how will the other meaning be recovered? Must both meanings be re-primed, and the second meaning be selected instead? Or will the second meaning not be recoverable, and the text be difficult to understand? Yet, if we discard active suppression altogether, how can we explain why inappropriate meanings are no longer primed when there is context, and are equally primed with no context? With the conditional retention hypothesis, these problems are solved.

4. The Experiments

Two experiments were conducted to test between the active suppression and the conditional retention hypotheses of cued reactivation. The first experiment had subjects read texts which consisted of sentence pairs. The first sentence had an ambiguous word in it, with the text preceding the word either biased toward one of the meanings of the word or biased toward neither meaning. The second sentence of the texts either did not bias toward either word-sense, continued biasing toward the same meaning as the first sentence, or biased toward the other meaning. Thus, the subjects saw four types of text with two sentences in each: (1) no biasing context, (2) bias follows the ambiguous word, (3) single bias surrounds the ambiguous word, and (4) two different biases on either side of the ambiguous word. Table 1 contains examples of each type of text.

After the subject read a text, he or she saw a string of letters and had to decide, as quickly as possible, if the string was a word. One of four strings of letters was presented after each text. One of the strings was a nonsense word. One of the strings was a word unrelated to either meaning of

No biasing context with ambiguous word (NONE):	They had someone check for bugs. It was a routine precaution.
Bias following ambiguous word (FOLLOWS):	They had someone check for bugs. The secretaries had reported seeing roaches.
Bias surrounding ambiguous word (SURROUNDS):	The CIA called in an inspector to check for bugs. The secretaries had reported seeing microphones.
Double bias around ambiguous word (DOUBLE):	The CIA called in an inspector to check for bugs. The secretaries had reported seeing roaches.

Table 1: Examples of the four types of text presented to subjects.

the word. Two other strings were words each related to one of the meanings of the ambiguous word. Table 2 contains examples of each type of letter string.

Nonsense string:	RUD
Unrelated word:	PEN
Related to meaning 1:	SPY
Related to meaning 2:	ANT

Table 2: Examples of the four string types presented after text containing the word "bugs".

If the initially inappropriate meaning is actively suppressed, then that meaning is no longer primed, and thus will have to be re-primed when further context reverses the appropriateness judgment made earlier in the text. If meanings are actively suppressed, ambiguity is essentially no longer present; therefore, the initially inappropriate meaning should need to be re-activated. Thus, a doubly-biased story should take more time to understand than an unbiased, a bias-following, or a bias-surrounding story because the latter three types do not initially bias inappropriately, and therefore do not force such a reinterpretation, whereas a doubly-biased story initially biases the understander toward one meaning which turns out to be inappropriate, forcing a reinterpretation of the ambiguous word. With active suppression, if both meanings are suggested in the text, we would expect a slower judgment of whether the string following the text was a real word than if only one meaning is suggested by the texts.

With the conditional retention hypothesis, because all meanings of a word are retained as long as further text is available, cued reactivation becomes a simple matter of selecting one meaning over the other using all available context, rather than initial context. Because there are still links left to the initially inappropriate meaning, there should be little difference in the time it takes to understand a story with both meanings of a word presented and a text with only one of the meanings presented.

The second experiment had subjects read the same texts as in the first experiment, but the subjects' task was slightly different: they had to decide as quickly as possible which of two words was most closely related to the meaning of the text. Each of the two words was related to one of the meanings of the ambiguous word in the text (see Table 2). If the active suppression theory is correct,

initially inappropriate meanings would have to be re-accessed, and initially appropriate meanings would then be suppressed. Thus, few mistakes would be made about which word was most closely related to the text. With the conditional retention hypothesis, both meanings would be retained, so that errors on the decision task should be fairly frequent.

5. Results and Conclusions

The data for the first experiment are not all collected. The results of the second experiment strongly suggest that a conditional retention hypothesis should be adopted.

sentence-pair types:	NONE	FOLLOWS	SURROUNDS	DOUBLE
percentage of errors:	0%	7%	7%	54%

Table 3: Percentage of errors made in word choice task.

When subjects were given unbiased texts which did not have an ambiguous word primed by text, they made no errors in judgment about which word was most associated with the text. When subjects were given bias-following texts with the first sentence unbiased and the second sentence biased toward one meaning of an ambiguous word, only 7% of the decisions were errors. When subjects were given bias-surrounding texts with both sentences biased toward the same meaning of the ambiguous word, only 7% of the decisions were errors. The difference between these conditions is not significant. However, when subjects were given doubly-biased texts with one sentence biased toward one meaning of an ambiguous word, and the second sentence biased toward a different meaning, 54% of the decisions were errors. The number of errors made in this condition is significant.

Although the results so far indicate that a conditional retention theory is correct, we are currently re-running our experiments. We are, among other minor changes, using a larger population to increase external validity.

These results have helped us to refine our model of inference decisions [Granger, Eiselt, & Holbrook, 1984]. We do not see these findings as applicable only to the lexical analysis level, because the ambiguous words were the basis for pragmatic inferences about the whole text. Therefore, we are incorporating our findings at both the lexical analysis level and the pragmatic inference level. We are using this model to predict the results of other experiments which we have designed.

6. Acknowledgements

The authors would like to thank Rogers Hall for his assistance. We also deeply appreciate the time and effort put into this project by Frank Kavanaugh.

7. References

Granger, R.H., & Holbrook, J.K. 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, New York, 1983.

- 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 on Artificial Intelligence*, Washington, D.C., 1983.
- Granger, R.H., Eiselt, K.P., & Holbrook, J.K. The parallel organization of lexical, syntactic, and pragmatic inference processes. *Proceedings of the First Annual Workshop on Theoretical Issues in Conceptual Information Processing*, Atlanta, Georgia, 1984.
- Lucas, M. Lexical access during sentence comprehension: Frequency and context effects. *Proceedings of the Fifth Annual Conference of the Cognitive Science Society*, Rochester, New York, 1983.
- Swinney, D.A. Lexical access during sentence comprehension. (Re)-consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, 18, 645-660, 1979.
- 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.

