JOURNAL OF MEMORY AND LANGUAGE 25, 348–368 (1986)

The Independence of Syntactic Processing

FERNANDA FERREIRA AND CHARLES CLIFTON, JR.[1]

University of Massachusetts, Amherst

Three experiments addressed the question whether semantic content or pragmatic context can direct the initial syntactic analysis assigned to sentences. Each experiment determined whether syntactic processing biases that have been observed in sentences presented in isolation can be overcome. In two experiments that measured eye movements, we found that the syntactic processing biases remained even when they resulted in thematically based anomaly or when they conflicted with discourse biases. In a third experiment, we used a self-paced reading task to replicate some of the results obtained using eye movement measures. We argue that the data support the existence of a syntactic processing module.

One approach to analyzing a complex system, such as the language processing system, is to decompose it into its separable components (cf. Simon, 1962). Fodor (1983) has elevated this methodological position to a theoretical claim. He argues that the mind is organized into two distinct types of systems: a number of input systems (modules), and a central processing system. The most important characteristic of an input module is information encapsulation: A module can only process information stated in its own representational vocabulary (Frazier, 1985), and it is insensitive to information stated in another, such as the representational vocabulary used by the central processor. Frazier and her colleagues (Frazier, 1978; Frazier & Fodor, 1978; Rayner, Carlson, & Frazier, 1983) have proposed a model of the language processor that is consistent with this modular approach. The language processor is viewed as consisting of a number of autonomously functioning components, and each component corresponds to a level of linguistic analysis (phonological, lexical, syntactic). These components are largely insensitive to nonlinguistic information sources. The present paper focuses on the claimed independence of the syntactic component from higher level, nonlinguistic knowledge sources.

If the syntactic processor (or parser) is modular, it should initially construct a syntactic representation without consulting nonsyntactic information sources, such as semantic or pragmatic information or discourse structure. The parser uses information stated in the vocabulary of syntactic rules, while semantics, pragmatics, and discourse structure reflect our knowledge of meaning and of what is plausible in the world. Information about plausibility is associated with the operation of the central processor. Notice, however, that the modular view does not imply that this higher level information is never consulted by the language processor. It is important to distinguish between *initial* and *eventual* use of nonsyntactic information. Such information could be consulted at a later stage in processing, perhaps after the parser has computed its initial representation. If, on the other hand, the language processor is nonmodular and information of any type—phonological, syntactic, knowledge of the

[1] Requests for reprints may be addressed to either author at the Department of Psychology, University of Massachusetts, Amherst, MA 01003.

0749-596X/86 $3.00

world or of one's conversational partner's beliefs, etc.—could in principle guide any decision about how to represent incoming language (Marslen-Wilson & Tyler, 1980; Tyler & Marslen-Wilson, 1977), then, initially, the most plausible analysis should generally be computed. This position will be referred to as the interactive position.

The operation of the parser can be revealed by observing how it behaves when faced with syntactic ambiguity, i.e., with more than one potential analysis of a portion of a sentence (Frazier, 1978). There is substantial evidence that the parser computes only one analysis, the first one available (Frazier, 1978; Frazier & Fodor, 1978; Frazier & Rayner, 1982; Rayner et al., 1983; cf. Berwick & Weinberg, 1984; Crain & Steedman, 1985; Marcus, 1980, for differing views). Frazier (1978) proposed that the parser follows the Minimal Attachment strategy: Attack incoming material into the phrase-marker being constructed using the fewest syntactic nodes consistent with the well-formedness rules of the language. If readers or listeners[2] see or hear the ambiguous string of words, "The editor played the tape. . . ," they will initially begin to construct a syntactic structure like that represented in Figure 1A (the simple active structure), since this is the structure with the smallest number of syntactic nodes for the string. Similarly, if the string is "Sam loaded the boxes on the cart," the reader will begin to construct a representation like that shown in Figure 1B, in which the prepositional phrase is attached to the VP, not to the NP "the boxes."

The Minimal Attachment strategy is efficient in terms of computational and memory load: Only one analysis at a time is constructed, and all incoming material is structured as it is received. However, if the syntactic analysis turns out to be incorrect and the parser is led down the garden path, then the parser will have to reanalyze the misanalyzed material. For instance, if the strings just discussed continued . . . *liked it* or . . . *onto the dock*, the reader would have to reanalyze the structure into the reduced relative clause structure shown in Figure 1C, or the structure in which the prepositional phrase is attached to the object NP, as in Figure 1D. Since it takes extra work to give up the initial simpler analysis and to construct representations with more syntactic nodes in them, the Minimal Attachment strategy predicts that people will take more time and have more difficulty processing these "Nonminimal Attachment" sentences than sentences which conform to the Minimal Attachment strategy. This prediction was supported by Frazier and Rayner (1982), who measured eye movements during reading. They found that subjects took more time to read Nonminimal Attachment sentences than Minimal Attachment sentences, and that fixation durations were longer in the region of the sentence that disambiguated the analysis than in other regions.

Rayner et al. (1983) examined the effects of within-sentence plausibility constraints on the Minimal Attachment strategy. They proposed that if the interactive model is correct, then subjects should not be garden pathed (i.e., initially compute the wrong syntactic analysis which must later be revised) in sentences semantically biased towards the Nonminimal Attachment reading, because they would always adopt the semantically preferred reading of the sentences. If, on the other hand, the modular model is correct, then garden paths would still occur in the Nonminimal Attachment sentences. Rayner et al. compared sentences like (1a) and (1b) (among other constructions). In the Minimal Attachment reading of these sentences, the prepositional phrase is attached to the verb

[2] While we studied only readers, not listeners, and while we recognize that prosodic cues may reduce some temporary ambiguities for listeners, we have no reason to doubt that the syntactic processing strategies used by listeners and readers would be the same, nor that content or discourse context would have similar effects for listeners and readers.

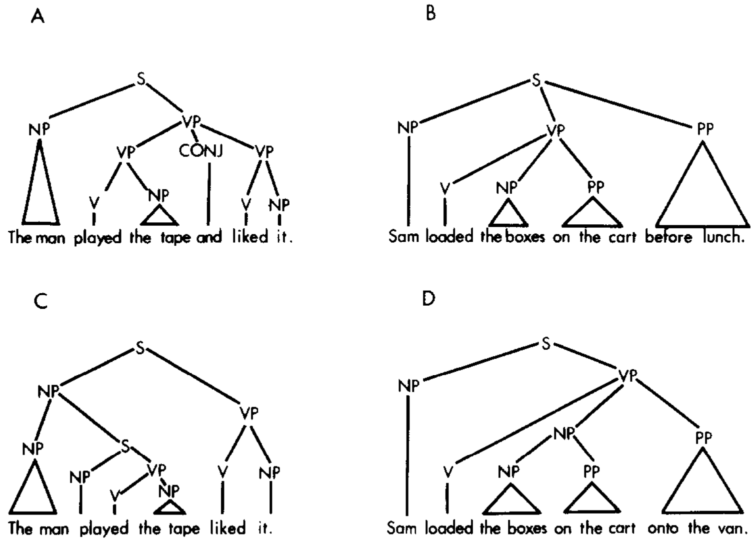

FIG. 1. Syntactic structures of minimal attachment (Panels A and B) and nonminimal attachment (Panels C and D) sentences.

phrase, indicating the instrument used in seeing (cf. Fig. 1B). This reading is plausible in (1a), but not in (1b).

(1a) The spy saw the cop with binoculars but the cop didn't see him.

(1b) The spy saw the cop with a revolver but the cop didn't see him.

Rayner et al. found that their subjects were biased by the pragmatics of the sentences so that they eventually understood "with the revolver" in (1b) as modifying the NP "the cop," while they took "with binoculars" as indicating the instrument of seeing. At the same time, Rayner et al. found evidence of garden pathing in (1b), indicating that pragmatic plausibility information did not override the initial syntactic preference for Minimal Attachment structures. However, pragmatic information *was* used eventually, showing that even a modular model must allow for some interaction between syntactic and nonsyntactic information sources. Rayner et al. (1983) proposed that another component of the language processing system, the thematic processor, examines alternative thematic structures listed for the heads of phrases and proposes plausible ones to the syntactic processor.

The Rayner et al. experiment convincingly tested and rejected the hypothesis that pragmatic information determines the first analysis of sentences. However, proponents of an interactive position may suggest alternative interpretations of these data. One potential suggestion is that only world knowledge important enough to be encoded in the grammatical processing system could guide parsing. For example, some verbs demand an animate agent, and result in anomalous or metaphoric sentences when used with an inanimate subject. A person presented with *The evidence examined by the lawyer* may interpret *the evidence* as a theme, not an agent, since *examined* requires an animate agent. *The lawyer* will then be interpreted as the agent of the verb. Consequently, the entire phrase will be interpreted as a relative clause. If thematic information could be used to guide parsing, then the difficulty of the above phrase should be eliminated or at least reduced, relative to *the defendant examined by the lawyer*. These predictions were examined in Experiment 1.

A second potential suggestion is that if the sentences were placed in appropriately biasing contexts, the effects of Minimal Attachment would not be observed. Many theories of language processing would make this prediction. Text-based theories such as that of Kintsch (Van Dijk & Kintsch, 1983; cf. Graesser and Riha, 1984; Haberlandt, 1984) emphasize the role of context in language comprehension. Such theories go beyond the noncontentious claim that context does affect interpretation to claims that context guides parsing and that readers rely on atypical and potentially disruptive strategies to encode sentences presented in isolation (Van Dijk & Kintsch, 1983, p. 32).

Crain and Steedman (1985) explicitly argue that syntactic ambiguities are resolved by semantic and discourse plausibility, rather than by syntactic strategies. They posit a syntactic processor which independently "proposes" alternative analyses, while a semantic processor "disposes" of them. They suggest that the different possible "parse routes" are constructed in parallel, so that appropriate context could in principle eliminate the difficulty associated with some syntactic structures. In particular, a given analysis could be eliminated if the context did not satisfy its presuppositions.

Complex NPs such as the relative clause *The editor played the tape* or the NP *the girl with the long hair* in *John hit the girl with the long hair* carry (according to Crain and Steedman) several presuppositions: (1) that a set of objects is in focus (e.g., editors or girls); (2) that the NP refers to one member of the set; and (3) that a single member can be identified on the basis of the information provided by the complex NP. Thus, an appropriately biasing context for these types of sentences is one which mentions a set of some objects, and allows one member of the set to be picked out. No garden paths will occur because after receiving, for instance, *the editor*, the language processor will look for further information specifying which editor is being discussed. When *played the tape* comes in, the phrase is taken as modifying the editor and allowing the right one to be picked out, rather than as describing an action taken by the editor.

An implication of these ideas is that "null" contexts (i.e., the situation in which target sentences appear in isolation), which are typically used in sentence-processing experiments to compare the relative contributions of syntactic and nonsyntactic information sources, are not really neutral. A truly neutral context must satisfy the presuppositions of the target sentence. If it does not, then there will be a purely pragmatic, not syntactic, basis for favoring one interpretation of the target sentence. A neutral context for *The editor played the tape* must specify a set of editors and provide some information that could be used to distinguish a particular editor from the others.

Crain and Steedman conducted some experiments to test these proposals by placing minimal attachment and nonminimal attachment sentences in various contexts and obtaining timed or untimed grammaticality judgments. However, since their experiments did not use an on-line measure of sentence processing, they could not distinguish between initial (or early) use of context, which is at issue here, from eventual use of context, which is not at issue. To determine whether context determines the initial analysis of ambiguous text, it is necessary to place the sentences in context and then measure processing difficulty using an on-line measure. This was done in the second and third experiments reported here.

Experiment 1

The first experiment was designed to assess the on-line operation of the parser when it has available thematic information that biases the interpretation of a syntactically ambiguous string. We used recorded eye movements to measure the time subjects took to read sentences with relative clauses. The first noun phrase was either animate or inanimate, and the relative clause was either reduced or unreduced. The sentences in (2) and (3) illustrate the four different possibilities. The slashes indicate the relevant phrases of the sentence. We will refer to the regions as c − 2 (c for critical), c − 1, c, c + 1, and c + 2.

(2a) The defendant / examined /
c − 2 c − 1
by the lawyer / turned out /
c c + 1
to be unreliable. (animate, reduced)
c + 2

(2b) The evidence / examined / *by the lawyer* / turned out / to be unreliable. (inanimate, reduced)

(3a) The defendant / that was / examined / by the lawyer / turned out / to be unreliable. (animate, unreduced)

(3b) The evidence / that was / examined / by the lawyer / turned out / to be unreliable. (inanimate, unreduced)

Sentences 2a and 2b have a temporary syntactic ambiguity that is essentially disambiguated by the *by* phrase and fully disambiguated by the verb.[3]

The sentences in (2) are Nonminimal Attachment sentences, in that the Minimal Attachment reading would take the first verb *examined* to be the main verb of the sentence. According to the garden-path theory described here (Frazier & Rayner, 1982), they should be difficult to process at the point of disambiguation (the prepositional phrase). This difficulty should not be shared by the syntactically unambiguous unreduced relative clause sentences shown in (3).

Consider the difference between (2a) and (2b). The first verb in these sentences requires an agent. The subject of (2a) is semantically a potential agent, but the subject of (2b), *evidence*, is not. If this semantic information can be used to guide the analysis of these sentences, the difficulty of (2b) should be reduced or eliminated. Specifically, subjects should be faster to read regions c and c + 1 for sentence (2b) than (2a).

On the other hand, if the parser still initially constructs a Minimal Attachment analysis, even in the face of the semantic information, then the most the semantic information can do is to hasten reanalysis, resulting perhaps in an earlier and less long-lasting disruption for (2b) compared to (2a). Specifically, subjects should take equally long to read region c and possibly c + 1 for sentences (2a) and (2b), and the reading times should be longer than for the control sentences (3a) and (3b). If the semantic information does aid reanalysis, subjects may be somewhat faster in regions c (and c + 1) for (2b) than (2a) (but still slower than for both (3) sentences). In addition, it is possible that region c − 1 will take longer to read for sentence (2b) than (2a) if subjects detect the anomalous nature of the phrase *The evidence examined*.

Method

Subjects. Sixteen students from the University of Massachusetts were paid $5 an hour to participate in the experiment. All subjects had normal uncorrected vision.

Apparatus. Subjects' eye movements were recorded via a Stanford Research Institute Dual Purkinje Eyetracker interfaced with a Hewlett–Packard 2100 computer that controlled the experiment. The eyetracker has a resolution of 10 min of arc, and the horizontal and vertical signals from the eye tracker were sampled every millisecond by the computer. Eye position was determined by comparing the signals each 4 ms with the prior 4 ms. A complete record of eye location, fixation duration, and fixation sequence was stored on a computer disk for each experimental item. The stories were presented on a Hewlett–Packard 1300A CRT, and the subject's eye was 46 cm from the CRT.

Eye movements were recorded from the

[3] While the string *The defendant examined by the lawyer* could continue *some very technical looking documents*, retaining its active (Minimal Attachment) reading, such a reading seems highly unlikely. We assume therefore that readers took the *by* phrase to signal a relative clause reading.

right eye and viewing was binocular. The brightness of the screen was adjusted to a comfortable level for each subject and adjusted whenever necessary during the experiment to maintain the subject's comfort. A black theater gel covered the CRT so that the letters would appear clear and distinct to the subjects.

The 16 sentences in this experiment were presented one at a time to subjects (randomly intermixed with 40 texts used in other experiments). The letters were presented in lower case, except for the first letter of the first word of a sentence, and the first letter of proper nouns.

Materials. Sixteen target sentences such as (2) and (3) were presented to subjects. Each subject saw four sentences in each of the four conditions. The four conditions were (1) Animate Reduced, (2) Inanimate Reduced, (3) Animate Unreduced, and (4) Inanimate Unreduced. The sentences were always presented together with a following filler sentence, which was included to prevent subjects from becoming sensitive to the target sentence. This filler was related to the first sentence so that together they formed a small text. However, this second sentence was not scored and, since it followed the target sentence, it did not affect reading of the target.

Four lists of items were constructed. In a single list, a sentence appeared in only one of the four conditions. Condition number was assigned to the items according to a Latin Square design. Each list was preceded by 2 practice items, consisting of about three sentences. Thus there were 56 items (2 practice, 16 experimental, and 40 texts used in other experiments) in each list.

Procedure. Each subject was prepared a bite bar that eliminated head movements during the experiment. The subject sat in front of the CRT and the experimenter explained the procedures and gave him or her the instructions. The eye tracking system was then calibrated (Rayner, 1978), and a calibration check was made before each text.

The subject then read the items. Following every fourth text, the experimenter asked the subject a single true/false question about some aspect of the item. The questions were included only to induce subjects to pay attention. The subject replied to the question by tapping the table in front of him with his finger, once if the statement was true, and twice if it was false. The subjects were asked to tap the table rather than to respond orally to prevent loss of calibration. However, the subject was free to come off the bite bar at any time between items.

Scoring regions. The target sentences were divided into scoring regions. All sentences were divided into (1) the first NP ($c - 2$), (2) the *that was* phrase, if it existed, (3) the first verb ($c - 1$), (4) the *by* phrase or prepositional phrase (c), (5) the main verb and any auxiliaries ($c + 1$), and (6) the rest of the sentence, if the sentence continued past $c + 1$ ($c + 2$). Only regions $c - 1$, c, and $c + 1$ were included in the data analyses.

Results

The results of this experiment will be analyzed in terms of first pass reading times and second pass reading times. Following Frazier and Rayner (1982) and Rayner et al. (1983), reading times (omitting 15% of the data lost due to track losses, blinks, and other malfunctions) were divided by number of characters (including character spaces and punctuation marks) in order to control for the effect of region length.[4] Also, first pass reading times were separated from total and second pass reading times. First pass reading times include only left-to-right fixations that were the first fix-

[4] We argue later that dividing reading time by number of characters is inappropriate for the self-paced reading task used in Experiment 3. However, the reasons given there do not apply to eye movement experiments. In such an experiment, we can expect the total time spent fixating on a region to be 0 if the length of the region is 0, and to increase in an essentially linear fashion with the number of characters in the region, making time per character an appropriate measure.

TABLE 1
MEAN FIRST PASS AND SECOND PASS READING TIMES PER CHARACTER (IN ms) EXPERIMENT 1

	First pass			Second pass		
	Region			Region		
Condition	c − 1	c	c + 1	c − 1	c	c + 1
Animate Reduced	33.3	40.4	31.9	15.3	8.2	12.8
Animate Unreduced	31.9	30.7	33.1	6.9	3.6	8.0
Inanimate Reduced	37.7	38.4	32.6	12.6	14.9	16.9
Inanimate Unreduced	30.1	30.3	28.6	0.0	0.0	2.5

ations on that character. There is one exception to this procedure: Any right-to-left movement from the first fixation on a line to a position closer to the beginning of that line was treated as occuring from left to right rather than right to left. This exception was made because such regressions are typically made to correct for overshoot of the eye after a return sweep, and most likely do not reflect aspects of cognitive processing (Rayner, 1978). Second pass reading times include only regressions and rereads of the material, and exclude all first pass reading times. This procedure of separating first pass reading time from second pass reading time allows one to isolate initial analyses of the linguistic stimuli from reanalysis. All effects are significant at the .05 level of significance unless otherwise noted.

The first pass reading times for the four different conditions in regions c − 1, c, and c + 1 are shown in Table 1. Subjects were significantly faster in the unreduced conditions, indicating that the nonminimal attachment sentences were difficult for subjects to understand. The effect of reduction was significant ($F_1(1,15) = 112.18$; $F_2(1,15) = 14.66$; $p < .01$). There was no interaction between reduction and animacy of the first noun phrase.

The Bonferroni t test was used to compare these differences in region c, the disambiguating region. For the animate NP sentences, the difference between the unreduced and the reduced versions was significant by items and subjects; for the inanimate NP sentences, the difference was significant by items and almost significant by subjects ($t = 1.68$; $t = 2.04$ needed for significance).

Considering just region c − 1, reading times were long when the verb followed the inanimate NP in the reduced condition. This difference was significant by items and almost significant by subjects ($t = 1.99$).

The second pass reading times are shown in the second panel of Table 1. Reading times appear to be longer in the reduced than the unreduced conditions. This effect was significant ($p < .05$) by both subjects ($F_1(1,30) = 4.77$) and items ($F_2(1,30) = 4.46$). Finally, the probability of a subject making a regression did not differ among the four conditions.[5]

Discussion

The nonminimal attachment sentences in this experiment were subject to the normal difficulty in the disambiguating *by* phrase region, reflecting garden pathing brought about by the Minimal Attachment preference. The presence of disambiguating syntactic information in the unreduced relative sentences eliminated this difficulty, indicating that the task was sensitive to the use of syntactic information. However, the difficulty of the reduced relative sentence

[5] Note that second pass reading times included times when a subject reread a whole passage, or a large section of it. These were not counted as regressions; regressions were limited to instances when the eyes moved from a disambiguating region to the preceding ambiguous region or an earlier point in the disambiguating region.

during the disambiguating *by* phrase persisted even when the Minimal Attachment analysis was blocked by the inanimacy of the subject NP, which made that NP unfit as the subject of the first verb. Readers thus did not use semantic category information to guide their syntactic analysis. The crucial point, however, is that this information was demonstrably available to them at the point the syntactic analysis was done. Reading times for the first verb (region c − 1) were long when the verb followed the inanimate NP, indicating that readers were sensitive to the fact that the preferred analysis resulted in an anomaly. This fact indicates that eye movements are sensitive in an immediate fashion to syntactically sensitive anomaly effects, providing further evidence for rapid on-line comprehension of sentences (cf. Just & Carpenter, 1980; Rayner, 1983). Nonetheless, the readers apparently did not resolve this anomaly on a semantic basis, but instead waited for syntactic information.

Experiment 2

The second experiment was conducted to determine whether the normal operation of the parser could be altered by the presence of contextual information that biases the interpretation of a syntactically ambiguous string. We recorded subjects' eye movements as they read target sentences that appeared in either Minimal Attachment or Nonminimal Attachment form, as in (4).

(4a) The editor played the tape / and agreed the story was big. (Minimal Attachment)

(4b) The editor played the tape / agreed the story was big. (Nonminimal Attachment)

Notice that the two sentences are identical except for the presence of the word *and* in (4a). The sentence is ambiguous up to the slash, and is disambiguated by the next one or two words (*agreed* vs *and agreed*). These are the active/reduced relative sentences.

In addition to this construction, we included sentences which involve a prepositional phrase attachment ambiguity, such as (5).

(5) Sam loaded the boxes on the cart.

One attachment of the PP (attachment as a daughter of the verb phrase) results in a Minimal Attachment structure, and another attachment (attachment as a daughter of the object NP) results in a Nonminimal Attachment structure. The two different structures are illustrated in Figure 1.

Notice that sentence (5) is fully ambiguous, and so, as it stands, does not force the subject to take a Minimal Attachment or Nonminimal Attachment reading. Therefore, it is impossible to know which reading the subject actually computed. To get around this problem, the target sentences were disambiguated after the critical portion, as in (6).[6]

(6a) Sam loaded the boxes on the cart / *before* his coffee break. (Minimal Attachment)

(6b) Sam loaded the boxes on the cart / *onto* the van. (Nonminimal Attachment)

For each sentence, the italicized word disambiguates the sentence. Notice that the Minimal Attachment and Nonminimal Attachment versions are identical up to the slash. The (6) sentences are disambiguated by the properties of the verb. *Loaded* subcategorizes a NP and a locative PP. In (6a), *on the cart* must be the locative PP since the preposition *before* is not locative. In (6b), *on the cart* cannot be the locative PP since *onto the van* is obligatorily locative.

[6] A second type of prepositional phrase attachment sentence, to be referred to as "conjunction control" sentences, was also used in the experiment. An example is *Sally broke the egg in the bowl, another egg onto the counter, and another egg all over the radio* (MA) or *Sally broke the egg in the bowl, a plate, and a glass of water* (NMA). However, these sentences seemed to confuse subjects badly; they yielded no evidence of recovery from garden paths in the first pass reading times and only some such evidence in the second pass reading times, and resulted in very poor question-answering performance. Therefore, we do not report results from them here. Full details are available in Ferreira (1985).

The phrase *on the cart* must be taken as modifying the boxes, thus producing a Nonminimal Attachment reading.

Because the sentences are disambiguated, the subject is forced to compute the Minimal Attachment or the Nonminimal Attachment reading of the critical portion of the sentence. In addition, the disambiguation permits the construction of questions that tap which reading the subject actually computed, and therefore whether the subject understood the sentence.

Target sentences such as those illustrated in (4) and (6) appeared as the next-to-last sentence in either a biasing or neutral context. The biasing contexts contained information that strongly selected for a particular reading of the target sentence. A biasing minimal attachment context contextually biased the interpretation of (for example) *The editor played the tape* towards a Minimal Attachment reading by mentioning and describing a single editor of whom it could be asserted that he played a tape. A biasing Nonminimal Attachment context biased the string towards a Nonminimal Attachment reading by explicitly mentioning two editors, one of whom was played a tape, the other of whom was shown some photos. The next sentence, the target sentence, referred to the editor who was played the tape. This context should strongly bias the reading of the target sentence, since the reduced relative structure is used to distinguish between two referents previously mentioned.

The neutral contexts contained information that *permits* both analyses, but does not strongly *select* for one or the other. The neutral contexts were constructed in accordance with the points made by Crain and Steedman (1985) concerning the presuppositions of complex syntactic structures. Continuing with the editor example, more than one editor was mentioned, but in contrast with the Nonminimal Attachment biasing contexts, the editors were only mentioned (but not given distinguishing contexts).

The PP attachment ambiguity sentences were contextually biased in a similar way. For example, for the Nonminimal Attachment sentence *Sam loaded the boxes on the cart onto the van*, two sets of boxes were explicitly mentioned, one pile of boxes on a cart and another pile resting on the floor. The Minimal Attachment version (6a) mentioned a set of boxes and mentioned the cart as a potential location to which the boxes could be moved.

Context–target passages appeared in four different forms: Minimal Attachment context–Minimal Attachment target (MA–MA), Nonminimal Attachment context–Nonminimal Attachment target (NMA–NMA), neutral context–Minimal Attachment target (N–MA), and neutral context–Nonminimal attachment target (N–NMA). The four conditions for target sentence (4) are shown in Table 2, and for target sentence (6) in Table 3.

We did not cross contexts and target sentences, in contrast with Crain and Steedman's second experiment. We assumed that crossing of the contexts and target sentences would have produced anomalous passages, and both the modular and interactive models predict that subjects would notice the pragmatic anomaly of crossed contexts and show long reading times. Therefore, the crossed conditions were not considered informative.

The predictions for this experiment are the following: If only the semantically/pragmatically preferred analysis is computed initially, the biased condition (MA–MA and NMA–NMA) should be faster than the neutral conditions (N–MA and N–NMA), since the presence of potentially useful information should facilitate construction of the proper analysis of the target sentence. The N–MA should be faster than the N–NMA condition (since in the absence of a biasing context, the simpler structure should be computed more quickly).

According to the Crain & Steedman (1985) model, in which all analyses are initially considered but the most plausible is

TABLE 2
EXAMPLE OF A RELATIVE CLAUSE AMBIGUITY PASSAGE IN THE FOUR EXPERIMENTAL CONDITIONS, EXPERIMENT 2

Condition	
NMA–NMA	John worked as a reporter for a big city newspaper. He sensed that a major story was brewing over the city hall scandal, and he obtained some evidence that he believed pretty much established the mayor's guilt. He went to his editors with a tape and some photos because he needed their approval before he could go ahead with the story. He ran a tape for one of his editors, and he showed some photos to the other. *The editor played the tape agreed the story was a big one.* The other editor urged John to be cautious.
MA–MA	. . . He gave a tape to his editor and told him to listen to it. *The editor played the tape and agreed the story was a big one.* The other editor urged John to be cautious.
N–NMA	. . . He brought out a tape for one of his editors and told him to listen carefully to it. *The editor played the tape agreed the story was a big one.* The other editor urged John to be cautious.
N–MA	. . . He brought out a tape for one of his editors and told him to listen carefully to it. *The editor played the tape and agreed the story was a big one.* The other editor urged John to be cautious.

very quickly selected, there is no reason to expect a difference between the MA–MA and the NMA–NMA conditions. Both should be fast, because the context in each case fully satisfies the presuppositions of the target sentences. However, either the N–MA or the N–NMA, or both, conditions should be slow. To the extent that the context satisfies the presuppositions of the target sentence less than completely, the syntactically incorrect interpretation of the target sentence will occasionally be chosen. Correcting this garden path will require extra time on some trials. Only if the contexts are absolutely neutral, and *if* a reader can carry along multiple interpretations of a sentence at no processing cost, will the neutral conditions be as fast as the biased conditions.

If only the Minimal Attachment analysis

TABLE 3
EXAMPLE OF A PREPOSITIONAL PHRASE AMBIGUITY PASSAGE IN THE FOUR EXPERIMENTAL CONDITIONS, EXPERIMENT 2

Condition	
NMA–NMA	Sam worked at a factory warehouse. His job was to make sure that boxes of merchandise were ready to be delivered. Sam had to fill up a van so it could go out. He had a pile of boxes on a cart and another pile on the floor. He knew some guys from another department needed the cart. *Sam loaded the boxes on the cart onto the van.* Then he was free to take a much needed break.
MA–MA	. . . Sam wanted to go for his coffee break, but his boss said Sam had to fill up one more cart before he could go. The boss knew some guys from another department needed the cart. *Sam loaded the boxes on the cart before his coffee break.* Then he was free to take a much needed break.
N–NMA	. . . Sam wanted to go for his coffee break, but his boss said he had a little more work to do. He wanted Sam to free up a cart for some guys in another department. *Sam loaded the boxes on the cart onto the van.* Then he was free to take a much needed break.
N–MA	. . . Sam wanted to go for his coffee break, but his boss said he had a little more work to do. He wanted Sam to free up a cart for some guys in another department. *Sam loaded the boxes on the cart onto the van.* Then he was free to take a much needed break.

is computed initially, the NMA–NMA condition should be slower than the MA–MA condition, and this difference in reading time should appear only in and after the disambiguating region of the sentence, not before. The N–NMA condition should be slower than the N–MA analysis. In crucial contrast to the hypotheses discussed earlier, the biased conditions should be no faster than the corresponding neutral conditions, at least early during the reading of the disambiguating region. The biased condition may, however, also be faster than the neutral condition after the disambiguating region of the sentence if contextual information is used quickly to revise an incorrect syntactic analysis.

Method

Subjects. Sixteen students from the University of Massachusetts were paid $5 an hour to participate in the experiment. All subjects had normal uncorrected vision.

Apparatus. The apparatus was identical to Experiment 1.

Materials. Sixteen reduced relative passages were presented to subjects along with eight double-argument prepositional phrase ambiguity passages (plus the eight conjunction control passages mentioned in footnote 6). There were thus 32 items in total, each presented in four different versions. Each target sentence was always followed by a filler sentence in order to prevent subjects from becoming sensitive to the target sentence. The target sentences were arranged so that, across the four conditions in which the sentence appeared, the sentence was in the exact same location on the screen and in the same location within the passage. The passages appeared one at a time to subjects. The entire passage appeared at once, and the passages varied between 6 and 10 lines in length.

A single true/false statement designed to assess the subject's interpretation of the target sentence was written for each passage. For example, subjects were asked to verify *Sam put the boxes on the cart.* The statement is true if the target sentence was Minimal Attachment and false if Nonminimal Attachment. For half the trials, the true/false answers followed this pattern, and for the other half, the true/false response was associated with the Nonminimal Attachment interpretation and a false response with the Minimal Attachment interpretation.

Four lists of passages were constructed. In a single list, a story appeared in only one of the four conditions. Condition numbers were assigned to the items according to a Latin Square design. The order of items within a list was randomized, except that in each list, the first two items were practice passages. There were 34 items in each list (2 practice passages and 32 experimental passages).

Four lists of questions were also constructed corresponding to each list. For 24 of the 32 experimental items, the question was about the target sentence, and for the other 8, the question was about some other part of the passage. (Only three-fourths of the questions were about the target sentence because we did not want subjects to become sensitive to them.) The assignment of target vs filler question was made randomly.

Scoring regions. All sentences were divided into (1) a region prior to any ambiguity, if such a region existed; (2) an ambiguous region; (3) a disambiguating region; and (4) a region after disambiguation, if there was one. Table 4 illustrates the regions of the different sentences. Region (1) is labeled c − 2, region (2) c − 1, region (3) c, and region (4) c + 1. To maintain comparability between the c regions of NMA and MA reduced relative sentences, only fixations on the word or words they had in common were scored. Fixations on the conjunction of the MA sentences (which would be expected to be very short, because of the high frequency of the conjunction; cf. Just & Carpenter, 1984; Rayner & Duffy, in press) were eliminated from consideration.

TABLE 4
SCORING REGIONS OF TARGET SENTENCES, EXPERIMENT 2

Sentence type		Region c − 2	c − 1	c	c + 1
Active vs Reduced Relative	NMA		The editor played the tape	agreed	the story was big.
	MA		The editor played the tape	and[a] agreed	the story was big.
Double arguments	NMA	Sam loaded	the boxes on the cart	onto the van.	
	MA	Sam loaded	the boxes on the cart	before his coffee break.	

[a] Fixations on *and* not scored.

Procedure. The procedure was identical to Experiment 1.

Results

The results of this experiment will be analyzed in terms of (a) first pass reading times, (b) second pass reading times, (c) regressive eye movements, and (d) question-answering data.

First pass reading times. Mean reading times per character for all conditions are given in Figure 2. Figure 2 shows generally longer reading times for Nonminimal Attachment than Minimal Attachment sentences in the disambiguating (c) and subsequent (c + 1) regions. Separate analyses of variance for the two sentence types were performed on the data. The statistical significance of the apparent differences was marginal for the reduced relative sentences (Fig. 2A). The main effect of attachment was significant in the items analysis, $F_2(1,15) = 5.67, p < .03$, and nearly significant in the subjects analysis, $F_1(1,15) = 3.80, p = .07$. The effect of regions was significant, $F_1(2,30) = 12.07, p < .01; F_2(2,30) = 3.90, p < .05$). The regions × attachment interaction did not approach significance. The statistical significance of the apparent effects for the double argument preposition sentences (Fig. 2B) was more satisfactory. There was an effect of regions, $F_1(1,15) = 19.62; F_2(1,7) = 20.71, p < .01$; no effect of attachment; and a significant effect of region by attachment, $F_1(1,15) = 14.58; F_2(1,8) = 9.84, p < .02$). No effect of context approached significance in any analysis.

Second pass reading times. Mean second pass reading times for all conditions are given in Fig. 3. As with the first pass data, separate analyses were conducted for each sentence type. For the active/reduced relative sentences, the effect of attachment was significant, $F_1(1,15) = 8.55; F_2(1,15) = 10.66; p < .01$). The effects of regions

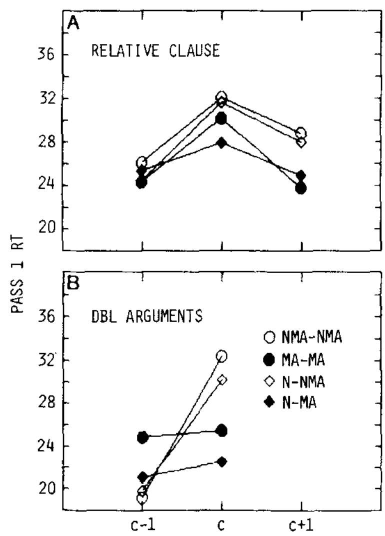

FIG. 2. First pass reading time per character (in ms), Experiment 2.

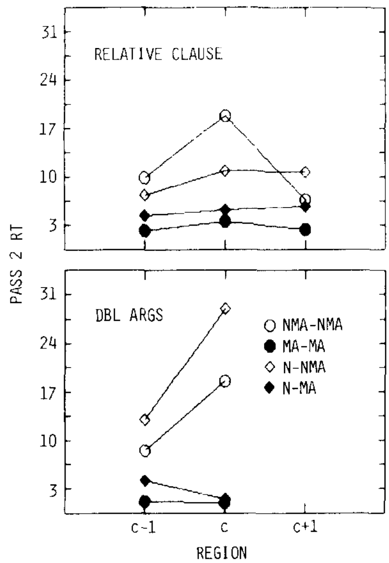

FIG. 3. Second pass reading time per character (in ms), Experiment 2.

and the regions by attachment interaction were significant in the subjects analyses ($F_1(2,30) = 5.31$ and 3.28, $p < .05$, respectively), but not in the items analyses ($F_2(2,30) = 2.49$ and 2.06, $p < .20$, respectively).

For the double argument sentences, the interaction between regions and attachment was significant, $F_1(1,15) = 15.31$; $F_2(1,7) = 16.22$; $p < .01$, as were the main effects of attachment, $F_1(1,15) = 21.46$; $F_2(1,7) = 11.08$; $p < .01$, and regions, $F_1(1,15) = 7.17$; $F_2(1,7) = 5.61$; $p < .05$. Again, no effect of context appraoched significance in any analysis, except for the by-subjects analysis of the context × attachment × regions analysis of the double argument preposition sentences, $F_1(1,15) = 4.20$, $p = .06$; context may have speeded reanalysis slightly.

Regressive eye movements. The number of regressions for each sentence was calculated for each subject. The only regressions that were included in the analysis were regressions from *in* and/or *past* the disambiguating region to a region to the left. The mean numbers of regressions per target sentence for each condition are presented in Table 5. Subjects made more regressions while reading Nonminimal Attachment than Minimal Attachment sentences. This was confirmed in a context by attachment analysis of variance, $F_1(1,15) = 35.94$; $F_2(1,23) = 31.39$; $p < .01$. The main effect of context was not significant, nor was the interaction between context and attachment.

Questions. The percentage correct answers for each condition are also presented in Table 5. Subjects were most accurate in the MA context–MA sentence condition, least accurate in the neutral context–NMA sentence condition, and intermediate in the other two conditions. This pattern of data resulted in a main effect of context, $F_1(1,15) = 6.33$; $F_2(1,23) = 5.52$; $p < .03$. The apparent effect of attachment was nonsignificant, $F_1(1,15) = 2.38$, $p < .20$; $F_2(1,23) < 1$.

Discussion

First pass reading times. Although the results for the reduced relative sentences were of marginal significance, it seems safe to conclude that subjects initially took longer to read the Nonminimal Attachment than the Minimal Attachment sentences. Inspection of Figure 2 makes it clear that this difference was limited to the regions after the disambiguating region, not before, and held regardless of context. This result confirms the predictions made by the modular model of syntactic processing. This model states that contextual information does not affect the initial syntactic decisions made by the syntactic processor, but is used to aid reanalysis of a misanalyzed string. Consistent with this model, longer

TABLE 5
MEAN NUMBER OF REGRESSIONS PER TARGET SENTENCE AND PERCENTAGE CORRECT, EXPERIMENT 2

Condition	Regressions	Percentage correct
NMA–NMA	0.715	78
MA–MA	0.194	89
Neutral–NMA	0.579	71
Neutral–MA	0.216	74

reading times were associated with the Nonminimal Attachment sentences, but not before the parser sent out its error signal. If the results had turned out differently, so that longer reading times were found *before* the disambiguating region, this would have constituted evidence that the parser was building the more complex nonminimal attachment structure on the basis of contextual information, before the parser had syntactic evidence that the Nonminimal Attachment structure was the right one. Since this result did not occur, there is no evidence that the parser altered its normal mode of operation. The parser computes the simplest structure that it can, and computes a more complex analysis only when it receives a syntactic error signal.

Second pass reading time. Subjects clearly took longer to read Nonminimal Attachment than Minimal Attachment sentences after they had made a regressive eye movement. Second pass reading was concentrated in the critical disambiguating regions of nonminimal attachment sentences, but did appear to extend back to the ambiguous region of these sentences, if to a lesser degree.

Regressive eye movements and question-answering. Regressive eye movements reflect (among other things) the parser's attempts to reanalyze a misanalyzed string. In this experiment, subjects made more regressions with the Nonminimal Attachment sentences than with the Minimal Attachment sentences, as predicted by the garden-path theory of sentence comprehension. The effect of context and its interaction with attachment was not significant, suggesting that the only variable that affected the probability of a regression was the syntactic structure of the sentence. In the case of the questions, context was the only variable that had any effect (83% correct in the context condition and 72% correct in the neutral condition). The question-answering data argue that while subjects did not use context to constrain their syntactic analysis of the target sentence, context was used eventually. The target sentence was better remembered when it appeared in a biasing context than when it appeared in a neutral context.

A possible objection to our interpretation of this experiment focuses on the comparison between the Nonminimal Attachment and Minimal Attachment sentences, rather than the crucial comparison between sentences in the neutral vs biasing contexts. It could be argued that the difficulty of the Nonminimal Attachment sentences is due not to the syntactic reanalysis they occasion, but to the fact that extra work must be done during the evaluation of the ambiguous noun phrase to establish its antecedent. To interpret *the boxes on the cart* as a noun phrase referring to some specific boxes on some specific cart, one must have a representation of those boxes in one's mental model. It may be necessary to infer their existence from some other information presented in the context, when faced with the possibility of a phrase that could refer to them. This inference time may be reflected in the slow reading times for Nonminimal Attachment sentences.

There are several difficulties with this objection. First, while this sort of inference might well be needed in the neutral context conditions, where only a vague mention of some boxes is made, no inference is needed in the context that biases toward the Nonminimal Attachment reading. The context asserted that there was a pile of boxes on a cart, and another pile of boxes on the floor. If some extra work is nonetheless required to find the referent of a Nonminimal Attachment phrase, a second very telling difficulty is apparent. Any work of inference or referent-finding should be done *during* the reading of the ambiguous noun phrase, according to a model like Crain and Steedman's, in which the semantic processor shadows the syntactic analysis very closely. There is no reason why the inference work would be delayed until after syntactically disambiguating information was presented—unless of course the inferencing *follows* the syntactically-based resolution of syntactic processing biases. However, no hint of slow reading times was found during the ambiguous phrases when they occurred in a neutral or Nonminimal Attachment context. The only effects appeared after unambiguous syntactic information had been presented.

In summary, this experiment offers strong support for Frazier's (1978) Minimal Attachment principle. The Minimal Attachment strategy is not an unusual strategy adopted by subjects in single-sentence experiments, but a principle which operates across a wide variety of situations. In addition, the experiment provides some data on the question of how reanalysis procedures are carried out, supporting the selective reanalysis hypothesis proposed by Frazier and Rayner (1982).

Experiment 3

This experiment serves as a replication of a part of the second experiment, using a different measure, phrase-by-phrase reading time in a self-paced reading task. The eye movement monitoring task used in the previous experiments, while sensitive, is costly and not available to all investigators, while a task in which text is presented in response to a subject's button-presses can be implemented on any microcomputer. If comparable results can be obtained with the two techniques, further research would be facilitated (cf. Just, Carpenter, & Woolley, 1982; Kennedy & Murray, 1984). Further, while the overall pattern of results from Experiment 2 was clear and orderly, the data from the theoretically much-discussed reduced relative clause sentences were of marginal significance. Therefore, Experiment 3 was conducted to explore the times taken to read such sentences, in the contexts used in Experiment 2, but using a task in which subjects controlled the presentation of segments of sentences using a button-press response.

Method

Subjects. Thirty-two students from the University of Massachusetts Psychology Department human subjects pool participated in this experiment.

Materials. The sixteen active/reduced relative texts from Experiment 2 were presented to subjects together with sixteen filler texts. Subjects saw only one version of each text, and they saw four texts in each of the four conditions. Presentation of experimental and filler text was randomized individually for each subject.

Each story was divided into segments or regions so that segment-by-segment reading times could be obtained. The segments ranged between one and five words, and corresponded roughly to phrases. Each target sentence was divided into regions as follows: Noun phrase, ambiguous segment (verb phrase vs reduced relative clause), disambiguating (critical) segment (conjunction plus verb vs verb), and phrase after the disambiguating segment. For example, in (7), c represents the critical region:

(7) The editor / played the tape /
$c - 2$ $c - 1$
(and) agreed / the story was big.
c $c + 1$

Each story (experimental and filler) was followed by three questions designed to assess the subjects' comprehension of the story. The second question asked them about their interpretation of the target sentence (e.g., *Who played the tape?*), and was written so that a different answer was correct for the minimal attachment and nonminimal attachment sentences. If the subject gave an ambiguous answer, he was probed for more information until he either gave an unambiguous answer or indicated that he didn't know the correct answer. The other question-answering data were not scored, and were included simply to prevent the subjects from becoming aware of the purpose of the experiment.

Procedure. Each of the 32 texts was presented to subjects on a CRT. The initial display for a text presented a dash in place of each letter, but preserved spaces. The subjects pressed a button to bring up the first region of the text. When they had read and understood the segment, they pressed the button again to bring up the next phrase, and the previous phrase was replaced with dashes. The subjects continued in this manner until all 32 texts were read. The reading times for each phrase were recorded.

Results

The mean reading times were computed for each segment of each target sentence, and are presented in Figure 4. These data are difficult to interpret since region length is confounded with structural analysis (minimal attachment sentences have one extra word in the critical region), and longer regions will generally take longer to read. It would be inappropriate simply to divide the reading time for a region by the number of characters in that region, because such a data transformation assumes that reading time is normally a linear increasing function of number of characters, with a value of zero when number of characters equals zero. Since the pacing button reaction time sets a lower bound on the zero intercept of this function, the assumption is clearly wrong. A more nearly adequate position assumes that reading time is normally a linear function of number of characters, with a zero intercept equal to the sum of the button-pressing reaction time and other constant times. A linear regression analysis could be used to estimate the slope and zero intercept of such a function, and thus to estimate the expected reading times for regions of varying lengths. Deviations from these expected times would indicate the existence of factors that speeded or slowed the reading of any given segment.

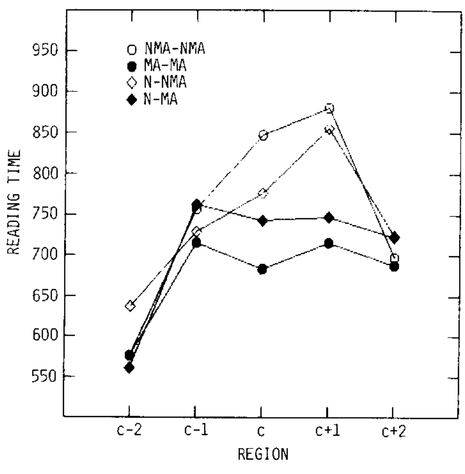

FIG. 4. Mean reading time (in ms) for target sentences, Experiment 3.

Such an analysis was performed by computing the linear regression equation expressing reading time for each segment in each experimental passage as a function of the number of characters in it for each subject. The correlation averaged over all subjects was .38. The regression equation was used to obtain expected reading time on the basis of number of characters alone for each segment. The expected reading times were then subtracted from the obtained reading times and the resulting difference scores were submitted to an analysis of variance.[7]

The mean reading time differences are presented in Table 6. Nonminimal Attachment sentences took longer to read than Minimal Attachment sentences, $F_1(1,31) = 10.98$, $F_2(1,15) = 14.61$; $p < .01$, and there was also a main effect of region, $F_1(3,93) =$

[7] The regression equation does not take into account the fact that the conjunction in the disambiguating region of the Minimal Attachment sentences is high in frequency as well as short. Therefore, it may be slightly biased toward finding longer reading times in region c for NMA than for MA sentences (although no such bias is present in region c + 1, which showed comparable differences). An analysis that treated unregressed reading times, as reported in Figure 4, and which was biased in the opposite direction because of the greater number of characters in region c for MA than for NMA sentences, was also done. This analysis showed a significant effect of attachment, $F_1(1,31) = 14.98$; $F_2(1,15) = 5.03$; $p < .05$, when performed on regions c and c + 1. No other effects were significant in this analysis.

TABLE 6
MEAN DIFFERENCES BETWEEN OBTAINED READING TIME AND READING TIME EXPECTED ON THE BASIS OF REGION LENGTH, TOGETHER WITH PERCENTAGE CORRECT, EXPERIMENT 3

	Region			
Condition	c − 1	c	c + 1	Percentage correct
NMA–NMA	1	240	158	81
MA–MA	−53	2	5	76
Neutral–NMA	−58	176	145	55
Neutral-MA	−23	67	47	63

11.43; $F_2(3,45) = 12.81$; $p < .01$. The interaction of these two variables was also significant, $F_1(3,45) = 4.45$; $F_2(3,93) = 5.87$; $p < .01$, which indicates that the Nonminimal Attachment sentences took more time to read in segments c and c + 1, while for the Minimal Attachment sentences, reading time was evenly distributed across the regions. Finally, context interacted with attachment, $F_1(1,45) = 12.36$; $F_2(1,93) = 5.70$; $p < .03$. The effect of attachment was *larger* for sentences placed in an appropriately biasing context than for those placed in a neutral context.

The Bonferroni t test was used to examine the difference in reading times between Minimal Attachment and Nonminimal Attachment sentences for the two context types ($p < .05$). In both the biased and neutral conditions, subjects took longer to read Nonminimal Attachment than Minimal Attachment sentences in regions c and c + 1, but not in region c − 1. These results were significant on both the subjects and items analyses.

The question-answsering accuracy was computed for each of the four conditions. These data are given in Table 6. The subjects understood the target sentence more often in the biased than in the neutral conditions (81% correct for the Nonminimal Attachment and 76% correct for the Minimal Attachment sentences in the biased condition; 55% correct for the Nonminimal Attachment and 62% correct for the Minimal Attachment sentences in the neutral condition). This effect of context bias was significant in the analysis of variance treating items as a random variable, $F_2(1,15) = 12.30$, $p < .003$, but the effect of sentence form and the interaction between the two variables were not.

Discussion

The results from Experiment 3 fully confirm the conclusions reached in Experiment 2. Nonminimal Attachment sentences are read more slowly in the critical disambiguating regions than are Minimal Attachment sentences, even in the presence of a context which has been predicted to guide sentence analysis and which demonstrably affects the eventual comprehension of sentences. We conclude that the effects observed earlier are robust, even for reduced relative clause sentences, and even using a simple and economical testing procedure.

GENERAL DISCUSSION

At an empirical level, the experiments reported here provide evidence that some syntactic preferences, such as the preference for Minimal Attachment analyses, operate largely independent of semantic, pragmatic, and contextual factors. Such independence of operation seems to hold for the rather subtle preference to attach prepositional phrases as a modifier of a verb phrase rather than as a modifier of a noun phrase (Experiment 2), as well as the more blatant preference for main clause readings over reduced relative clause interpretations (Experiments 1–3). While it remains an open question whether all syntactic preferences show such independence from nonsyntactic factors, we suggest that the present findings do constitute support for the existence of an informationally encapsulated syntactic processor. Following Frazier (1985), we propose that such a processor comprises one or more modules (Fodor, 1983), each of which can process only information stated in its own representational vocabulary. The syntactic processing module, we argue, is responsible for the effects observed here and must be sensitive to phrase structure information (Frazier, 1978; Frazier & Rayner, 1982) and verb subcategorization information (Clifton, Frazier, & Connine, 1984), but it is not sensitive to such nonsyntactic information as the semantic properties of nouns that make them appropriate for certain thematic roles associated with verbs, or the pragmatic information provided by a discourse context.

The second and third experiments showed that the minimal attachment principle holds even in "natural" language settings such as coherent discourse. Proponents of interactive views of language processing charge that studies of syntactic processing in which sentences are presented in isolation are not representative of normal language processing (cf. Van Dijk & Kintsch, 1983). According to this view, in normal language processing all but discourse-initial sentences appear in some discourse context, and the context is crucial for understanding the sentence. Syntactically complex constructions occur in fairly constrained contexts that satisfy the presuppositions of the constructions (Crain & Steedman, 1985), and if they do not occur in such contexts, normal language strategies cannot operate and people resort instead to atypical strategies. The finding from the present study that the minimal attachment strategy operates in discourse challenges this common view. Strategies proposed on the basis of experiments using single sentences as stimuli are not artifactual but appear to operate during discourse comprehension.

While this study has shown that nonsyntactic information is not used during the initial syntactic analysis of a sentence, it is obvious that this information is consulted at some point during sentence comprehension. Any adequate model of sentence comprehension would have to specify at what stage of analysis the information is used, and how it is integrated with other information sources. An important component of such a model is the thematic processor proposed by Frazier (1985: Rayner et al., 1983). The thematic processor operates on the output of the syntactic processor (verbs and their arguments) and, if necessary, proposes alternative analyses that are more semantically or contextually appropriate. The vocabulary of thematic roles is a plausible intermediate representation to mediate between linguistic and nonlinguistic vocabularies. However, the question of how the thematic processor uses information contained in a discourse remains unanswered. It is possible that a discourse processor that represents discourse and sentence topics communicates with the thematic processor and recommends plausible verb arguments based on these topics. Alternatively, discourse information may be represented in a mental model (Johnson-Laird, 1983) which holds in a nonverbal form a structural analogue of the entities and relationships among them described in the discourse.

A second question suggested by this research is how inconsistencies between the outputs of the syntactic and thematic processor are resolved. Some sort of reanalysis of the sentence must take place, but it is not known whether this is a purely syntactic process, or whether higher level information influences the parser at this stage of syntactic analysis. For the present experiments, the reanalysis findings can be explained by a simple syntactic mechanism, since for all the sentences, a syntactic error signal triggered the reanalysis. It is possible that when the parser receives an error signal, it may be primed for some short period of time to compute the less preferred analysis. For example, when the parser analyzes *The editor played the tape* *agreed the story was big*, it receives an error signal in the disambiguating region and consequently reanalyzes the sentence using the only other legitimate analysis it has available.

Nevertheless, a mechanism as simple as this one cannot account for the reanalysis of sentences such as *John saw the girl with the red hair*. The implausible minimal attachment analysis is syntactically well formed, but yet the parser computes the other more plausible analysis. Higher level information can thus cause the parser to compute an alternative analysis of a well-formed sentence, in effect sending out a semantic error signal, and suggesting a different syntactic analysis. It does not seem plausible that the parser would use high level information only to reanalyze fully ambiguous sentences and ignore the information when it has available a syntactic error signal. Instead, the parser uses the error signal as an additional clue to aid in the construction of a different analysis. Of course, these remarks are highly speculative, but begin to address the complex question of how sentences are reanalyzed.

While the issue of reanalysis has yet to be resolved, it is clear from the present experiments that it is at the stage of reanalysis, not initial analysis, that nonsyntactic information sources are consulted. It may seem inefficient for the syntactic processor to ignore potentially useful information contained in discourse during its initial construction of a syntactic structure. However, this is not so. Although a discourse or situational context may be highly suggestive of a complex syntactic structure, it would almost never be anomalous to use the simple active structure instead. It is extremely hard to predict what sort of a syntactic structure a person will use in any context, and there is therefore little to be gained from trying to predict the structure since it will come up shortly anyway. Furthermore, as Fodor (1983) has argued, a processor that consults a restricted information base can in principle operate more quickly than a processor that must retrieve and consider all possible sources of information. Assuming that a syntactic representation must almost always be assigned to a sentence, and given the time pressures associated with language comprehension, it seems more efficient to design the language processing system such that mandatory operations on stimuli can be performed quickly and automatically, and then revised if necessary.

Appendix 1

Critical Sentences Used In Experiment 1

The *baby/skin* (that was) felt by the blind man was very soft and delicate.

The *defendant/evidence* (that was) examined by the lawyer turned out to be unreliable.

The *woman/bill* (that was) paid after the end of the month had worried Mary a great deal.

The *visitor/ship* (that was) sighted by the lookout probably brought bad news.

The *woman/sign* (that was) painted by the artist was very attractive to look at.

The *singer/song* (that was) listened to by the spellbound audience was the most beautiful ever heard.

The *tramp/trash* (that was) smelled by the dog was laying on the sidewalk.

The *man/car* (that was) towed from the parking lot was parked illegally.

The *woman/trees* (that was/that were) expected by the gardeners didn't arrive on time.

The *man/message* (that was) recorded on the tape recorder could not be understood.

The *man/meal* (that was) brought to the high priest could hardly be described as appealing.

The *slave/car* (that was) sold illegally to the plantation owner turned out to be worth very little.

The *man/questions* (who was/that were) asked about the murder could not be answered.

The *children/stories* (that were) told about the incident were a great source of concern.

The *author/book* (that was) read by the student was very hard to understand.

The *girl/postcard* (that was) mailed inside the big box was a warning from the mafia.

Appendix 2

Minimal Attachment (a) and Nonminimal Attachment (b) Target Sentences

Reduced Relatives (Experiments 2 and 3)

1a. The man expected to die but would not give up easily.

b. The man expected to die would not give up easily.

2a. The editor played the tape and agreed the story was big.

b. The editor played the tape agreed the story was big.

3a. The man taught the new method but thought the standard method might be better.

b. The man taught the new method thought the standard method might be better.

4a. The troll brought the princess and thought she looked good enough to eat.

b. The troll brought the princess thought she looked good enough to eat.

5a. The horse raced past the barn and fell in a puddle.

b. The horse raced past the barn fell in a puddle.

6a. The woman told the joke but didn't think it was funny.

b. The woman told the joke didn't think it was funny.

7a. The union sued for damages but didn't expect the settlement to be large.

b. The union sued for damages didn't expect the settlement to be large.

8a. The companies mailed the information but decided to discontinue certain services immediately.

b. The companies mailed the information decided to discontinue certain services immediately.

9a. The man sold the Vega but knew he wasn't getting a very good deal.

b. The man sold the Vega knew he wasn't getting a very good deal.

10a. The woman served the caviar and almost fell into the pool.

b. The woman served the caviar almost fell into the pool.

11a. The man ordered the drink but refused to drink it.

b. The man ordered the drink refused to drink it.

12a. The woman paid the money and left the store immediately.

b. The woman paid the money left the store immediately.

13a. The man read the story and said it just confirmed his suspicions.

b. The man read the story said it just confirmed his suspicions.

14a. The man asked for directions and pulled out a map of New England.

b. The man asked for directions pulled out a map of New England.

15a. The woman delivered the letter and then suddenly got upset.

b. The woman delivered the letter suddenly got upset.

16a. The company awarded the contract and was anxious for the project to get started.

b. The company awarded the contract was anxious for the project to get started.

Double Arguments (Experiment 2)

17a. Katie laid the dress on the floor after her mother yelled at her.

b. Katie laid the dress on the floor onto the bed.

18a. George placed the record on the shelf after they listened to it.

b. George placed the record on the shelf onto the turntable.

19a. Leslie positioned the dress on the rack after shaking out the wrinkles.

b. Leslie positioned the dress on the rack onto the display.

20a. Laura dragged the doll behind the bed before closing the door.

b. Laura dragged the doll behind the bed into the closet.

21a. The clerk put the saleslip in the bag after taking her money.

b. The clerk put the saleslip in the bag into her hand.

22a. Mary set the flowers on the table before her guests arrived.

b. Mary set the flowers on the table onto the cabinet.

23a. The sheriff locked the suspect in his office before he called the FBI.

b. The sheriff locked the suspect in his office into the jail cell.

24a. Sam loaded the boxes on the cart before his coffee break.

b. Sam loaded the boxes on the cart onto the van.

ACKNOWLEDGMENTS

This research was supported by NIH Grants HD-18708 and HD-17246 to the second author, and by a Natural Sciences and Engineering Research Council of Canada Postgraduate Scholarship to the first author. Some of the research reported here was presented as a Master's thesis by F. Ferreira under the guidance of C. Clifton. She would like to thank the members of her committee, Keith Rayner and Lyn Frazier, for their helpful advice and encouragement. She would also like to thank Susan Duffy for technical advice and assistance, Christa Chapman for help in scoring the data, and John Henderson for many useful discussions. The authors would also like to thank Mark Steedman for his helpful criticisms of an earlier report of portions of the present paper, and Marcel Just for valuable suggestions toward the analysis of Experiment 2.

References

Berwick, R. C., & Weinberg, A. S. (1984). *The grammatical basis of linguistic performance*. Cambridge, MA: MIT Press.

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

Crain, S., & Steedman, M. (1985). On not being led up the garden path: The use of context by the psychological parser. In D. Dowty, L. Karttunen, & A. Zwicky (Eds.), *Natural language parsing*. Cambridge: Cambridge Univ. Press.

Ferreira, F. (1985). *The role of context in resolving syntactic ambiguity*. Unpublished master's thesis, University of Massachusetts.

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

Frazier, L. (1978). On comprehending sentences: Syntactic parsing strategies. Bloomington: Indiana University Linguistics Club.

Frazier, L. (1985). Modularity and the representational hypothesis. *Proceedings of the Northeastern Linguistics Society, November 1984*. Providence, RI: Brown University.

Frazier, L., & Fodor, J. D. (1978). The sausage machine: A new two-stage parsing model. *Cognition*, **6**, 291–326.

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.

Graesser, A. D., & Riha, J. R. (1984). An application of multiple regression techniques to sentence reading times. In D. E. Kieras & M. A. Just (Eds.), *New methods in reading comprehension*. Hillsdale, NJ: Erlbaum.

Haberlandt, K. (1984). Components of sentence and word reading times. In D. E. Kieras & M. A. Just (Eds.), *New methods in reading comprehension*. Hillsdale, NJ: Erlbaum.

Johnson-Laird, P. N. (1983). *Mental models*, Cambridge, MA: Harvard Univ. Press.

Just, M. A., & Carpenter, P. A. (1980). A theory of reading: From eye fixations to comprehension. *Psychological Review*, **57**, 329–354.

Just, M. A., & Carpenter, P. A. (1984). Using eye fixations to study reading comprehension. In D. E. Kieras & M. A. Just (Eds.), *New methods in reading comprehension research*. Hillsdale, NJ: Erlbaum.

Just, M. A., Carpenter, P. A., & Woolley, J. D. (1982). Paradigms and processes in reading comprehension. *Journal of Experimental Psychology: General*, **111**, 208–238.

Kennedy, A., & Murray, W. S. (1984). Inspection times for words in syntactically ambiguous sentences under three presentation conditions. *Journal of Experimental Psychology: Human Perception and Performance*, **10**, 833–849.

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

Marslen-Wilson, W., & Tyler, L. (1980). The temporal structure of spoken language understanding. *Cognition*, **8**, 1–72.

Rayner, K. (1978). Eye movements in reading and information processing. *Psychological Bulletin*, **85**, 618–660.

Rayner, K. (1983) *Eye movements in reading: Perceptual and language processes*. New York: Academic Press.

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.

Rayner, K., & Duffy, S. (In press) Lexical complexity and fixation times in reading: Effects of word frequency, verb complexity, and lexical ambiguity. *Memory & Cognition*.

Simon, H. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, **106**, 467–482.

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

Van Dijk, T. A., & Kintsch, W. (1983). *Strategies of discourse comprehension*. New York: Academic Press.

(Received July 10, 1985)
(Revision received November 18, 1985)