11 PARSING AND COMPREHENSION: A MULTIPLE-CONSTRAINT VIEW

Roman Taraban and James L. McClelland

Carnegie-Mellon University

INTRODUCTION

A theory of parsing must explain how sentences are processed. Specifically, it must explain how a serially presented surface string is analyzed into its underlying representation. Two central components of the underlying representation include (a) a specification of how any particular constituent is configured with other constituents, which we refer to as syntactic attachment; and (b) an assignment of semantic roles to constituents, which we refer to as thematic role assignment. This representation for a sentence generally specifies what goes with what in a sentence, as well as the semantic relations holding between one constituent and another. Constructing this representation under the limits imposed by the serial order of a sentence string constitutes a large part of what is required to read and understand a sentence.

Consider the two parse trees shown in Figure 11.1. These trees consist of the same constituents, but they differ in how the constituents attach into the tree. In the first case, *clearly* attaches to the sentence node (S), and in the second case it attaches to the verb phrase node (VP). Does this have any consequences? Certainly. The point of syntactic attachment constrains the possible thematic roles that a constituent can fill. In the example, when *clearly* is attached to S, it fills the role of a parenthetical comment of the speaker; when it is attached to VP, it modifies the verb, specifying the manner in which the writing was done (Jackendoff, 1972). Of course, constraints in the other direction may also hold. That is, the thematic role that a constituent is to fill will

determine the place where it attaches onto a parse tree. Thus, if one wanted to attach a parenthetical comment into a parse tree, it would get attached to S to provide sentential context rather than to VP to specify something about the verb; if one wanted to specify something about the action expressed by the verb, like the manner in which an action was performed or the instrument that was used, then the constituent would get attached to the VP node.

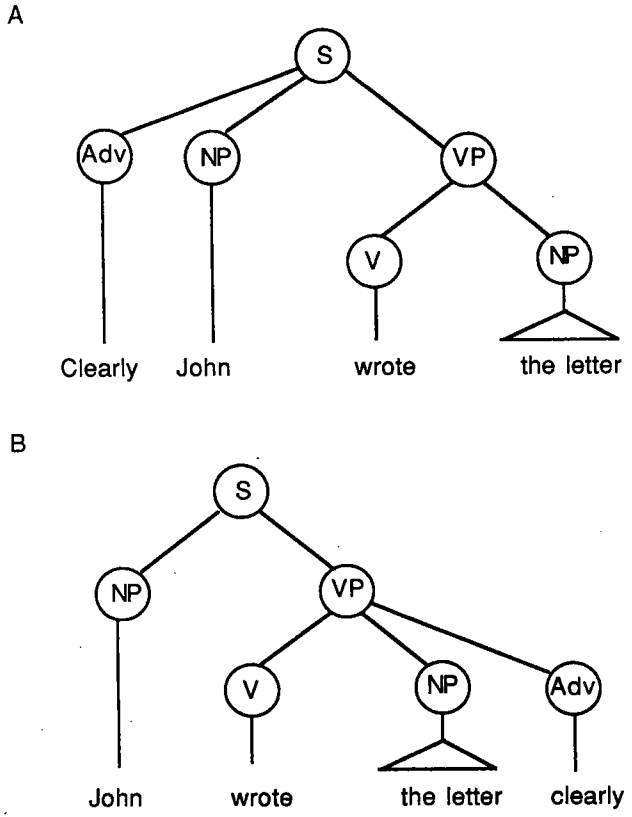

FIGURE 11.1. Examples of parse trees.

This simple example exemplifies the major theoretical concern to be tackled in this chapter, which can be summarized thus: In reading, can factors associated with thematic roles – specifically, semantic factors – constrain the initial syntactic attachment of a constituent? The reason this question is particularly important is because the initial constraints in sentence processing have often been viewed as originating in the syntax. In this chapter we question this view. We do not intend to deny the possibility that syntactic considerations may play some role. Rather, we seek evidence for the influence of semantic factors in the context of a model in which syntactic and semantic factors jointly influence the course of constructing the initial representation of a sentence.

TWO VIEWS OF PARSING

The relationship between syntax and semantics cannot be dealt with independently of a consideration of the cognitive architecture in which these potential sources of influence have their effect, that is, without a consideration of parsing models. This is because the architecture places limits on the directionality of influence within a model. For example, given the configuration shown in Figure 11.2A, Processor B could never directly influence Processor A; Processor B could only use whatever output Processor A decided to send it. If Processor B could generate a feedback signal, as shown in Figure 11.2B, then it could affect Processor A, but only after Processor A had produced some initial output and Processor B had evaluated it. Therefore, whatever contribution factors associated with Processor B might make to processing, this influence will always be with respect to the output of Processor A.

The Standard View: "Syntax Proposes and Semantics Disposes"

If we label Processor A as the *syntactic* processor and Processor B as the *semantic* processor, then Figure 11.2A and 2B schemata can be used to describe the general constraints on information flow in one class of parsers, which we will refer to as syntax-first parsers. Processor A does syntactic attachment; Processor B instantiates and evaluates the thematic roles of the attached constituents. In parsers of this class, semantic processes can filter syntactic output but cannot influence initial syntactic processing (Bresnan & Kaplan, 1982; Ford, Bresnan & Kaplan, 1982; Forster, 1979; Kaplan & Bresnan, 1982; Rayner, Carlson, & Frazier, 1983; Winograd, 1971; Woods, 1972).

The operation of the syntactic processor in syntax-first models is subject to two types of guidance, which are, by definition, in the syntax. These two types of guidance allow us to identify roughly two subclasses of syntax-first parsers.

1. *Principled.* On the one hand, the parser could be guided by general syntactic principles, that is, by a set of rules that specify the attachment of a constituent based strictly on the grammatical category of words and phrases (e.g., N V Adv NP PP VP S′) and the partial configuration of the parse tree when the principle is applied. Examples of these sorts of principles are Right Association (Kimball, 1973) and the Canonical Sentoid Strategy (Fodor, Bever, & Garrett, 1974). Frazier and colleagues (Frazier, 1978; Frazier, 1987; Frazier & Fodor, 1978; Frazier & Rayner, 1982; Rayner et al., 1983) have proposed various versions of syntax-first models that are

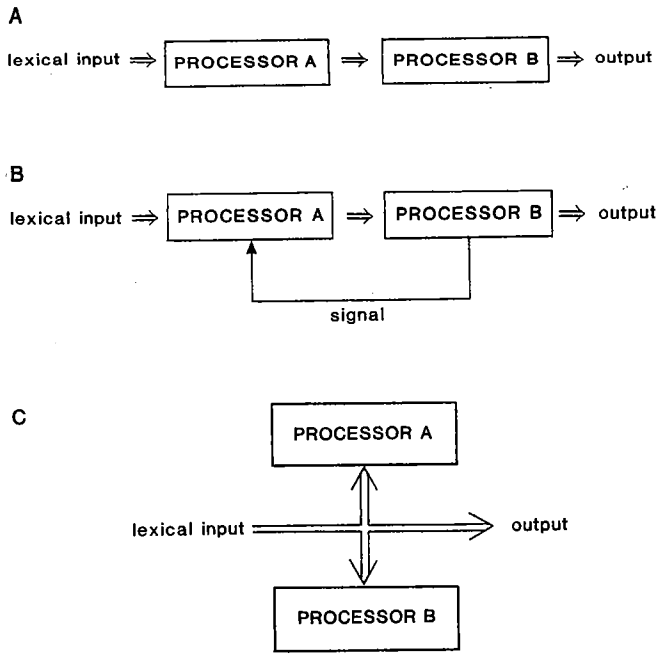

FIGURE 11.2. Syntax-first schemata (A and B) and multiple-constraint schema (C).

especially important exemplars of this subclass as the models configure phrases into a parse tree largely through guidance from principles (e.g., Minimal Attachment and Late Closure). These principles are presumably so powerful that the syntactic processor in these models initially computes just a single structurally preferred analysis of a sentence consisting of phrasal configurations specified by these general principles.

2. *Lexical.* Another subclass of syntax-first models consists of models that use lexically specified arguments and complements associated with heads of phrases (Bresnan & Kaplan, 1982; Chomsky, 1981; Ford et al., 1982; Kaplan & Bresnan, 1982; Mitchell & Holmes, 1985) – referred to here simply as subcategorization frames – in addition to syntactic principles. Ford et al. (1982), for example, use information about the preferred grammatical functions associated with verbs in order to make initial decisions about syntactic attachment. These subcategorization frames are stored in the lexicon and are activated when the verb is initially encoded. The syntactic processor in a sense looks at the verb and asks, What do I expect from this word syntactically? If the preferences that these frames represent are strong, they

influence initial syntactic analysis; if there are no strong lexically based preferences, default syntactic preferences direct parsing.

The Alternative View: Multiple Constraint Satisfaction

An alternative class of models consists of those in which construction of syntactic and conceptual representations occurs in parallel, with syntax and semantics exerting mutual influence on each other. For purposes of discussion, we need only extend the syntax-first schemata as shown in Figure 11.2c to illustrate that in this class of models syntactic principles, lexically encoded information about arguments and complements of heads of phrases – that is, subcategorization frames – and conceptual knowledge can influence parsing from start to finish. Models in this class include interactive models of the type proposed by Just & Carpenter (1980), MacWhinney (1986), Marslen-Wilson and Tyler (1980), McClelland (1987), St. John & McClelland (in press), Thibadeau, Just, and Carpenter (1982), and Tyler and Marslen-Wilson (1977). In these models, it is generally assumed that graded activation and competition allow information from multiple syntactic and semantic sources to work together or to compete with each other, with the interpretation that is most consistent with the information and most internally consistent winning out over other alternatives. Syntactic processes could dominate in some cases, but semantic information could dominate in other cases, depending on the relative strength of each source.

How is Comprehension Related to Parsing?

According to syntax-first models, syntactic attachment is initially done without considering semantic information. Of course, it is true that the syntactic relevance of semantic information has been noted in lexically based models and that these models incorporate into the syntax features like animacy and a few other selectional features. Detailed consideration of roles and their plausible fillers is, however, against the spirit of syntax-first proposals. Thus, we would not include within this class those models in which the actual semantic characteristics of a phrase or the thematic roles it might fill with respect to other constituents could influence the initial decisions made by the syntactic processor.

Multiple-constraint models include syntactic constraints on initial processing. A major difference is that they also include semantic constraints. The syntactic attachment of a constituent that is being read and the role that is assigned to this constituent are immediately subject to what that constituent and other constituents in the context are all about. For example, in the sentence fragment *Joe stirred the coffee with a . . .*, there is probably an expectation for an instrument that is like a spoon; if the fragment were *Joe stirred the paint with a . . .*, the expectation would probably still be for an instrument, but one that was somewhat different from a spoon. And if the fragment were *Joe chose the paint with a . . .*, the expectation for the prepositional phrase

would presumably be altogether different. Thus, according to this view, comprehension and parsing cannot be separated.

THE RIGORS OF READING

The analysis of a sentence in a left-to-right serial fashion is typical of reading. Time and time again, the reader faces uncertainties about attachment and assignment because information is not yet available to force one or another decision. For instance, upon reaching *clearly* in *John clearly wrote the letter*, *clearly* could be interpreted in accord with either Figure 11.1a or 11.1b. The ambiguity could be lessened by content preceding the constituent – for example, *As the handwriting expert himself has testified, John clearly wrote the letter* – or it could be resolved by content following the ambiguity – *John clearly wrote the letter and received an A for style.* In some cases, the ambiguity encountered and manner of initial resolution lead to disastrous consequences, as in the well-known garden path sentence *The horse raced past the barn fell.* Here, the late occurring final verb is a cue that *The horse* fills the thematic role of patient rather than agent, and that *raced* is not the main verb but rather a participial form, which requires major revisions that some people are unable to grasp even after considerable thought (Warner & Glass, 1987).

Processing Expectations

What does a reader do under these conditions of uncertainty? On the one hand, it appears that readers are sensitive to the alternative constructions possible when faced with ambiguities, as discussed in the research on verb-complexity (Fodor, Garrett, & Bever, 1968; Holmes & Forster, 1972; Shapiro, Zurif, & Grimshaw, 1987). Further, readers appear biased in the initial syntactic attachment of constituents, as discussed in part in the research on verb preferences (Clifton, Frazier, & Connine, 1984; Connine, Ferreira, Jones, Clifton, & Frazier, 1984; Holmes, 1984, 1987; Mitchell & Holmes, 1985). It was our view that readers were sensitive to what a sentence was about, and the content of a sentence suggested itself as a potential source of processing preferences. It was our goal to show that the ongoing process of constituent attachment and role assignment was influenced by prior content. This content set up *expectations* for further processing and provided an important source of guidance.

Our studies centered on a particular syntactic ambiguity that has been the focus of previous studies that purported to find evidence for a syntax-first model and, in particular, for the view that initial parsing decisions are guided by general syntactic principles. Our intuitions, however, suggested to us that, in this particular construction, syntactic preferences were in fact relatively weak and that attachment and role assignment seemed more susceptible to semantic guidance based on the specific content of the sentence. We chose to study this construction as

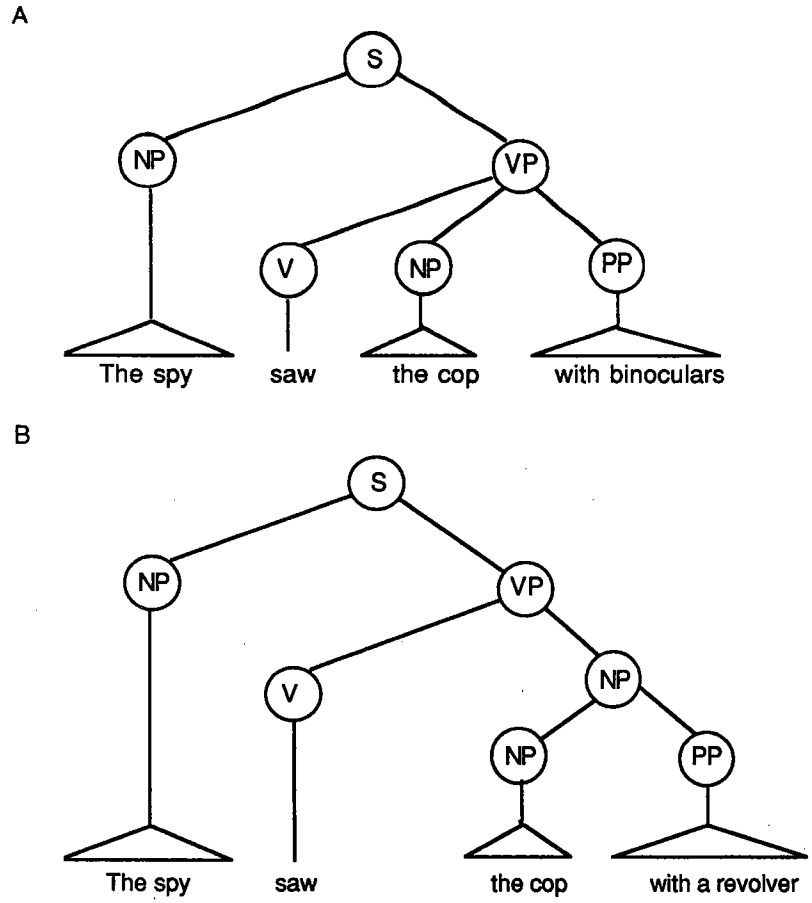

FIGURE 11.3. Examples of parse trees. A. Minimally attached prepositional phrase. B. Nonminimally attached prepositional phrase.

a test case, not because we felt that syntactic influences on processing are necessarily weak in general, but because we felt that content-based influences on initial processing would be most easily observed in the absence of strong syntactic bias.

The ambiguity in question is the ambiguity in the attachment of prepositional phrases in sentences like (1) and (2). The target constituent in both cases is the prepositional phrase (PP) that follows the object noun phrase (Object NP) in the first clause. This PP could attach either to the VP node or to the Object-NP node, as shown in Figure 11.3. As shown in Examples 1 and 2, the target phrase is followed by a conjunction signaling the beginning of a new clause.

1. The spy saw the cop with *binoculars* but the cop didn't see him.
2. The spy saw the cop with a *revolver* but the cop didn't see him.

IS PROCESSING GUIDED ONLY BY GENERAL SYNTACTIC PRINCIPLES?

Minimal Attachment

In the first experiment, we wanted to simply show that general syntactic principles were not adequate by themselves for accounting for the types of guidance people used for syntactic attachment. We therefore contrasted expectations based on general syntactic principles to those based on sentence content. The prediction that attachment and assignment of the critical prepositional phrase would be influenced by the specific content in the sentence comes directly from the view in multiple-constraint models that other sources besides general syntactic principles influence initial processing.

The syntactic principle that is relevant to guiding the attachment for phrases like *with binoculars* or *with a revolver* in Examples 1 and 2 is Frazier's minimal attachment principle (Frazier, 1978; Frazier & Fodor, 1978; Frazier & Rayner, 1982; Rayner et al., 1983), which states that initial syntactic decisions will favor the simplest attachment of a phrase into the phrasal representation of the sentence; that is, the syntactic processor favors the structure with the minimum number of nodes. Figure 11.3 makes this notion concrete. According to minimal attachment, the initial attachment of the PP *in both sentences* will be to the VP node – that is, it will function as a complement to the verb *saw*, as shown in Figure 11.3a. The reason is that such an attachment of the prepositional phrase does not, on Frazier's account, require the construction of new nodes in the syntactic tree that represents the structure of the sentence. Attachment as a constituent of a complex noun phrase – as in *the cop with a revolver* – on the other hand, does, according to Frazier, require the construction of a new node – the one that represents the complex noun phrase as a whole, under which the simple noun phrase *the cop* and the prepositional phrase are both attached, as shown in Figure 11.3b. If minimal attachment makes an implausible decision, as *with a revolver* attached to VP would be, a thematic processor can veto this decision, based on its access to likely thematic arguments of verbs and world knowledge, and could require syntactic reanalysis. This reasoning fits a syntax-first model of the Figure 11.2b type.

Experiment 1

Although results in Rayner et al. (1983) for sentences like 1 and 2 fit the predictions of minimal attachment and supported a syntax-first model, it seemed intuitively possible, and consistent with reasoning according to a multiple-constraint model, that the specific content preceding the prepositional phrase in their stimuli predisposed so-called minimal attachment. We therefore used the original pairs of matched sentences from Rayner et al. and constructed an additional set of sentence pairs for which we felt the content preceding the prepositional phrase predisposed subjects towards non-minimal attachment, as illustrated in Examples 3 and 4.[1]

3. The couple admired the house with a *friend* but knew that it was overpriced.
4. The couple admired the house with a *garden* but knew that it was overpriced.

In order to quantify subjects' expectations for either minimal or nonminimal attachment, we submitted the Rayner et al. stimuli and the Taraban and McClelland stimuli to two tests of 'expectedness.' One was a cloze task in which subjects completed sentence *frames* (that part of the test sentence up to the noun in the prepositional phrase) with the first completion that came to mind; the other test was a rating task in which subjects previewed *frames* and rated prepositional phrase completions using a scale worded in terms of 'expectations' (these completions were used in a reading task with another group of subjects). The results from both of these tests of the stimuli clearly showed that subjects' expectations for Rayner et al. frames were for minimal attachment of the prepositional phrases, whereas expectations for the Taraban and McClelland frames were for nonminimal attachment of the prepositional phrases.

We then collected word-by-word reading times for the sentences, using a self-paced task in which subjects answered a comprehension question after each sentence that they read. Our main goal was to determine the amount of guidance provided by the minimal attachment principle and the amount contributed by the specific content of the sentence, which, as noted, are two distinct sources of influence. The results for the Rayner et al. stimuli are shown in Figure 11.4a, indicating that minimally attached phrases had a significant total reading time advantage of 94 ms compared to matched nonminimally attached phrases, computed over the noun-filler and the three words that followed (e.g., *binoculars but the cop* vs. *revolver but the cop*). This replicated the major finding in the Rayner et al. study. The results for the Taraban and McClelland stimuli produced just the opposite effect on reading times,

[1] The complete set of stimuli for this experiment and for the next experiment, as well as a full description of the procedure and the statistical results, are provided in Taraban and McClelland (1988).

with nonminimally attached phrases showing a total significant reading time advantage of 69 ms when compared to matched minimally attached phrases, as shown in Figure 11.4b.

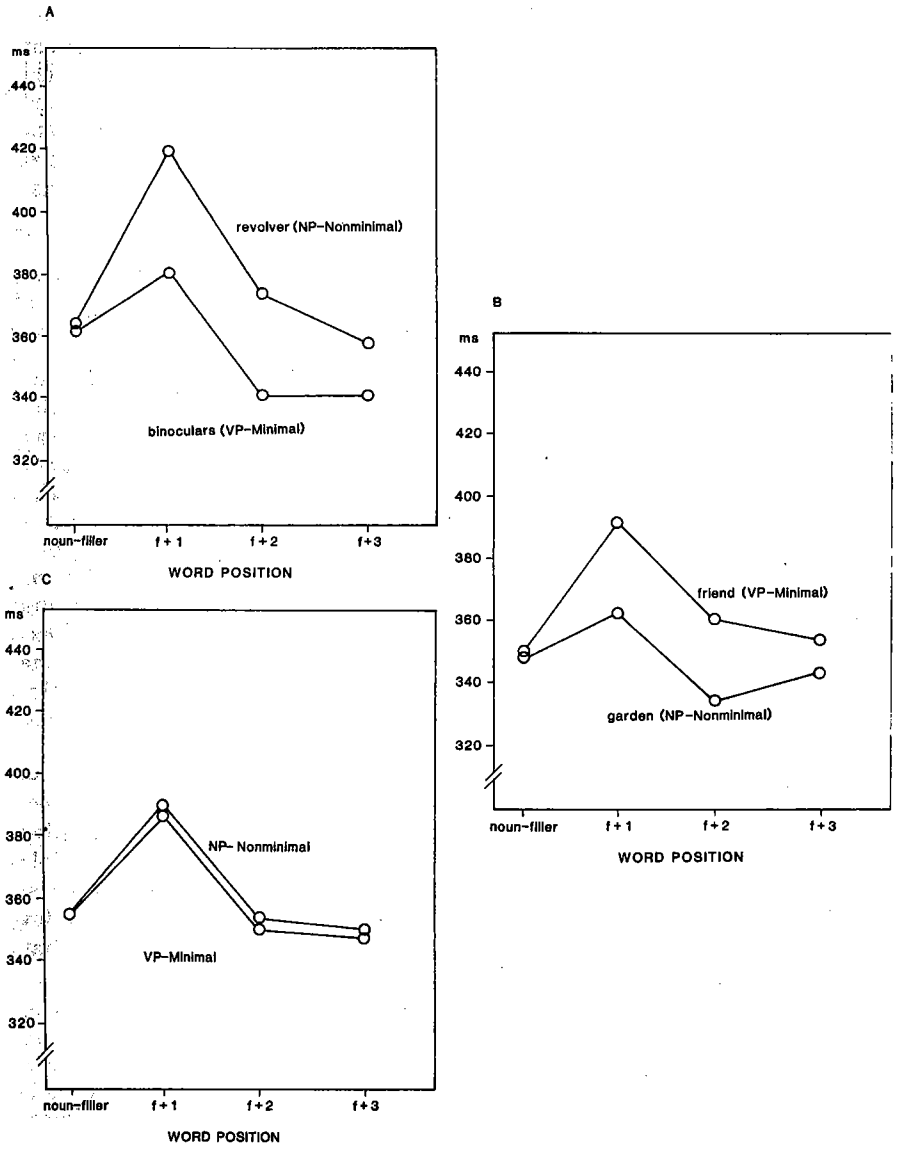

FIGURE 11.4. Experiment 1: Reading times (ms) for the noun-filler and the three words that followed for Rayner et al. stimuli (e.g., *binoculars but the cop*) (Panel A); for Taraban and McClelland stimuli (Panel B); and for minimally-attached versus nonminimally-attached stimuli (Panel C).

Content, Not General Syntactic Principles, Makes the Difference

The expectation ratings and cloze task showed that one set of sentence frames set subject's expectations to favor VP attachment of the prepositional phrase, whereas another set of sentence frames set expectations to favor Object-NP attachment. The reading times showed that these expectations were not epiphenomenal with respect to reading; rather, they produced significant differences in reading time. When attachment expectations fit the attachment required by the content of the sentence, reading times were faster compared to cases in which expectations did not fit. To determine whether there was any evidence that minimal attachment played any role in performance with these stimuli, we simply combined the results from both sentence sets and looked at the overall advantage of minimally versus nonminimally attached sentences. As shown in Figure 11.4c, the principle itself made little difference in reading times, producing an insignificant total difference of 12 ms over the four critical word positions.

Although these results do not strictly rule out the minimal attachment principle as one source of guidance, they reflect the influence of other sources besides general syntactic principles on the attachment that a subject initially expects. Specifically, subjects are influenced to expect a particular attachment for a constituent based on the specific content that precedes that constituent in the string. Thus, an important factor in the on-line processing of sentences is the degree to which the ultimate attachment of constituents in the sentence actually matches the subject's expectations for these constituents. When these expectations are violated, subjects experience difficulty relative to cases in which expectations are fulfilled. It should be noted that the particular interaction between expectations, attachment, and reading difficulty that we have reported here cannot be accounted for by *any* general syntactic principle of which we are aware – that is, by any principle that does not consider the content of the sentence – as the expectation effects occurred in sentences that differed in the content, and not the syntactic constituents, of the sentence frames. Minimal attachment may have played a small role below the level of detection possible in our design. For the remaining studies, though, we continued to focus on the role of content.

THE SEPARATE CONTRIBUTION OF THREE TYPES OF CONTENT-BASED EXPECTATIONS

The violation of content-based expectations appears to cause processing difficulty, as we found in Experiment 1. In that experiment, we tried to quantify expectations for syntactic attachment, in accord with the specific type of prediction that minimal attachment makes. But several aspects of subject's expectations may have been violated in that experiment, because presumed violations of attachment expectations covaried with possible violations of thematic role expectations for the

prepositional phrase and also with possible expectations for the noun-filler for the prepositional phrase. For example, the critical phrases in Examples 1 and 2 require different attachments, but they also receive different thematic roles (instrument of seeing in 1 and possession of the cop in 2), and these are instantiated with different noun-fillers (binoculars vs. revolver). Perhaps it was the violation of the expected thematic role of the prepositional phrase that determined processing difficulty, rather than the violation of the expected attachment per se. Or perhaps it was neither the violation of the expected attachment of the prepositional phrase nor of the expected thematic role of the prepositional phrase that produced the effects but simply a violation of a subject's expectations for a particular noun-filler. Subjects may have had a select pool of candidates, or perhaps a single candidate, in mind for the noun-filler, given the prior content of the sentence. For a verb phrase attachment, for example, and a particular role, like instrument, there are clearly better and worse instances of appropriate instruments in the context of a particular sentence. The data from Experiment 1 do not help to separate out these various sources of influence. In fact, the two conditions for the Rayner et al. and Taraban and McClelland stimuli confound all three sources. It was important to ascertain that we were indeed tapping into factors associated with syntactic attachment and role assignment, and not simply factors associated with particular noun-fillers. If processing was in fact subject to all three influences, we wanted to know their relative effects.

Experiment 2

We needed to consider, then, how to measure the effects of violations of expectations for each of these sources of influence without confounding them with the others. One way was to identify a set of sentence frames such that each one evoked a consistent expectation for an attachment and role for the prepositional phrase. We could then find four different noun-fillers for each frame, creating four different prepositional phrases. These four prepositional phrases associated with the same frame would differ according to the way in which they violated subject's expectations. The following is an example set of four sentences, with the labels used for the experimental conditions shown in parentheses:

5. The janitor cleaned the storage area with the
   a. broom (Fully Consistent)
   b. solvent (Less-Expected Filler)
   c. manager (Less-Expected Role)
   d. odor (Less-Expected Attachment)
   because of many complaints.

Phrases in the first condition are (1) consistent with subject's expectations for the attachment and role of the prepositional phrase, and (2) the particular word used for the noun-filler is actually quite good for that role and attachment. (See Taraban & McClelland, 1988, for a discussion of how thematic roles were identified.) To determine (1) we used a cloze task identical to the one in Experiment 1. The set of frames selected for the experiment showed 90% agreement between the attachment and role associated with completions in the cloze task and the attachment and role required of noun-fillers in the fully consistent condition. To determine (2) a separate group of subjects rated the expectedness of the noun-fillers, as in Experiment 1, and they additionally rated the plausibility of the fillers in the context of the sentence frames.

The sentences in the remaining conditions violated expectations for the filler; filler and role; and filler, role, and attachment, as follows. The second condition used less-expected and less-plausible noun-fillers, according to the results of the rating tasks, but fillers that were consistent with the expected attachment and role for the phrase. The sentences in the remaining conditions included the expectation violations of the previous conditions and added a new one: the thematic role of the prepositional phrases in condition (c) did not fill a role subjects were expecting; the attachment of prepositional phrases in condition (d) was not the attachment subjects were generally expecting, in addition to the roles being unexpected, as in (c). The results from the rating task for the noun-fillers showed a significant difference in expectedness and plausibility between conditions (a) and (b), and no differences between (b) and (c) or between (c) and (d), as was hoped for. As we closely matched the plausibility and expectedness of noun-fillers in conditions (b), (c), and (d), the comparison of conditions (b) and (c) in the reading task was a relatively pure indication of the processing cost of violating role expectations, and the comparison of (c) and (d) was a relatively pure indication of the processing cost of violating attachment expectations, over and above the cost of role expectation violations. A comparison of conditions (a) and (b), on the other hand, which differed in terms of the plausibility and the expectedness of the noun-fillers, indicated the effects of manipulating these factors for the noun-filler itself, while holding everything else constant.

Expectations for Thematic Roles and for Thematic Role Fillers Count for a Lot

The reading data were collected in a manner identical to Experiment 1. An examination of mean reading times by position in Figure 11.5 shows significant effects for two types of expectations. One is an expectation for the noun-filler that instantiates a thematic role. High expectedness and high plausibility for noun-fillers in the fully consistent condition and moderate expectedness and plausibility for noun-fillers in the less-expected filler condition produced a significant difference in reading time of 40 ms, summed over the noun-filler and the three words following it. That is, there was a total advantage of about 40 ms over all four words when noun-fillers were more highly expected and plausible (e.g., *broom because of many*) than when they were not (e.g., *solvent*

because of many). This effect was produced in sentences that were consistent both with respect to subject's expectations for attachment of the prepositional phrase and thematic role assignment of the prepositional phrase.

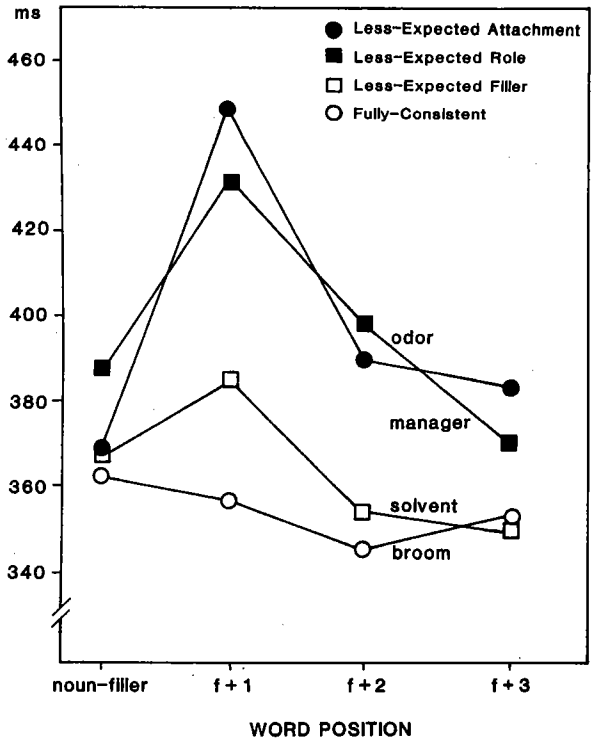

FIGURE 11.5. Experiment 2: Reading times in ms for the noun-filler and for each of the three words that followed (e.g., *broom because of many*).

The second type of expectation that was important was for the thematic role of a phrase. A violation of this expectation produced the most dramatic effect in this study. The less-expected filler condition shows reading times for sentences in which the noun-fillers disambiguated the role of the prepositional phrase in accord with subject's expectations for the thematic role of the phrase, whereas the less-expected role condition shows reading times when the role required by a sentence for the prepositional phrase was inconsistent with expectations. Attachment was held constant for sentence pairs in these conditions, and filler expectedness and plausibility were controlled. Phrases that could be assigned a role in accord with expectations (less-expected filler) produced a highly significant total advantage of about 120 ms over all four words (e.g., *solvent because of many*) compared to those phrases for which the role required by the sentence for the phrase was not in accord with role

expectations (less-expected role) (e.g., *manager because of many*). Relative to on-line reading effects in general, this violation of thematic role expectations produced a substantial slowdown in processing.

Quite surprisingly, violations of subject's expectations for the syntactic attachment of prepositional phrases produced little additional processing difficulty over and above that already produced by thematic role violations. The phrases that were neither in accord with attachment expectations nor thematic role expectations (less-expected attachment) showed a nonsignificant total additional disadvantage of about 16 ms over the critical four words (e.g., *odor because of many*) compared to the condition in which attachment was in accord with expectations but the thematic role required by the sentence for the prepositional phrase was not (less-expected role) (e.g., *manager because of many*).

Semantic Violations?

The evidence so far is consistent with the view that the guidance that content provides requires comprehension processes. This is most clearly so for the instantiation of a thematic role filler, which resulted in clear differences in difficulty. Fillers like *broom* or *solvent*, for example, cannot be evaluated for their value as instruments unless one also considers what they are instruments for and who is using them. It seems that comprehension involves just these sorts of intricate webs of weightings between agents, objects, and actions.

A similar case can be made for thematic role violations. The cloze task that we used to examine subject's expectations for thematic roles gave subjects ample time to form a conceptual representation of the frame, and it is presumably on this basis that subjects provided completions in that task. The predictive value of the cloze data is clear in the reading data. This combination of data suggests that the expectations for thematic roles in the on-line reading task, like the expectations for the fillers of those roles, were activated in the course of comprehending the sentence (cf., Schank, 1972, 1975), rather than through processes that did not require comprehension, like the activation of lexical subcategorization frames.

Finally, the time course for all the violations – noun filler, thematic role, and attachment – provides additional support for the view that the effects depend on semantically based expectations that guide the first attempts at integrating the prepositional phrase into a representation for the sentence. Figure 11.5 shows that regardless of the type of violation, a slowdown in processing occurs almost immediately to some aspect of the completion that does not fit. The type of violation determines how much additional processing is required for resolution. With more accurate tracking equipment, we could perhaps find a detectable difference in the point where the various violations are noticed by the subject. But for now, the effects seem to fall primarily on the word immediately following the word that instantiates a filler and disambiguates the attachment and role assignment of the prepositional phrase. Again,

it is easier to argue that what all these violations have in common is based on readers' attempts to make sense of a sentence as they go along – that is, on comprehension processes – rather than arising from lexical subcategorization frames.

THE STRONG PREDICTION OF MULTIPLE-CONSTRAINT MODELS

Lexical models and multiple-constraint models differ in how they use semantic information to guide initial processing. Multiple-constraint models allow for the influence of lexically encoded information on initial processing, but as we have stressed, they also allow for the influence of semantic information that is not easily encoded in a lexicon, like constraints arising from the interaction of specific participants and things described in a sentence. In contrast, in lexical models constituents are initially configured without semantic guidance. Specifically, information that is stored in subcategorization frames about likely arguments and complements for heads of phrases can be accessed and used in order to configure the syntactic attachment of the constituents in a sentence and, additionally, to label their thematic roles, without giving these constituents a semantic interpretation. This allows for the major work of sentence representation to be done in the syntax and reflects the rationale behind syntax-first models.

Chomsky (1981) provides the theoretical impetus for this view, by providing a strictly feed-forward connection between the lexicon and grammar that construct the syntactic representation of a sentence and that component that interprets the sentence. The Marcus (1980) parser is closely related to this sort of thinking in that syntactic representation is viewed as relatively foolproof without requiring close communication with interpretive mechanisms. The syntactic processor passes its output on to a case frame processor in a strictly feed forward manner, and any syntactic ambiguity that requires semantic resolution is handled with an interrupt to the syntactic processor. Finally, the lexical-functional grammar (LFG) parser of Kaplan and Bresnan (1982) fits a similar mold. Thematic role labels are accessed in the lexicon and are associated with the respective constituents in a representation of structure and grammatical function, but the role information is sufficiently independent of syntactic processing to be passed on and interpreted by a separate mechanism (Halvorsen, 1983). Indeed, processing in an LFG parser is intricately worked out without much concern for semantic/conceptual interaction (Bresnan & Kaplan, 1982; Ford et al., 1982; Kaplan & Bresnan, 1982).

As long as one limits one's view to verb-based subcategorization frames (cf. Ford et al., 1982), syntax-first guidance appears to be manageable yet powerful under some sort of priority ordering system for examining all possible expansions of the verb phrase. A multiple-constraint model predicts that guidance based on the words that appear

in a sentence is not always that simple but may depend on the influence of constituents beyond just the verb head. Therefore, in the next experiment, we sought to demonstrate the effects of another constituent – the Object NP – on expectations for attachment and role assignment in on-line reading, as before.

Experiment 3

This experiment used the same sequence of syntactic phrases as the previous experiments, and the constituent of interest again was the post-verbal prepositional phrase but with the following difference. In this experiment, we constructed sets of four sentences that held the subject noun phrase, main verb, and preposition constant, but varied the object noun phrase, as shown in the following examples:

6. The dictator viewed the masses from the
   steps (verb phrase attachment in locative role; *consistent* with subject's expectations)
   city (noun phrase attachment in source role; *inconsistent* with subject's expectations)
   but he was not very sympathetic.

7. The dictator viewed the petitions from the
   prisoners (noun phrase attachment in source role; *consistent* with subject's expectations)
   podium (verb phrase attachment in locative role; *inconsistent* with subject's expectations)
   but he was not very sympathetic.

With one object the attachment and role for the prepositional phrase that are consistent with subject's expectations are different from the attachment and role that are consistent with the other object. For example, *masses* evokes an expectation for verb phrase attachment in the role of location, whereas *petitions* evokes an expectation for noun phrase attachment in the role of source, when used in the context of the sentence frame in the example. For each object, there were two prepositional phrase noun-fillers: one that was consistent with expectations given that particular object and one that was inconsistent with those expectations.

The goal in using these stimuli was to shift expectations from one particular attachment and role assignment for the prepositional phrase to an alternative attachment and role assignment by changing the object noun phrase. If expectations actually shift with a change in this constituent, then we should predict faster reading times in an on-line reading task for all sentences in the consistent conditions compared to those in the inconsistent conditions. A stronger test of the expectation shift would require one analysis using the sets of consistent and inconsistent sentences for which subjects are expecting verb phrase attachment

(VP-expectation) and a separate analysis using the sets of consistent and inconsistent sentences for which subjects are expecting noun phrase attachment (NP-expectation). After all, the change in the object is posited to modulate expectations for attachment regardless of the site of attachment. The VP-expectation and NP-expectation sets each provide for a test of consistency over frames that use the same verb, and could show whether there is actually only one consistent cell over the four conditions associated with a verb or whether there are truly two consistent cells, as required for a clear test of the hypothesis at hand.

Eighteen stimulus quadruples like Examples 6 and 7 were used for this experiment.[2] The sequence of pretesting and reading was identical to Experiment 2.

The Object Noun Phrase Influences Expectations and Reading Times

When we compared reading times for noun-fillers and the four words that followed in sentences in the consistent condition to those in the inconsistent condition there was a significant total net advantage of 90 ms for consistent prepositional phrases. Word-by-word reading times for the overall comparison of consistent to inconsistent sentences are shown in Figure 11.6a, and separate times for objects that evoked an expectation for prepositional phrase attachment to the verb phrase and those that evoked an expectation for prepositional phrase attachment to the Noun Phrase are shown in Figures 11.6b and 11.6c.

In the overall analysis, faster reading times could not be dependent on a single specific attachment, because the consistent conditions cross attachments. In the analysis by type of attachment expectation (VP- or NP-expectation), the advantage of consistent sentences could not depend exclusively on a particular verb frame, since the verb was held constant across the VP- and NP-expectation sets. In fact, minimal attachment and verb-based lexical models would both predict null effects for Consistency, for these reasons. Faster reading times for the consistent conditions were found though and can reasonably be attributed to the modulating effect of the particular object noun phrase on the attachment and role assignment that subjects were expecting for the prepositional phrase.

In pointing out that the object noun phrase is a source of influence on attachment and role assignment, we definitely do not want to suggest that it has this influence on its own. Although the other constituents in a frame were held constant for purposes of this demonstration it is, we believe, fairly clear that the particular objects had their influence by virtue of the other constituents that appeared with them in the frame.

[2] A full description of this experiment and a list of the stimuli is in Taraban (1988), and can be obtained from the first author.

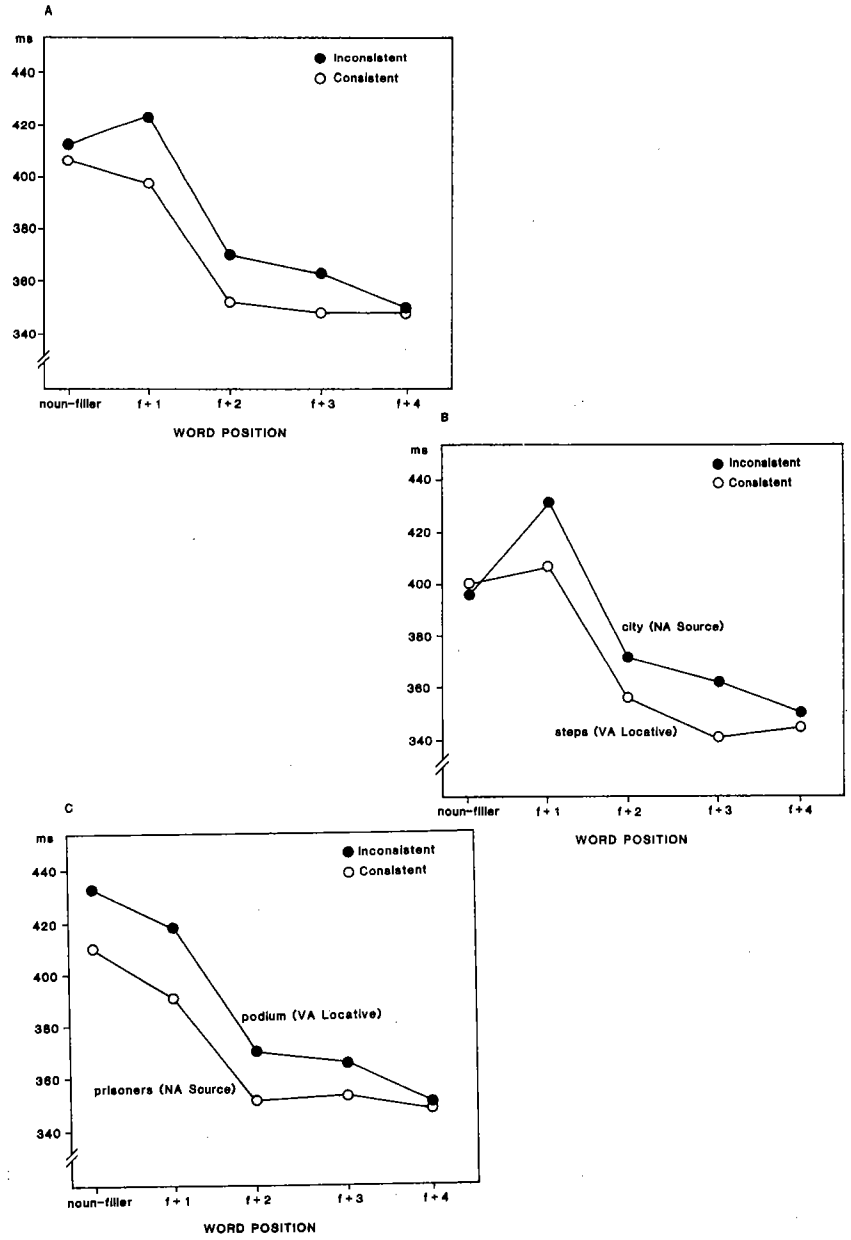

FIGURE 11.6. Experiment 3: Reading times in ms for the noun-filler and the four words that followed (e.g., *steps but he was not*) (Panel A); for sentence frames with verb phrase attachment expectations (Panel B); and for sentence frames with noun phrase attachment expectations (Panel C).

The point of this experiment was to show that the main verb in a sentence is not necessarily the sole bearer of information about the likely arrangement of constituents in a sentence, but that that information may in fact be more widely distributed. With even just a single other potential source of influence, any sort of priority ordering of promising attachments and assignments to adopt in advance of definite information, as suggested in verb-based lexical models, becomes quite unwieldy and inefficient. If it turns out that information is as widely distributed as suggested, then a priority ordering is most definitely out of the question. Therefore, not only is the verb questioned as a sole source of influence by the experimental results at hand, but the method of implementing information from that particular source, that is, through subcategorization frames, is also to be questioned.

This experiment showed the influence of just a single particular constituent – the object noun phrase – on the attachment and assignment of another constituent. It is reasonable to suppose that there are other sources. Some of these might be major constituents, like the subject noun phrase, as suggested by some off-line work by Oden (1978). Other sources could be outside the major clause in a subordinate clause (e.g., Tyler & Marslen-Wilson, 1977) or could be quite abstract, like the referential presuppositions of a phrase (Altmann, 1987; Altmann & Steedman, 1988; Crain & Steedman, 1985). Thus the finding here, although limited in scope, does fit the pattern of other findings and falls within the predictions one would make with any of a number of models that fall in the class of multiple-constraint models.

THE MAJOR ISSUES CONSIDERED

Did Subjects Use Prior Content to Guide the Initial Attachment and Role Assignment of a Constituent?

Throughout this chapter we have argued that our findings suggest that the sentence content preceding a constituent evokes expectations for the thematic role and attachment of the constituent. Is this argument justified?

"Having an Expectation" Reconsidered

The argument remains somewhat opaque and difficult to evaluate without first considering what it means to have an expectation. There are various ways of construing the capacity of an expectation to exert its influence. We consider two distinct ways, which we very loosely label here as the *active* model and the *latent* model. According to the active model, expectancies about attachment and role assignment are activated *prior* to reaching the disambiguating material that determines the fit of these expectancies with the disambiguating material.[3] When the fit is good, the pre-activated expectations and the subsequent activation of constituents composing the disambiguating material go through easily, relative to the case in which the fit is not so good. The distinguishing feature of this view is that the expectations that a person holds are activated to some degree prior to being used to process forthcoming material. In the latent model, this is not so. A person reads through the content of a sentence and only attempts to form a representation for content as it is made available. Multiple possible syntactic attachments and role assignments for constituents do not come into play until the person reaches the disambiguating material and must process it. At this point, some role assignments and attachments go through more easily than others given the particular preceding content, not because they are already active but because the preceding content affords some constructions of the constituent more readily than others.

The end result, empirically, of the active and latent models would be the same, at least in the experimental paradigm that we adopted. In both cases, processing would be slower in the region of disambiguation when the disambiguating material did not fit together very well with the way that the prior content disposed subjects to initially process it. The data that we collected do not help to distinguish between these two possibilities. In fact, it is not clear what paradigm would.

The force of the argument made in this chapter, though, does not in the end depend on which model – active or latent – most accurately describes the course of processing. The thesis is that the content prior to the disambiguating information sets people up to initially attach a constituent and assign it a role according to one set of possibilities rather than another, and all the data are consistent with this claim.

Fit of the Filler as a Possible Confound

There is a possible alternative explanation of the data that has nothing to do with expectations. As an alternative explanation would weaken the thesis, it deserves careful consideration.

Perhaps, it might be argued, we were not so much observing a violation of expected roles and attachments in the data but a difficulty in integrating the noun-filler of a prepositional phrase into the attachment and role to which it ultimately fits. To make this argument concrete, consider Sentences 8 and 9 from Experiment 1.

8. The doctor examined the patient with a stethoscope . . .
9. The doctor examined the patient with a toothache . . .

Subjects process Sentence 8 more quickly than Sentence 9, it might be argued, not so much because the prepositional phrase in 9 actually

[3] Kurtzman (1985) described a number of possible models of this sort, that vary according to the point and extent of activation, and the processes that are included, like the *abandonment* and reinstatement of hypotheses.

violates their expectations for role and attachment – they might test all roles and attachments in parallel without any preference for one over another – but because *toothache* fits less well into the role of modifier-of-the-patient than *stethoscope* fits into the role of instrument-of-the-verb.

This interpretation seems quite consistent with the results of Experiment 1, because in that experiment, it is possible that (on the average, anyway) the prepositional phrase noun-fillers that supposedly violated subjects' role assignment and attachment expectations fit less well into their ultimate role given the prior content of the sentence.[4] This being the case, it could well be that the difficulty subjects had with the unexpected versions of the prepositional phrases in that experiment could have reflected a difficulty finding a coherent interpretation of the entire sentence. Therefore, beginning with Experiment 2, we treated the *fit of a noun-filler* for a particular attachment and role as a factor distinct from the *fit of the attachment and role assignment* of a phrase; the first factor was quantified using expectation and plausibility ratings for the noun-fillers (e.g., *broom* vs. *solvent* vs. *manager* vs. *odor*) and the second factor was quantified using a cloze task for sentence frames (e.g., *The janitor cleaned the storage area with the* __).

In Experiment 2, we showed that noun-fillers that are judged to fit less well with the prior content result in longer reading times relative to fillers with a better fit. However, the remaining results of Experiments 2 and 3 are difficult to reconcile with the view that all we found were effects of this sort. For in these experiments, we found that prepositional phrases whose noun-fillers were matched for plausibility and indeed also for rated expectedness could still differ in the amount of slowing they produced, as a function of whether the attachment and role that the prepositional phrase was interpreted as filling matched the attachment and role that the subject expected based on the prior content of the sentence. This fit-to-expectation effect for attachment and role assignment was thereby shown to be separate from problems with integrating the noun-filler into its ultimate role and attachment.

Summary

In summary, the results of Experiments 2 and 3 are all consistent with the general conclusion that subjects generate expectations for the attachment and role assignment of upcoming constituents and use these expectations to guide the processing of those constituents. The differences we obtained between conditions in which the prepositional phrase attached and took a role as expected and those conditions in which it did not seem to indicate that the prior content of the sentence leads the subject to set up some mental structure that is more ready to accommodate a prepositional phrase filling one attachment and role than a prepositional phrase filling another attachment and role. When a completed prepositional phrase fits these expectations – even if it does not fit it particularly well – processing goes through much more quickly than it does when the prepositional phrase requires a different attachment and role.

[4] Kurtzman (1985) provided an analysis of the Rayner et al. (1983) stimuli that is compatible with this view. In a rating study he conducted, he found that the nonminimal completions were "more unlikely, unusual, unexpected, or implausible" (p. 213) than minimal completions.

Did Subjects Use General Syntactic Principles for the Initial Attachment of Constituents?

The major goal in Experiment 1 was to see how much of the total processing difference between matched pairs of sentences could be accounted for by the specific content composing the test sentences. But we also wanted to see whether any effect remained to be accounted for by general syntactic principles that could arrange constituents into the ultimate configuration they would hold. Principles such as these would be useful for guiding processing. The relevant principle was Frazier's minimal attachment principle. As we showed in Experiment 1, the prediction made by this principle did not hold up to the data, specifically, the principle did not appear to play a role in guiding subjects' attachment decisions. The evidence we provided, admittedly, is only relevant to the prepositional phrase attachment at hand. Evidence for minimal attachment in other cases, or evidence for other principles, needs to be evaluated on its own merits.

What is the Locus of Content-Based Guidance?

The hypothesis that we were pursuing here suggested that the locus of content-based guidance would not be limited to the influence of the main verb in a sentence. Once one entertains the possibility that other constituents besides the verb can influence attachment and assignment expectations, one is left with a multitude of additional possible sources, for example, the subject, the object, modifiers of these phrases, or even a definite *the* vs. indefinite *a* form for these phrases. We were unable to examine the influence of each potential additional source and chose rather to attempt to show the influence of the object noun phrase. Demonstrating the influence of this constituent would at least show that the verb is not the sole source of guidance. In Experiment 3, we therefore constructed pairs of sentence frames (the content up to and including the preposition) that differed in the object used but that were otherwise identical. If the object were not a source of guidance, the same attachment and role should have been expected for both frames and that attachment and role should have outperformed other possible attachments and roles in on-line reading. We showed, however, that for these particular test sentences, it was not possible to predict which particular attachment and role would be facilitated in on-line reading by simply looking at the main verb. There are findings

in the literature (Altmann, 1987; Altmann & Steedman, 1988; Crain & Steedman, 1985; Oden, 1978) indicating influences of other factors, though not all of these studies showed their effects on on-line processing.

Are Thematic Roles Assigned by the Syntax?

We know that violations of thematic role expectations produce slower processing relative to the case in which role expectations are fulfilled. This was a finding in Experiment 2. A question of major theoretical importance is whether the thematic role information had been semantically interpreted at the point when the slower processing times occurred, as an answer to this question is relevant to determining what the underlying cognitive architecture must be like.

One could place thematic role information associated with heads of phrases in the lexicon. This information as well as information about the arguments and modifiers of these lexical items could then be used to construct a structural representation of a sentence. Chomsky (1981, 1986) does so within core grammar. Kaplan and Bresnan (1982; Bresnan & Kaplan, 1982; Ford et al., 1982) use lexical entries similar to Chomsky's. In both cases, thematic role information is passed on to a semantic processor, but the thematic roles themselves can be viewed initially, at least, as part of syntax. That is, constituents are placed into a structural representation with various labels, like agent and instrument, but these labels are initially assigned without consideration for the meaning of the constituents filling these roles. According to this view, semantic interpretation is not required in order to benefit from the guidance provided by thematic role information, resulting in a model like Figure 11.2a or 11.2b.

Do the longer reading times in our data reflect a mislabeling of constituents based on a projection of role labels from individual lexical entries? It is possible, one could argue, that each lexical head could have preferred structural and thematic assignments stored with it. After a role is assigned and submitted to semantic processing, reassignment of roles may take place based on world knowledge, for example, that a *manager* is not a good instrument of cleaning in *The janitor cleaned the storage area with the manager because . . .* There are a number of problems with this view, although we will not totally reject this possibility. The implications of this view, though, are worth pondering.

In the first place, heads may have multiple sets of attachments and roles associated with them. This is evident for verbs, which often can be transitive, intransitive, or associated with sentential complements. The same could be true of noun heads in noun phrases, which could have modifiers that are likely to appear with them to one degree or another, for example, modifiers like *with a garden* for *house* or *with pepperoni* for *pizza.* Thus, the lexicon could provide information about individual items to a syntactic processor, but (a) this would most probably need to be a weighted list of possibilities, and (b) these would probably have

to be evaluated in parallel in order for the processor to eventually find a mutually consistent representation for all the constituents.

The notion of a weighted list of possible expansions for heads of phrases presents a possible paradox for a lexically based, head-driven model. This is because many heads of phrases may not *in themselves* strongly support a preference but may nevertheless contribute to a strong expectation about forthcoming constituents when taken together with other information in the sentence. For example, *married,* as head of a verb phrase, allows many possible attachments and roles, so that if one were to simply look at the distribution of possibilities for *married,* there might not be a single strong contender to rely on for guidance. However, hypotheses about the right role and attachment for a *with*-PP, in advance of the complete prepositional phrase, may be strong by virtue of all the preceding constituents. In the context of *The woman married the man with . . .* the preposition reliably evokes an expectation for syntactic attachment of the prepositional phrase to the noun phrase *the man* as a descriptor of a possession, like *money* or *a Corvette.* If the verb were *beat* and the object *the man,* we might expect *with a mop,* and if the verb were *kissed* and the object *the man,* we might expect *with affection.* As is evident, these three verbs are associated with varied attachments and roles. The verb itself may in many cases constrain the possibilities for how the sentence will be elaborated. By itself, it may only weakly constrain the possibilities; in concert with the other constituents it may strongly constrain the possibilities. Complex associations between lexical items are not typically considered to be part of a lexicon so it would probably fall to the work of the syntactic processor to sort through the weighted lists for each lexical item in order to find a representation for the constituents. But facts about women liking men with Corvettes seem to be outside the purview of a syntactic processor and an associated grammar. Therefore, a model that relies on a weighted list of options for individual entries and a pure syntactic processor that uses this information would probably fail to account for the effects of conjunctions of constraints emanating from combinations of heads of phrases.

One way out of this might be to posit complex lexical entries that captured intricate weightings between agents, actions, and objects. Then, conjunctive effects of the sort we found could be attributed to these complex lexical specifications that are semantically uninterpreted. We argue here, that although possible, this sort of position is unattractive because it is tantamount to requiring that people introduce everything that they know about the world into subcategorization frames, simply for the sake of analyzing language in the syntax. A model that is consistent with our data and that is more attractive in light of these considerations is a multiple constraint satisfaction model.

Are Prepositional Phrases Arguments?

Our answer to the question "Are thematic roles in the syntax?" is that they may be, but if so, they have a tenuous existence within the syntactic processor, which, it seems, can only call on other processors to evaluate them. One objection that could be raised against us is that prepositional phrases of the sort that we have been using are not really arguments of the verb, and other mechanisms may be necessary for processing them, but that a core grammar applies to true arguments of verbs – like a PP-locative for *put* and a PP-goal for *gave*, and their thematic roles are safely "in the syntax."

One way of responding to this claim is to invoke Occam's razor and ask why one should separate verbs according to the consistency with which they take labels. If there is a mechanism or some combination of mechanisms that can do attachment and role assignment for modifiers and complements, that is, the less-intrinsic associates of a constituent why must another device be postulated for processing the more-intrinsic associates, that is, the so-called arguments? Multiple-constraint devices, as described below, can presumably make assignments for both arguments and complements, without additionally requiring any sort of preclassification of verbs.

Processing considerations do not bode well for the idea that arguments can be assigned by the syntax without input from other sources. Consider the verb *put* for example, which is oft touted as requiring NP-patient and PP-locative arguments. Can such information be applied with no concern for what the constituents mean? Sentences like 10 and 11, which follow the sequence *put NP on-PP in-PP* illustrate that it cannot. Upon reading *put* the person can reliably expect a locative argument to follow, but (a) cannot know how many potential locatives will appear, and (b) if more than one potential argument appears, cannot be positive about how to organize them. Any single PP in Sentences 10 or 11 could fill the required slot for a locative for *put*, because all the prepositions heading these phrases are capable of filling a locative role. Combinations of PPs could also fill the required slot for *put*. Therefore, although *put* seems to represent a very predictable case, this is only true for single-PP cases. It might still be argued that *put*'s lexical representation proposes an initial parse of such sentences which must then be checked by semantic processes and potentially undone when subsequent information becomes available. Further research is required to distinguish this possibility from the one that we prefer, in which we assume the processor forms a conceptual representation as it goes along and lets this influence initial decisions about constituent attachment in all cases.

10. The cook put the roast on the table in the kitchen.
11. The cook put the roast on the table in the oven.

QUESTIONS OF MODELING

What is the Right Mechanism for Modeling Thematic Role Information?

The experiments here provide a real boost for the importance of expectations for thematic roles per se in on-line reading. In order to model this influence in a computational mechanism, it would be good to get a firm grip on exactly what a thematic role is. We have handled this question throughout using no more than a handful of broad distinctions, like instrument, that have an established place in the literature. Yet it seems that language is capable of nuancing any distinction that one would care to make in a way that could influence the way we process information that falls into these categories. There are, for example, a doctor's instruments, a musician's instruments, and an aviator's instruments, and for each we may have certain beliefs that can be influenced by the particular type of action associated with them, the agent, the circumstances, and so on. The point is simply that taxonomic thematic roles may not be fine-grained enough for most of the distinctions that we commonly make, and a role for any particular constituent could be shaded by the other constituents that it appears with. There is only suggestive evidence for this view here. The fit of noun-fillers, though, in Experiment 2 has some bearing on this issue. For example, *Hanukkah* is not such a good filler in *The choir sang the carol on Hanukkah*, but it would probably be considered a good filler in *The rabbi said the prayer on Hanukkah*. In both instances, the prepositional phrase locates the action in time, though a subcategory of these temporal phrases is obviously more appropriate for the rabbi saying a prayer than for the choir singing a carol.

The intricate nature of thematic roles is evidenced in the limited success linguists have had in devising tests for the presence of any particular role. One might expect that a role like agent would be one of the easiest to handle. One claim that has been made about agency, which is discussed in Lyons (1968), is that process verbs, like *die*, cannot have agentive subjects. This may be true in Sentence 12, but as Cruse (1973) points out, it is not true in Sentence 13. Another claim is that agents are supposed to be animate. If we compare Sentences 14 and 15, neither of which has animate subjects, we see that we could accept an agency argument for *wind* in 14 but not for *stone* in 15. This is because a *stone* would most typically be used by someone to break a window, so it is more instrumental than agentive. The wind, though, could be considered agentive, because it is using its own energy in carrying out the action (Cruse, 1973). If this is true, then elaborating 15, as in 16, now gives the stone critical features of agency.

12. Christ died.
13. Christ died in order to save us from our sins.
14. The wind opened the door.
15. The stone broke the window.

16. As a result of the explosion, a stone flew across the road and broke the window.

These examples show that it may be difficult to find a small set of features that defines a thematic role like agency. What also becomes clear is that a constituent is not necessarily endowed with a thematic role by virtue of its own features but gets part of its role from features of the constituents that it appears with. A striking example of this is found in the comparison of Sentences 17 and 18, where the critical constituent is an adverb, which is not usually considered to be of much importance in terms of *governing* other constituents (see, e.g., Chomsky, 1981). Yet, here the adverb plays a critical role in modulating the thematic role structure for the sentence. In Sentence 17, the adverb *accidentally* negates critical features of agency that John might have had. In Sentence 18, *carefully* affirms John's agency (Cruse, 1973).

17. John accidentally pushed the door open.
18. John carefully pushed the door open.

Given the empirical results from our experiments demonstrating the importance of thematic role information in on-line processing, a plausible model would have to provide a mechanism for using this information early on in the process of representing a sentence. Given the kinds of subtle effects one can glean from the example sentences, this mechanism could not limit itself to the features of individual constituents in finding a role for that constituent; rather, the role for any particular constituent would emerge both from the features of the constituent itself and from the features of other constituents that it appeared with in a sentence. The appropriateness of a thematic role filler thereby becomes dependent upon an intricate interaction of the semantic features of all the constituents in a sentence. To our knowledge, the McClelland and Kawamoto (1986) model is the best to-date for amalgamating these fine distinctions from a large multitude of input features. Knowledge of which features interact with which others and the implications of these interactions is gained through exposure to sentences that describe events in the real world. Within this model, there may be gross regularities in the co-occurrence of features, one of which might be labeled the instrumental role. A finer-grained analysis could presumably reveal clusters within these broader distinctions, corresponding say, to the typical instruments of a doctor as opposed to a policeman.

A Multiple Constraint Satisfaction Network

Some years ago, Forster (1979) postulated a model of sentence processing that consisted of a linear chain of processors, with each processor accepting input from one, and only one other processor. This was an autonomy of syntax model, and Forster considered it more promising theoretically than a model in which any processor could be influenced by any other processor. As it appears from our data that syntactic processing can be influenced by thematic roles and related conceptual

information, there is reason to suspect that syntax is not autonomous and further, that alternative models should be seriously considered. We think that models that have been proposed in recent years that belong to the class of multiple-constraint models provide an attractive alternative to those models in which syntax shuts itself off from higher-level sources of influence and is insensitive to the guidance that these sources of information afford.

St. John and McClelland (in press) provide one example of a multiple constraint model. The model consists of a pool of densely interconnected units, that is, a parallel distributed processing network. In the course of processing a corpus of sentences, the network eventually learns to assign thematic roles to constituents based on syntactic and semantic constraints, and is able to disambiguate ambiguous words, instantiate vague words, and elaborate implied roles. As each constituent of a sentence is encountered, the model updates its representation of the sentence as a whole, and if information early in a sentence indicates one interpretation that later proves to be incorrect based on subsequent input, the model adjusts its representation. In such cases, however, there is a more dramatic change in representation than in cases where subsequent input is consistent with the initial interpretation. Such changes take time, allowing the model to account for effects of violations of expectations on processing.

The various capabilities of the St. John and McClelland model correspond in some sense to the mutual influence of multiple sources of information on each other. Interestingly, however, the St. John and McClelland model does not in fact contain several processors which execute distinct types of rules, for example, inference rules, syntactic rules, and semantic rules. Information is encoded internally in the connections in an integrated network, and these connections allow the network to act as if it knew the rules. At first glance it may seem that such a mechanism is even less principled and structured than any that Forster might have envisioned and certainly rejected. However, the network is actually carrying out an extremely delicate and subtle weighing of the various sources of evidence. If they are all treated together, it is only because all of them are required for the network to solve the problem of finding (through learning) a set of factors that, when appropriately weighted, allow it to correctly interpret the sentences that it is asked to process. In fact, the learning procedure that governs the adjustment of connection strengths is implementing what we take to be an extremely central and basic principle: that multiple sources of information must be taken into account and appropriately weighted so that the model's interpretations of sentences are minimally discrepant from the interpretations that sentences actually have in the world. It is this principle of minimizing discrepancy, rather than some artificial principle of decomposability, that governs processing in the St. John and McClelland model.

There are encouraging signs both on experimental and computational fronts for multiple-constraint satisfaction networks. At this juncture, it

seems as promising to seek out more empirical evidence for multiple constraints in sentence processing and to explore the computational aspects of these data as to hold fast to the view that syntactic processing is encapsulated from other sources of influence.

CONCLUSION: A THEORY OF PARSING NEEDS A THEORY OF READING COMPREHENSION

If the evidence actually suggested that syntax-first models were adequate, then it would be possible to develop a theory of parsing that did syntactic attachment independently of other systems. The parser would be sensitive to feedback from other processors, in the case of implausible constructions, for instance, but the parser would maintain its autonomy. Reading would consist of parsing plus something else, specifically, the additional independent outputs of conceptual and discourse processors.

The picture indicated by our data suggests something different. Parsing, in the narrow sense of constructing a syntactic representation indicating the relationship of grammatical constituents to one another, appears to keep close company with conceptual processes. Syntactic attachment is not solved autonomously from the process of constructing the conceptual representation of the sentence. Our data consistently showed that the same type of constraints that govern the final fully interpreted representation govern on-line processing. The final representation is clearly based on a consideration of real-world objects and relations. The close relation between data from two independent sources – (a) the cloze data and rating data in which subjects could fully interpret the sentences and (b) the closely monitored on-line reading times suggest that subjects are drawing on the same body of information in both cases. Therefore, it appears that the computation of syntactic structure depends on semantics.

Multiple-constraint models predict that expectations that guide processing can be generated in response to characteristics of other constituents of the sentence besides the verb, for which we have found some evidence. We expect that more evidence of this nature will be forthcoming. These models also predict guidance from pre- or extrasentential context. Therefore, it should be possible to demonstrate the effects of the discourse context in which a sentence falls on the immediate expectations subjects hold for constituents that they are in the process of representing (cf. Altmann & Steedman, 1988; Perfetti, this volume). The truism that reading is for understanding makes sense in light of a multiple-constraint model, for this model suggests that the conceptual representation for a sentence and the discourse are used by the reader to structure expectations for those parts of the sentence that remain to be encountered. If it is true that parsing draws on the kinds of information and representations that reading comprehension is all about, then to understand parsing one must also understand comprehension.

ACKNOWLEDGMENTS

We would like to thank Chuck Clifton, Keith Rayner, and G. B. Flores d'Arcais for a number of valuable clarifications of parts of this chapter. The first author wishes to thank Nelson Caro for assistance with programming the experiments. This work was supported by ONR Contract N00014-82-C-0374, NSF Grant BNS 88-12048, and by a NIMH Research Scientist Development Award MH00385 to the second author.

REFERENCES

Altmann, G. (1987). Modularity and interaction in sentence processing. In J. Garfield (Ed.), *Modularity in knowledge representation and natural language understanding* (pp. 249 - 258). Cambridge, MA: MIT Press.

Altmann, G., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition, 30*, 191 - 238.

Bresnan, J., & Kaplan, R. M. (1982). Introduction: Grammars as mental representations of language. In J. Bresnan (Ed.), *The mental representation of grammatical relations* (pp. xvii - lii). Cambridge, MA: MIT Press.

Chomsky, N. (1981). *Lectures on government and binding.* Dordrecht: Foris.

Chomsky, N. (1986). *Knowledge of language.* New York: Praeger.

Clifton, C., Frazier, L., & Connine, C. (1984). Lexical expectations in sentence comprehension. *Journal of Verbal Learning and Verbal Behavior, 23*, 696 - 708.

Connine, C., Ferreira, F., Jones, C., Clifton, C., & Frazier, L. (1984). Verb frame preferences: Descriptive norms. *Journal of Psycholinguistic Research, 13*, 307 - 319.

Crain, S., & Steedman, M. (1985). On not being led up the garden path: The use of context by the psychological syntax parser. In D. R. Dowty, L. Karttunen, & A. M. Zwicky (Eds.), *Natural language parsing: Psychological, computational, and theoretical perspectives* (pp. 320 - 357). Cambridge: Cambridge University Press.

Cruse, D. A. (1973). Some thoughts on agentivity. *Journal of Linguistics, 9*, 11 - 23.

Fodor, J. A., Bever, T. G., & Garrett, M. F. (1974). *The psychology of language: An introduction to psycholinguistics and generative grammar.* New York: McGraw-Hill.

Fodor, J. A., Garrett, M. F., & Bever, T. G. (1968). Some syntactic determinants of sentential complexity, II: Verb structure. *Perception & Psychophysics, 3*, 453 - 461.

Ford, M., Bresnan, J., & Kaplan, R. M. (1982). A competence-based theory of syntactic closure. In J. Bresnan (Ed.), *The mental representation of grammatical relations* (pp. 727 - 796). Cambridge, MA: MIT Press.

Forster, K. I. (1979). Levels of processing and the structure of the language processor. In W. E. Cooper & E. C. T. Walker (Eds.), *Sentence processing: Psycholinguistic studies presented to Merrill Garrett* (pp. 27 - 85). Hillsdale, NJ: Lawrence Erlbaum Associates.

Frazier, L. (1978). *On comprehending sentences: Syntactic parsing strategies.* Unpublished doctoral dissertation, University of Connecticut, Storrs. (Also published by the Indiana University Linguistics Club, 1979).

Frazier, L. (1987). Theories of sentence processing. In J. Garfield (Ed.), *Modularity in knowledge representation and natural language understanding* (pp. 291 - 307). Cambridge, MA: MIT Press.

Frazier, L., & Fodor, J. D. (1978). The sausage machine: A new two-stage parsing model. *Cognition, 6*, 291 - 325.

Frazier, L., & Rayner, K. (1982). Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology, 14*, 178 - 210.

Halvorsen, P. (1983). Semantics for lexical-functional grammar. *Linguistic Inquiry, 14*, 567 - 615.

Holmes, V. M. (1984). Parsing strategies and discourse context. *Journal of Psycholinguistic Research, 13*, 237 - 257.

Holmes, V. M. (1987). Syntactic parsing: In search of the garden path. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp. 587 - 599). Hillsdale, NJ: Lawrence Erlbaum Associates.

Holmes, V. M., & Forster, K. I. (1972). Perceptual complexity and underlying sentence structure. *Journal of Verbal Learning and Verbal Behavior, 11*, 148 - 156.

Jackendoff, R. S. (1972). *Semantic interpretation in generative grammar.* Cambridge, MA: MIT Press.

Just, M. A., & Carpenter, P. A. (1980). A theory of reading: From eye fixations to comprehension. *Psychological Review, 87*, 329 - 354.

Kaplan, R. M., & Bresnan, J. (1982). Lexical-functional grammar: A formal system for grammatical representation. In J. Bresnan (Ed.), *The mental representation of grammatical relations* (pp. 173 - 281). Cambridge, MA: MIT Press.

Kimball, J. (1973). Seven principles of surface structure parsing in natural language. *Cognition, 2*, 15 - 47.

Kurtzman, H. (1985). *Studies in syntactic ambiguity resolution.* Unpublished doctoral dissertation, Massachusetts Institute of Technology. (Also published by the Indiana University Linguistics Club).

Lyons, J. (1968). *Introduction to theoretical linguistics.* Cambridge: Cambridge University Press.

MacWhinney, B. (1986). *Toward a psycholinguistically plausible parser.* Paper presented at ESCOL conference, University of Pittsburgh.

Marcus, M. P. (1980). *A theory of syntactic recognition for natural language.* Cambridge, MA: MIT Press.

Marslen-Wilson, W., & Tyler, L. K. (1980). The temporal structure of spoken language understanding. *Cognition, 8*, 1 - 71.

McClelland, J. L. (1987). The case for interactionism in language processing. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp. 3 - 36). London: Lawrence Erlbaum Associates.

McClelland, J., & Kawamoto, A. (1986). Mechanisms of sentence processing: Assigning roles to constituents. In D. E. Rumelhart, J. L. McClelland, & the PDP research group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 2: Psychological and biological models* (pp. 272 - 325). Cambridge, MA: MIT Press.

Mitchell, D. C., & Holmes, V. M. (1985). The role of specific information about the verb in parsing sentences with local structural ambiguity. *Journal of Memory and Language, 24*, 542 - 559.

Oden, G. C. (1978). Semantic constraints and judged preference for interpretations of ambiguous sentences. *Memory & Cognition, 6*, 26 - 37.

Rayner, K., Carlson, M., & Frazier, L. (1983). The interaction of syntax and semantics during sentence processing: Eye movements in the analysis of semantically biased sentences. *Journal of Verbal Learning and Verbal Behavior, 22*, 358 - 374.

Schank, R. C. (1972). Conceptual dependency: A theory of natural language understanding. *Cognitive Psychology, 3*, 552 - 631.

Schank, R. C. (1975). *Conceptual information processing.* Amsterdam: North-Holland.

Shapiro, L. P., Zurif, E., & Grimshaw, J. (1987). Sentence processing and the mental representation of verbs. *Cognition, 27*, 219 - 246.

St. John, M. F., & McClelland, J. L. (in press). Learning and applying contextual constraints in sentence comprehension. *Artificial Intelligence.*

Taraban, R. (1988). *Content-based expectations: One source of guidance for syntactic attachment and thematic role assignment in sentence processing.* Unpublished doctoral dissertation, Carnegie-Mellon University, Pittsburgh.

Taraban, R., & McClelland, J. L. (1988). Constituent attachment and thematic role assignment in sentence processing: Influences of content-based expectations. *Journal of Memory and Language, 27*, 597 - 632.

Thibadeau, R., Just, M. A., & Carpenter, P. A. (1982). A model of the time course and content of reading. *Cognitive Science, 6*, 157 - 203.

Tyler, L. K., & Marslen-Wilson, W. D. (1977). The on-line effects of semantic content on syntactic processing. *Journal of Verbal Learning and Verbal Behavior, 16*, 683 - 692.

Warner, J., & Glass, A. L. (1987). Context and distance-to-disambiguation effects in ambiguity resolution: Evidence from grammaticality judgments of garden path sentences. *Journal of Memory and Language, 26*, 714 - 738.

Winograd, T. (1971). *Procedures as a representation for data in a computer program for understanding natural language.* (Project MAC-TR 84). Massachusetts Institute of Technology, Cambridge, MA.

Woods, W. A. (1972). *The lunar sciences natural language information system.* (BBN Report No. 2378). Cambridge, MA: Bolt, Beranek and Newman.