A Restricted Interactive Model of Parsing

M. Anne Britt, Gareth Gabrys, and Charles A. Perfetti
Learning Research and Development Center
University of Pittsburgh
Pittsburgh, Pennsylvania 15260
britt@vms.cis.pitt.edu

Abstract

Much of the controversy surrounding the autonomy of syntax issue has focused on whether prepositional phrase (PP) attachments can be influenced by prior discourse context. We briefly review four studies that have produced results supporting either the autonomy view or the interactive view, and then describe the results of a recent series of experiments that identify the conditions under which a reader will be garden pathed when encountering a structurally ambiguous PP. The results of these experiments suggest that a reader can avoid being led down the garden path by a discourse that successfully creates referential ambiguity, but only in sentences where the verb does not require the PP in order to be grammatical. We also describe a restricted interactive parser that can account for these empirical results. The parser divides its task between two processors that use limited forms of semantic and discourse information when making attachments.

Introduction

For many years, the possibility that syntactic processing proceeds independently of a sentence's meaning and prior discourse has captivated psycholinguists. While semantics and pragmatics clearly influence a sentence's final interpretation, theorists that postulate a modular cognitive architecture (e.g., Fodor, 1983) have proposed that a reader's initial commitment to a syntactic structure occurs without influence of meaning. Frazier (1979) has proposed that the syntactic processor can accept input only from logically prior processes, such as lexical access, whose results are required for syntactic analysis but whose functioning does not depend on the outcome of syntactic assignments. The syntactic processor cannot accept input from processes whose functioning depends to any extent on the outcome of parsing, such as semantic or pragmatic analysis. For ambiguous sentences, where more than one syntactic structure can be built, Frazier (1979) has proposed that the syntactic processor is guided by preference rules that resolve the ambiguity in a consistent direction without the assistance of semantic or pragmatic information. This position has become known as the autonomy view.

Support for the autonomy view has come from research on garden path sentences. In the case of lexically ambiguous prepositional phrase (PP) sentences such as 1a and 1b, readers appear to prefer a reading consistent with the PP supplying a goal location to the action rather than one in which the PP modifies the noun phrase (NP). Several studies of people reading these and other forms of structurally ambiguous sentences have found that people appear to resolve ambiguities in the direction of a verb attachment. Observations that attachment preferences persist even in the context of a discourse that supports the non-preferred structure have provided support for the autonomy view. (Ferreira & Clifton, 1986; Frazier, 1979; Rayner, Carlson, & Frazier, 1983).

1a. He dropped the book on the chair before leaving.

1b. He dropped the book on the battle onto the chair.

Recently, however, empirical evidence has accumulated against the autonomy view. Several experiments have shown that the non-preferred reading can be induced by placing an ambiguous sentence in the appropriate discourse context (Altmann & Steedman, 1988; Britt, 1991; Britt, Perfetti, Garrod, & Rayner, 1992). While some psycholinguists have argued that this finding supports an interactive view of syntactic processing, we wish to argue that semantic information can influence syntactic decisions only in limited circumstances.

In the next sections, we look at the evidence supporting and opposing the autonomy view for ambiguous PP attachments. Our purpose is to understand the conditions under which people appear to exhibit a preference for one syntactic structure and when this preference is overridden by discourse information. We briefly present the results of an experiment that investigates this question, and describe a parsing model to explain how the syntactic processor might act according to a default preference under some conditions but make use of discourse information under others.

Parsing of ambiguous PP's in context

Four experiments have directly examined whether discourse information can influence the selection of an initial syntactic structure for ambiguous PP sentences. Our review of these studies looks at two factors: 1) whether the referential ambiguity of the NP preceding the ambiguous PP (hereafter referred to as the target NP) was successfully established by the preceding discourse, and 2) whether the verb preceding the target NP required the ambiguous PP as an obligatory argument. We believe that these two factors can influence whether or not a reader can immediately use discourse information to select an interpretation of an ambiguous PP

Referential ambiguity. Out of context, readers tend to prefer interpreting an ambiguous PP as a verb argument than as a NP modifier (however, cf Taraban & McClelland, 1990). For example, in sentence 1a readers prefer to interpret the phrase "on the chair" as a goal location for the action "dropped" rather than as a modifier of the target NP "the book." This interpretation is supported by the PP in sentence 1a but leads the reader down the garden path in sentence 1b. Referential ambiguity is manipulated by placing these sentences in a passage that establishes either a single antecedent for the target NP (e.g., The book was about an ancient bloody battle) or more than one antecedent for the target NP (e.g., One of the books was about an ancient bloody battle, and the other was about diplomacy). When the reader encounters the target NP following a passage that contains more than one antecedent, she will fail to establish a unique antecedent for the target NP and may then attempt to use the ambiguous PP as a NP modifier to resolve the ambiguity, thus overriding the preference to attach the ambiguous PP to the verb.

One reason why some experiments have found context effects and others have not may be due to the success of the referential ambiguity manipulation. Of the four experiments to have directly studied ambiguous PP attachments, two experiments confounded this manipulation with other factors. Ferreira and Clifton (1986), whose studies found no effect of context on initial attachment preference, used several passages that inadvertently eliminated referential ambiguity in the two referent passages by foregrounding one of the two possible antecedents prior to the target sentence. When the subjects encountered the target NP, they could have unambiguously assigned it to the foregrounded antecedent, thus eliminating any need to use the ambiguous PP as a NP modifier.

The other study to confound the referential ambiguity manipulation is Altmann and Steedman (1988). This study did find an influence of context on initial PP attachments, but their manipulation of referential ambiguity has been criticized for systematically mentioning the object of the Low NP reading more recently than the object of the High VP reading (see Clifton & Ferreira,1989).

If these confounding factors are eliminated, referential ambiguity does appear to influence PP attachment. Britt et al. (1992) presented subjects with sentences similar to those of Ferreira and Clifton (1986) but in the context of passages that established either one antecedent or two referentially equivalent antecedents for the target NP, controlling for recency of mention. The result was that subjects were garden pathed on sentences such as 1b in the presence of one referent context but not in the presence of two referent context. This result was observed for reading times measured using both self-paced and eye-tracking procedures. Thus, the Britt et al. experiments show that when the target NP is referentially ambiguous with respect to a prior discourse, the preference for attaching an ambiguous PP to the verb is overridden in favor of resolving the ambiguity.

Argument structure. The second factor we look at in the prior research is whether the target sentence's verb requires the PP to fill an argument in order for the sentence to be grammatical. For example, the verb "dropped" can take a PP as a goal argument, as in sentence 2a, but is also grammatical without the PP as in sentence 2c. In contrast, the verb "put" is grammatical with the PP, as in sentence 2b, but is not grammatical without it, as in sentence 2d. Thus we say that the goal is an optional argument of "dropped" but is an obligatory argument of "put."

2a. He dropped the book on the chair.

2b. He put the book on the chair.

2c. He dropped the book.

2d.*He put the book.

Of the four studies to examine ambiguous PP attachment, the two that found no immediate influence of discourse on parsing used mostly obligatory verbs, whereas the two that found override used mostly optional verbs. Ferreira and Clifton (1986), for example, used many verbs that require obligatory goal arguments. A recent study by Rayner, Garrod and Perfetti (1992) is also subject to this criticism. These authors presented subjects with ambiguous PP sentences and manipulated whether the discourse context placed the target NP in focus or out of focus. Rayner et al. found that focus had no immediate influence on the attachment decision, as measured by first pass reading times, but did appear to affect the repair process, as measured by total reading times (which includes any eye-movements back to a word). It is difficult, however, to conclude that focus had no effect because many of the verbs used in the experiment required obligatory goal arguments. If the parser must fill obligatory arguments in order for a sentence to be

grammatical then this need may take precedence over the need to reinstantiate an out-of-focus referent or resolve an ambiguous reference.

The two studies that found an immediate influence of discourse context on the assignment of an ambiguous PP used mostly verbs that take PPs as optional arguments. Altmann and Steedman (1988), for example, used verbs that always took PPs as optional instrument arguments (or as some have argued, adjuncts). The Britt et al. (1992) experiments also used verbs that take PPs as instruments as well as some verbs that take optional goal locations.

This analysis suggests that two factors, the successful manipulation of referential ambiguity and whether a verb's argument structure requires an argument in order to be grammatical, may explain why some experiments have found context effects and others have not.

Identifying the conditions of override

The potential modulating role of verb argument structure was recently tested in a series of experiments by Britt (1991). These experiments directly manipulated both the referential ambiguity of a target NP and whether the verb preceding the NP preferred an obligatory or optional goal argument. Subjects in this experiment read verb attached and NP attached sentences preceded by passages that established either a single antecedent for the target NP or two antecedents for the target NP. Additionally, these sentences were presented with a verb that possessed either an optional goal argument (e.g., dropped) or an obligatory goal argument (e.g., put). Examples of these sentences are shown as 1a to 1d.

1c. He put the book on the chair before leaving.

1d. He put the book on the battle onto the chair.

Using a word-at-a-time moving-window procedure, Britt (1991) found that the NP attached sentences took longer to read than the verb attached sentences for both optional context, but this difference was only found for obligatory and obligatory verbs when presented in single referent verbs when the sentences were presented in two referent context. These results, shown in Figure 1, suggest that referential ambiguity does influence the parser's initial assignment of an ambiguous PP, but only when allowed to do so by a verb's argument structure. When the verb requires a goal argument in order to be grammatical, the ambiguous PP is taken to fill that argument even if the target NP is referentially ambiguous.

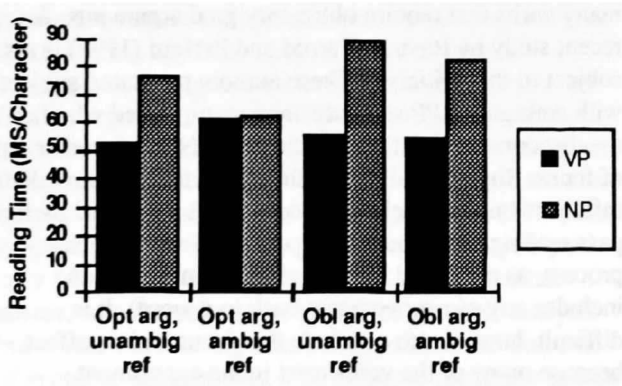

Figure 1. Results from Experiment 2 of Britt (1991).

A Restricted Interactive Model

A parser that uses only lexical class information and simple assignment principles is clearly incapable of accounting for the above empirical results. While a fully interactive parser could account for the results, we prefer to take a more conservative approach--to try to preserve as much autonomy of processing as possible while still accounting for the data. We call the result of this effort a *restricted interactive parser*. Presently, the only purpose of this parser is to explain how PP assignment decisions can be influenced by verb argument structure and referential ambiguity. Eventually, we intend to expand the parser's capability. In the remainder of this section we describe the structure of the parser, the predictions it makes for the sentences studied by Britt (1991), and how well it fits those reading time results.

The restricted interactive parser

The restricted interactive (RI) parser is based on ideas proposed by Abney (1991), Britt (1991), and Perfetti (1990) that divide parsing into two processes: *constituent building* and *argument filling*. The first process constructs small phrase units such as NPs and PPs. The second process assigns completed phrases to open arguments or attaches them as adjuncts. This approach to parsing has been described as "frame-based" as compared with models that rely on phrase structure rules.

The RI parser begins processing by instantiating a sentence frame. Parsing then proceeds bottom-up. Trigger words, such as determiners, nouns, and prepositions, cause phrases to open and new words are added to open phrases until a word is read that is incompatible with the phrase (i.e., late closure). Completed phrases are attached to open arguments or as adjuncts. The primary source of arguments is the verb's lexical structure, though noun arguments are also possible.

The constituent builder. At the parser's current level of development, the constituent builder constructs only a few forms of NPs and PPs. Simple NPs and PPs are constructed from legal combinations of prepositions, adjectives, and nouns using only lexical class information. Complex NPs, referred to here as DPs after

Abney (1991), are constructed from D NP PP sequences.

The constituent builder also handles the attachment of an ambiguous PP. After obtaining a NP, the constituent builder must decide whether to close the DP and assign it to a verb argument or delay closure so the PP can be attached as a NP modifier. The constituent builder chooses to close the DP early when: 1) the preposition can head a phrase that can fill an obligatory verb argument, or 2) the preposition can head a phrase that can fill an optional argument or adjunct and the NP is referentially unambiguous. This decision procedure creates a preference for using a PP as a verb argument or adjunct rather than as a NP modifier. In cases of referential ambiguity, however, the NP modifying function of the PP is preferred but only if the PP is not required as an obligatory argument.

The argument filler. The argument filler assigns completed phrases to arguments associated with thematic roles or attaches them as adjuncts. We distinguish arguments from adjuncts according to whether the role is part of a thematic expectation associated with a word or whether the role is supplied only by the semantics of the PP. For example, in sentence (5) we consider the phrase "on the bed" to be an argument because the verb "threw" establishes the role of goal location for the action. In contrast, in sentence (6) we consider the phrase "on Tuesday" to be an adjunct because no lexical item creates an expectation for a temporal modifier role.

(5) Billy threw the book on the bed.

(6) Billy borrowed the book on Tuesday.

The argument filler prefers to fill verb arguments over adjuncts. Whenever a phrase meets the lexical-semantic restrictions of an open argument, the argument filler will prefer this assignment to that of constructing an adjunct relationship.

It is important to note that the argument filler's preference for arguments does not extend to noun arguments. The decision to treat a PP as a NP modifier is made by the constituent builder before reaching the argument filler. One consequence of this distinction is to make the parser consistent with the observations of Clifton, Speer and Abney (1991) who found that people show an immediate preference for Verb arguments over Noun arguments and verb adjuncts, but no preference for noun arguments over verb adjuncts. However, contrary results have been obtained by Taraban and McClelland (1990) who found that for sentences whose verbs do not take the ambiguous PP as an argument, people prefer the NP attachment to the adjunct attachment. We believe that more empirical work is needed to resolve this issue.

The influence of lexical-semantic information

The decision to assign a PP as an argument or NP modifier requires more information about the preposition than its lexical class. One would not expect a reader to be garden pathed by sentence (7), for example, even though it has the same syntactic structure as sentence 1b.

(7) He threw the book of poetry.

In addition to lexical class, the RI parser takes into account certain semantic properties of both the preposition and the argument. For the sentences we use in our modeling effort, we assume that the parser matches prepositions to argument restrictions by such properties as whether a preposition can specify a particular type of relationship (e.g., location, topic).

The Influence of Discourse Information

The decision to assign a PP as a NP modifier occurs only under conditions of NP ambiguity and only when permitted by verb argument structure. The referential status of a pronoun or definite NP is obtained from an anaphoric search back through the text model. Three outcomes of this search are possible: no referent, a unique referent, or an ambiguous referent. In the case where no referent is found, an object is created in the text model, thus making no referent equivalent to finding a unique referent. In the ambiguous referent case, the parser postpones closing the DP if the next word can be part of a phrase that could resolve the ambiguity. The exception to this rule is if the phrase can also fill an obligatory verb argument.

Repair

Our modeling effort is intended to explain uninterrupted reading, that is, occasions when a reader is not led down the garden path. However, in order to obtain testable predictions for the sentences we are studying, we developed a simple repair process. When the parser detects a mismatch between a phrase and an argument, it tries to recover by: 1) finding the last phrase it assigned, 2) dissassigning the phrase, and 3) returning to the current context. The effect of the repair operation is to backtrack from the last decision and see what happens.

"Simulating" the parser

We are currently in the process of realizing the model as a computer program. The decisions made by the constituent builder and argument filler have been expressed as production rules. These rules test for patterns of lexical items stored in working memory, matches between the lexical features of words and

NP Attached Sentence in One Referent Context				NP Attached Sentence in Two Referent Context			
He	1:LexAccess 2:Open DP Open NP 3:Attach N-NP	**the**	1:LexAccess 2:Open DP 3:Attach D-DP	**He**	1:LexAccess 2:Open DP Open NP 3:Attach N-NP	**the**	1:LexAccess 2:Open DP 3:Attach D-DP
threw	1:LexAccess 2:Close NP Open VP 3:Attach NP-DP Close DP Attach V-VP 4:Attach DP-S	**battle**	1:LexAccess 2:Open NP 3:Attach N-NP	**threw**	1:LexAccess 2:Close NP Open VP 3:Attach NP-DP Close DP Attach V-VP 4:Attach DP-S	**battle**	1:LexAccess 2:Open NP 3:Attach N-NP
the	1:LexAccess 2:Open DP 3:Attach D-DP	**onto**	**1:LexAccess** **2:Close NP** **Open PP** **3:Attach NP-DP** **Attach P-PP** **4:Close DP** **5:Attach DP-PP** **Close PP** **6:Attach PP-VP** **7:Attach Fails** **8:Disassign DP** **9:ReOpen DP** **10:Attach PP-DP** **11:Close DP** **12:Attach DP-VP**	**the**	1:LexAccess 2:Open DP 3:Attach D-DP	**onto**	**1:LexAccess** **2:Close NP** **Open PP** **3:Attach NP-DP** **Attach P-PP** **4:Close DP** **5:Attach DP-PP** **Close PP** **6:Attach PP-DP** **7:Close DP** **8:Attach DP-VP**
book	1:LexAccess 2:Open NP 3:Attach N-NP			**book**	1:LexAccess 2:Open NP 3:Attach N-NP		
on	**1:LexAccess** **2:Close NP** **Open PP** **3:Attach NP-DP** **Attach P-PP** **4:Close DP** **5:Attach DP-VP**			**on**	**1:LexAccess** **2:Close NP** **Open PP** **3:Attach NP-DP** **Attach P-PP**		

Table 1. Operations of restricted interactive parser on two sentences from Britt (1991).

arguments, and NP ambiguity. For the present, we describe the results of hand simulations.

The parser has three basic operations: Open, Close, and Assign. Two additional operations, DissAssign and ReOpen are used in repair. Some actions can take place in parallel, others must occur serially. Table 1 shows the operations that take place on each word of the Britt (1991) target sentences. The actions of the constituent builder are represented by the "Open" and "Close" operations; the actions of the argument filler are represented by the "Assign" operations. Repair operations are signaled by a failed attempt to assign a phrase based on a mismatch between the features of a constituent and an argument. Operations that occur in parallel are part of the same numbered segment in Table 1.

Predictions. The model predicts two small but interesting differences in the amount of processing performed on different words. For sentences where DP closure and assignment are delayed because of referential ambiguity (optional VP and NP attached sentences in two referent context), the model predicts that, relative to the other conditions, the first preposition "on" should take slightly less time to read and that the second preposition "after/onto" should take slightly longer to read. This is because the close and assign operations are postponed at the first preposition but must occur at the second preposition. Interestingly, the model predicts that the reader will be garden pathed on optional VP attached sentences in 2 referent context, but that repairing this misparse does not produce longer reading times because no missassignments were made that need to be corrected.

Testing the Fit of the Model

Reading time predictions for each word were obtained by counting the number of cycles that result from applying the model to the sentences from Britt (1991). The number of operations was then entered into a regression equation predicting mean reading times for each word adjusted for length and overall condition differences. Because our goal was to explain the assignment of PPs, the analysis included only words from the verb to the second preposition. Figure 2 shows actual and predicted values for the 8 sentence conditions. The outcome of the regression analysis was that the overall fit of the model was quite good ($R^2 = .66$). The coefficient for number of operations was 22 msec.

With respect to the predictions described earlier, the error variance for subject reading times was too large to test for small trends in the data. The predicted curves provide a reasonable match to certain trends, especially those of the second preposition where the model accurately predicts the direction and size of the garden path effects. However, the predicted faster reading time on the first preposition for optional sentences in two referent context did not appear in the actual times. Accurate testing of these predictions must await empirical work that specifically examines this hypothesis.

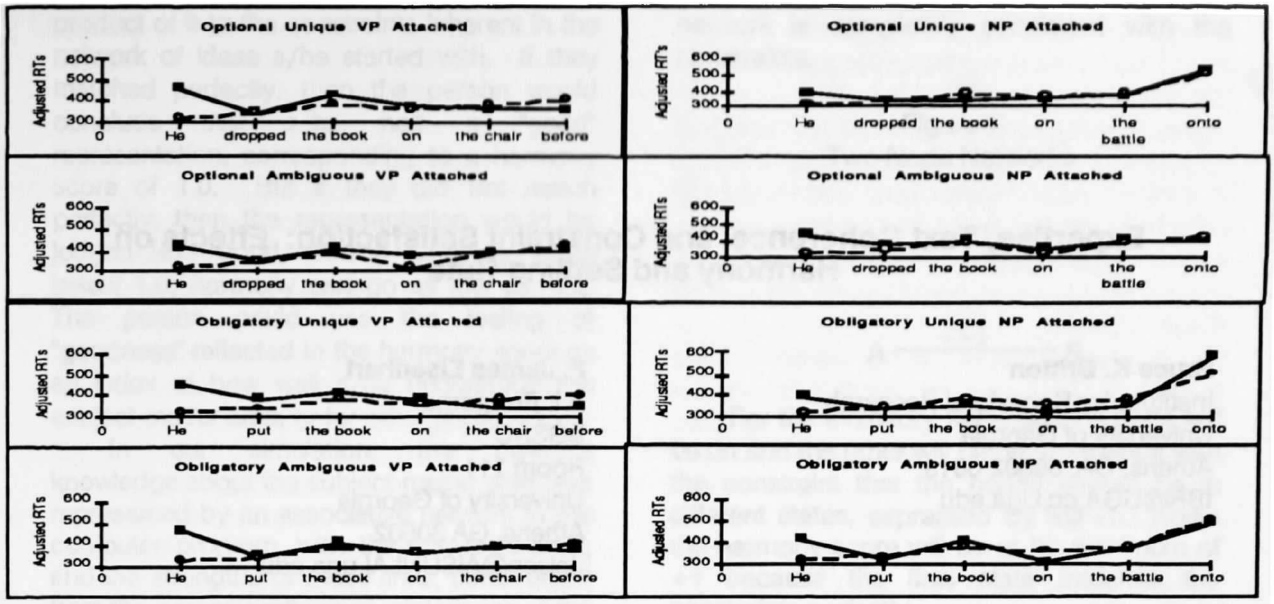

Figure 2. Mean adjusted reading time (solid lines) from Britt (1991). Dotted lines are predicted by the model.

Conclusions

We have argued that two factors, the successful creation of referential ambiguity and verb argument structure, control when override of a verb attachment preference occurs. Britt (1991) has supplied empirical support for this hypothesis, and in this paper we offer a theoretical explanation for these results. Our explanation, the restricted interactive model, acknowledges that not all syntactic processing is modular and that room must be made for the influence of discourse factors such as referential ambiguity. At the same time we reemphasize the overall control of strictly lexical and syntactic information. Discourse information can only exert an influence over syntactic processing when no overriding syntactic need exists.

References

Abney, S.P. (1991). Parsing by chunks. In R.C. Berwick, S.P. Abney, & C. Tenny (Eds.) *Principle-based parsing: Computation and psycholinguistics* (pp. 257-278). Boston, Kluwer Academic Publishers.

Altmann, G., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition, 30*, 198-238.

Britt, A. (1991). *The role of referential uniqueness and argument structure in parsing prepositional phrases.* University of Pittsburgh Doctoral Dissertation, Pittsburgh, PA.

Britt, A., Perfetti, C.A., Garrod, S. & Rayner, K., (1992). Parsing in discourse: Context effects and their limits. *Journal of Memory and Language, 31*, 293-314.

Clifton, C.,JR. & Ferreira, L. (1989). Ambiguity in context. *Language and Cognitive Processes, 4*, 77-104.

Clifton, C.,JR., Speer, S., & Abney, S.P. (1991). Parsing arguments: Phrase structure and argument structure as determinants of initial parsing decisions. *Journal of Memory and Language, 30*, 251-271.

Ferreira, F., & Clifton, C. (1986). The independence of syntactic processing. *Journal of Memory and Language, 25*, 347-368.

Fodor, J.A. (1983). *Modularity of mind.* Cambridge, MA.: M.I.T. Press.

Frazier, L. (1979). *On comprehending sentences: Syntactic parsing strategies.* University of Connecticut doctoral dissertation. Bloomington: Indiana University Linguistics Club.

Perfetti, C. A. (1990). The cooperative language processors: Semantic influences in an autonomous syntax. In D.A. Balota, G.B. Flores d'Arcais, & K. Rayner (Eds.), *Comprehension processes in reading* (pp. 205 - 230). Hillsdale, NJ: Lawrence Erlbaum.

Rayner, K., Garrod, S., & Perfetti, C.A. (in press). Discourse influences during parsing are delayed. *Cognition.*

Taraban, R., & McClelland, J.L. (1990). Parsing and comprehension: A multiple-constraint view. In D.A. Balota, G.B., Flores d'Arcais, & K. Rayner (Eds.) *Comprehension processes in reading* (pp. 231-263). Hillsdale, NJ: Lawrence Erlbaum.