25 Sentence Processing: A Tutorial Review

Lyn Frazier
Linguistics Department
University of Massachusetts
Massachusetts, U.S.A.

ABSTRACT

Considerable empirical evidence indicates that perceivers construct a grammatical representation of sentences during comprehension. The principles underlying constituent structure analysis have been investigated in some detail, and their predictions are beginning to be tested in languages other than English. If the results of these investigations continue to be positive, this will encourage the view that psycholinguistics is indeed making progress in the endeavour to develop a theory of human language processing, not just a theory of processing English. However, even within English, the interaction of structural principles with item-specific lexical preferences and with discourse constraints continues to be debated.

Several relatively neglected areas of research are now receiving attention, including theories of recovery from misanalyses, the role of thematic relations in comprehension, and the processing of various types of long-distance dependencies. Conclusions in these areas remain tenative. However, the hypotheses being explored clearly indicate psycholinguistics has changed considerably. The question is no longer whether the human language comprehension system is structured or whether it uses various broad classes of information (e.g. the grammar of the language): Instead, the focus of attention is on fairly detailed and articulated hypotheses about the nature of that structure and the principles underlying the co-ordination of the myriad information sources implicated in language comprehension.

PRELIMINARIES

An adequate theory of language comprehension must do more than describe the means by which the individual sentences of a text are processed and integrated into a coherent structure representing the meaning of the entire text. It must identify the principles determining the analysis of the input, e.g.

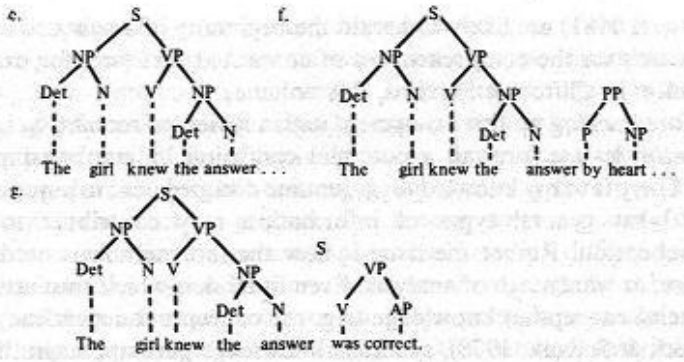

2. Minimal attachment: Do not postulate any potentially unnecessary nodes.
3. Late closure: If grammatically permissible, attach new items into the clause or phrase currently being processed (i.e. the phrase or clause postulated most recently).

This minimal attachment strategy predicts that the transition from (1e) to (1f) should be a smooth one; by contrast, the transition from (1e) to (1g) will require a revision of analysis (addition of the circled "S" node) since the phrase "the answer" will have been taken incorrectly to be the simple direct object of the verb "know." The other strategy implicated in constituent structure analysis is the late closure strategy which adjudicates in cases where two equally minimal attachments exist. It will favour attachments to preceding items (over attachment to subsequent items) and typically will favour attachments to phrases lower in the phrase-structure tree rather than to phrases higher up.

There are several important points to note about this model. The first concerns the generality of the strategies. Many of the construction-specific strategies that had been noted in the psycholinguistic literature follow as a special case of these strategies (see Frazier, 1979, for extensive discussion of this point). For example, Bever's (1970) main clause strategy specifies that perceivers preferentially adopt a main-clause analysis of an input rather than a subordinate-clause analysis. This preference may be viewed as one specific case of minimal attachment. As illustrated in (4), choosing the main clause analysis (4a) over a subordinate clause analysis (4b) follows from choosing the structure with minimal nodes.

4. a. S b. S
S

Other examples of minimal attachment include choosing VP attachment over NP attachment of the PP in (5), main clause analysis of the VP in (6), and the

NP-conjunction analysis of (7). Late closure operates to choose the direct object analysis in (8), low attachment of the adverb in (9) (where it modifies *left*), and of the PP in (10) (where *in the library* will modify *reading* rather than *put*). (This list is by no means exhaustive.) These strategies will determine the analysis of ambiguous strings; however, if a string is locally disambiguated (e.g. by punctuation or by clear prosodic effects) then by definition there will be only one permissible analysis of the input and we would expect perceivers to construct that analysis.

5. John hit the girl with a book.
   (cf. Rayner, Carlson & Frazier, 1983)
   (John hit the girl with a book with a bat.)
6. The horse raced past the barn fell.
   (cf. Bever, 1970; Rayner et al., 1983; Ferreira & Clifton, 1986)
7. Ernie kissed Marcie and her sister. . .
   (cf. Frazier, 1979)
   (Ernie kissed Marcie and her sister laughed.)
8. Since Jay always jogs a mile this seems like a short distance to him.
   (cf. Frazier & Rayner, 1982; Kennedy & Murray, 1984)
   (Since Jay always jogs a mile seems like a short distance to him.)
9. Joyce said Tom left yesterday.
   (cf. Frazier, 1979; Kimball, 1973)
10. Jessie put the book Kathy was reading in the library. . .
    (cf. Frazier & Fodor, 1978)

It is presumably no accident that the construction-specific parsing preferences exhibited by the human comprehension mechanism can be unified and viewed as the consequence of just two general strategies. One can easily imagine a system in which there would be absolutely no relation between the ranking of options at one choice point, and the ranking of options at some other choice point; i.e. a system in which it is necessary to add a new strategy each time a new construction is studied. In light of this possible variation, the systematicity exhibited by the sentence comprehension mechanism is really quite impressive. It is unlikely that it is accidental.

This generality in itself argues that language comprehension is not the result of a haphazard collection of whatever clues can be gleaned from superficial analysis of the lexical string, together with whatever general conceptual or world knowledge might influence the plausibility of various analyses. The experimental evidence presented in support of the strategies provides further evidence of the orderly nature of sentence analysis. For example, recording subjects' eye movements as they read sentences like (11), Frazier and Rayner (1982) show that average reading time per character is longer for (11b) than for (11a). There is also a higher probability of making a

regressive eye movement in (11b) than in (11a), and, as predicted by minimal attachment, the disruption in (11b) is associated with the disambiguating words (underlined in [11b]). Hence this region takes longer to process in (11b) than in the corresponding region in (11c) which is disambiguated by "that" (cf. Rayner & Frazier, in press). In short, the predictions of immediate minimal analysis are confirmed in detail.[1]

11. a. I suppose the girl knows the answer to the physics problem.
(Minimal attachment)

b. The girl knows the answer to the physics problem was correct.
(Nonminimal)

c. The girl knows that the answer to the physics problem was correct.
(Nonminimal-unambiguous)

The fact that the complexity associated with a nonminimal attachment sentence such as (!1b) shows up immediately (including on the very first fixation in the disambiguating region) argues strongly for the view that the processor systematically constructs a linguistic analysis of the input as the words of the sentence are encountered. If no analysis had been assigned to preceding items in (11b) there is no reason whatsoever to expect long fixations and regressive eye movements to be associated with the disambiguating region of the sentence.

Before turning to preferences based on nonstructural factors, let me emphasise that the minimal attachment and late closure strategies are not arbitrary; one can understand, for example, why different individuals should each adopt these particular strategies and not, say, their inverses. Both minimal attachment and late closure may be viewed as the result of adopting the first analysis available to the processor. Minimal attachment analyses will be available earlier than nonminimal ones due to the relative number of phrase structure rules that must be accessed for the two analyses, assuming that accessing more rules takes more time (see discussion in Frazier & Fodor, 1978). Late closure analysis permits earlier structuring of new unstructured items than would its opposite (i.e. early closure), since new items may be structured together immediately with already processed material. Hence, the structuring of the current item need not be delayed until subsequent items are received and processed.

Assuming that the need to structure material quickly is related to

[1]In Holmes, Kennedy, and Murray (in press), also the first experiment in Holmes (this volume), no difference in the complexity of (11b) and its unambiguous counterpart (11c) is observed, contrary to the predictions of minimal attachment. However, reading times are very long in this study. When subjects simply read for comprehension, significant garden-path effects are observed, cf. Rayner and Frazier, in press.

restrictions on human immediate memory capacity, we might expect all humans to adopt the first available constituent structure analysis. If so, we expect the minimal attachment and late closure strategies to be universal. Ideally we should be able to remove the grammar of English from our theory of sentence processing, plug in the grammar of some other language, and obtain the correct theory of the processing of that language. To the extent that this is so, the theory is in fact a theory of human sentence comprehension, and not merely a theory of the processing of English. And, if it should turn out that language-specific parsing differences do exist, this fact along with the detailed differences will require explanation.

It wouldn't be too surprising if the constituent structure parsing of all head-initial languages (in which the heads of phrases precede their complements) is the same as in English. Indeed, in a head-initial language like Italian, for example, there is at least intuitive evidence supporting the prediction of the above strategies. Marica deVincenzi informs me that there is a preference for noun-phrase conjunction, over sentential conjunction, in cases of temporary ambiguity (e.g. the Italian counterpart to [7]), as predicted by minimal attachment. Frazier (Note 6) investigates Dutch. Dutch is a "verb-second" language exhibiting head-final order in all embedded verb phrases (i.e. the verb follows its objects). Frazier presents both intuitive evidence and initial experimental evidence confirming the predictions of minimal attachment in Dutch. Ueda (Note 13) uses intuitive evidence to argue for the operation of minimal attachment in Japanese, a consistent left-branching (head-final) language (see also discussion in Frazier & Rayner, in press). Clearly more cross-language processing evidence is needed. However, the currently available evidence is encouraging in that it consistently supports the predictions of the strategies, even in languages which are typologically distinct from English.

Viewed from the perspective of this model, the question of how nonstructural information types influence sentence processing is a question concerning the interaction of information types. For example, given the basic model outlined here, one can ask how the processor is influenced by lexical, thematic, or world knowledge which disconfirms the structurally preferred analysis of the sentence. Viewed from the perspective of competing models (e.g. Crain & Steedman, 1985; Ford, Bresnan, & Kaplan, 1983), the question of how nonstructural information types influence processing may be taken to be a question concerning the nature of the decision principles determining the analysis of the input. The evidence which is crucial for distinguishing these two views is whether the information influences the selection of the *initial* syntactic analysis assigned to the string. If so, then the purely structural decision principles discussed above are either wrong altogether or they operate in more restricted circumstances than suggested here. The role of lexical preferences and of discourse constraints will each be taken up directly.

We shall then turn to questions concerning the revision of an initially computed structure.

Lexical Preferences

Ford et al. (1983) propose a series of principles to guarantee that a sentence is initially analysed in accordance with the strongest or preferred lexical form of the verb (and other heads of phrases). Indeed, they show that the ultimately preferred analysis assigned to the sentences in (12) differ, presumably because "want" occurs more often with just a single argument in its complement whereas "position" typically occurs with both a theme and a locative phrase in its complement.

12. a. The woman positioned [the dress] on that rack.
(Simple NP preferred)
b. The woman wanted [the dress on that rack].
(Complex NP preferred)

Numerous experimental studies show that lexical preferences influence some stage of language comprehension (e.g. Clifton, Frazier, & Connine, 1984; Holmes, this volume; Kurtzman, Note 9; Mitchell & Holmes, 1985). What these studies do not show is whether lexical preferences are used to guide the initial selection of an analysis or, alternatively, are used only later (e.g. as a filter confirming or disconfirming an analysis identified on other grounds, or contributing to the ease and speed with which a temporary misanalysis is corrected). I shall refer to the former possibility as "lexical proposal" and to the latter as the "lexical filter" view.

Holmes (this volume) discussed a grammaticality-judgement study conducted with Lauri Stowe which she takes to be evidence unambiguously favouring the lexical-proposal view. The basic finding is that a numerically smaller garden-path effect occurs in sentential-complement sentences like (11b) containing verbs, which are biased towards a sentential complement (e.g. *claim*) than in sentences with verbs biased toward a noun-phrase complement (e.g. *hear*). On Holmes' account, perceivers first test a sentential complement analysis for verbs like *claim*, but not for verbs like *hear*. For some reason (I actually don't understand why, since sentential complements usually begin with subject noun phrases) this sentential complement is abandoned when a noun phrase follows the verb. Thus, when the noun phrase is followed by a verb, perceivers must give up their current structural hypothesis and return to their initial (sentential-complement) hypothesis. This reanalysis will be relatively cost-free for a verb like *claim* since it involves re-establishing a hypothesis already considered; for a verb like *hear*, the sentential-complement analysis will not be considered until the disambiguating information is encountered and thus the garden-path effect is larger.[2] Notice, however, that an alternative account of these data is possible. Imagine for a moment that perceivers tend to minimally attach the postverbal noun phrase into the verb phrase following both types of verbs. We must then assume it takes longer to revise a lexically confirmed analysis than to revise an analysis which is inconsistent with the preferred or predominant usage of verbs. In short, this study seems compatible with either a lexical-proposal or a lexical-filter hypothesis (also see footnote 2).

Mitchell (this volume) presents evidence supporting the use of lexical preferences to filter or evaluate the structural analysis of a sentence. In a self-paced reading task, he shows that reading the first display (underlined here) took longer in (13b), containing an obligatorily intransitive verb, than in (13a), containing an optionally intransitive verb. In the second display (not underlined), the a-form takes longer than the b-form.

13. a. After the audience had applauded the actors/sat down for a well-deserved drink.

b. After the audience had departed the actors/ sat down for a well-deserved drink.

According to a lexical-proposal view, we would have expected perceivers to construct the correct intransitive analysis of (13b), whereas in (13a) perceivers would on occasion at least be expected to construct the incorrect analysis. Instead, the pattern of data suggest that a transitive analysis is initially taken in both sentence forms. The incorrect (transitive) analysis is reanalysed during the first display in (13b) but not until the second display in (13a). The lexical-proposal hypothesis does not provide a straightforward account of these data.

---

[2]Holmes (this volume) takes another aspect of her data to support the lexical-proposal hypothesis. She finds reading times for the word *the* (relative to its "unambiguous" control) to be significantly longer following S-bias verbs like *claim*, but not following NP-bias verbs like *hear*. According to Holmes, subjects expect *that* following *claim* and thus are garden-pathed when *the* appears. There are two problems with this argument. First, the appearance of *the* is perfectly consistent with the hypotheses that *claim* takes a sentential complement (i.e. many clauses begin with a subject noun phrase); hence the appearance of *the* does not disconfirm the reader's current structural hypothesis even on the lexical-proposal hypothesis. If anything, the occurrence of the word *the* disconfirms a prediction about the category of the following word, *not* a prediction about the structural type of the verb's complement. Second, if probabilistic predictions about the category of an immediately following word influence the nature or timing of lexical categorisation of ambiguous lexical items (e.g. *that*), then relatively long reading times associated with the word *the* in the control condition, not relatively short times in the reduced complement of NP-bias verbs, might be responsible for the effect. In short, there are several reasons to think the observed effect may not be directly related to the processor's hypothesis about the structural type of the complement.

Two additional types of evidence suggest that lexical proposal of a structural analysis is not correct, at least not in any fully general form. Sentences like (10) ("Jessie put the book Kathy was reading in the library ...") suggest that lexically guided analysis would have to be based on local lexical preferences. Intuitions indicate that the prepositional phrase in (10) is initially interpreted as a sister to "read" despite the fact that "read" does not require a locative phrase but "put" does.

In (10), it might suffice to say that it is only *local* lexical preferences which guide analysis. But in sentences like (14a), one preferentially interprets "to whom" as binding the gap after "say," despite the fact that "say" does not in general seem to occur with an overt goal or any other prepositional phrase. Thus, assuming local lexical preferences guide analysis would not account for this preference.

14. a. To whom did you say that you admitted it?
    b. To whom did you admit that you said it?

Further, since the preferred analysis of (14b) also has "to whom" binding a gap in the matrix clause (the gap following "admit"), the preference in (14a) cannot be attributed to some factor like the relative frequency of say + PP vs. admit + PP. In short, the intuitive evidence in (14) is really not consistent with a lexically guided parsing system, whether lexical guidance is construed as a purely local phenomenon or as a global competition between the lexical preferences of all heads of phrases.[3]

One further bit of evidence weighing against the use of lexical preferences to propose (rather than filter) an initial structural analysis derives from a study of head-final phrases. If the structural analysis of a phrase is initially governed by the lexical preferences of the head, then analysis of verb-final verb phrases should be delayed, i.e. one can't determine the lexical preferences of the verb before the verb is encountered. In other words, the lexical proposal view entails the existence of some language-specific parsing differences. Thus, we may either assume that initial structural analysis is based on

[3]Tanenhaus, Stowe, and Carlson (1985) examined word-by-word reading times for sentences like those in (a)

a. The sheriff wondered which horse/rock the cowboy raced _____ down the hill.

b. The sheriff wondered which horse/rock the cowboy raced desperately past _____.

They observed a plausibility effect following the verb (e.g. "raced the horse" is more plausible than "raced the rock") only for preferred transitive verbs. This might be taken as evidence in support of the lexical proposal view. For the plausibility relations of only the lexically preferred analysis to be important is entirely expected on the view that only this analysis is computed. However, as Mike Tanenhaus (pers. communication) has pointed out to me, the finding is also consistent with the lexical filter view on the assumption that subcategorisation filtering precedes semantic interpretation and/or pragmatic evaluation.

lexical preferences *only* in head-initial phrases; or we may assume parsing is lexically guided in all languages, but structural analysis occurs immediately, without waiting for the head, only in head-initial languages.

At present there is only extremely limited evidence available concerning the parsing of head-final constructions. However, the evidence does not indicate delays in the analysis of head-final constructions (cf. Frazier, Note 6). This, in turn, favours the view that lexical preferences are used to filter rather than propose structural anlayses since it argues against any universally operative lexical-proposal strategy (which incorrectly predicted delays). Of course, ultimately it may turn out that differences in the parsing mechanisms for distinct languages do exist; but we're not forced to that position yet.

Discourse Constraints

The relation between sentences in a discourse is complex. Sequences of sentences typically maintain not only referential coherence, but temporal and causal coherence. Sentences early in the sequence satisfy the presuppositions of later sentences, determine or constrain the topic/focus structure of subsequent sentences, and provide the context or partial model relevant for the interpretation of later sentences, thereby mediating the operation of Gricean principles. Thus, like lexical preferences, discourse biases are important determinants of the final understanding of a sentence. As with lexical preferences, the question of immediate interest is precisely when and how these biases operate; e.g. do they influence the initial choice of a grammatical analysis of the current input item.

Crain and Steedman (1985) propose various principles intended to guarantee that the structure assigned to the current sentence is whatever structure is maximally compatible with the current discourse model, i.e. whatever analysis requires the fewest revisions or additions to the current discourse model. Specifically, they propose a *principle of a priori plausibility* (choose the most plausible reading in terms of world knowledge and the universe of discourse), a *principle of referential success* (favour readings referring to an entity already established in the perceiver's model), and a *principle of parsimony* (other things being equal, choose a reading that carries fewer unsatisfied but consistent presuppositions).

The observation that plausibility, parsimony, and referential success govern some stage of language processing is widely accepted. What is novel about Crain and Steedman's proposal is the idea that the consequence of applying these principles can be determined immediately, essentially as each word or two is encountered. If true, this would permit the processor to construct basically just a single analysis of an input (except for very short one- or two-word stretches of the sentence), while maximising the chances of

computing the correct/intended/contextually appropriate analysis of the sentence.

Having received and analysed an entire sentence, it is certainly possible to identify each of the locally possible analyses of the sentence and determine the presuppositions carried by each, the relative a priori plausibility of each, and the relative referential consistency of each with preceding discourse context. But under what circumstances can this be determined on a word-by-word basis? Traditionally, at least, it seems to have been assumed that one must interpret a structure in order to determine the presuppositions it carries.

With respect to plausibility differences, it is often unclear what would count as a sufficient difference in the plausibility of different readings to permit early resolution of ambiguity. For example, in (15), is the difference in the likelihood of "answer" vs. "duck" as the direct object of "knows" sufficiently great that a simple direct object analysis is selected in (15a) but a sentential complement reading in (15b)? And how could the parser (even in principle) evaluate the relative plausibility of the direct object and sentential complement reading without yet knowing the identity of the embedded verb? Clearly it cannot simply assume, say, that the higher a phrase is on some animacy hierarchy, the more plausible it is as a subject; this might work in (15), but not in general (e.g. consider [16]).

15. a. John knew the answer ...
    b. John knew the duck ...
16. a. John heard the answer ...
    b. John heard the duck ...

Without knowing what constitutes a decisive bias during the left-to-right ongoing analysis of a sentence, we do not know how the above principles apply. We cannot simply appeal to the principles post hoc whenever it is convenient to do so. Further, without some rough idea of the answers to these questions, we do not know when the processor is predicted to assign just a single analysis to an input and when it must compute and maintain several alternative analyses because relative plausibility (referential success or parsimony) differences are not sufficiently great or clear to be decisive.

Despite these questions, there are some circumstances in which clear local plausibility or parsimony differences may be identified. For example, Crain and Steedman note that the local relation between "teacher" and "taught" and between "children" and "taught" will result in clear plausibility differences in (17), i.e. the correct reduced relative-clause analysis will be more plausible in (17b) than in (17a). Indeed, in a grammaticality judgement experiment, Crain and Steedman show that subjects are more likely to call (17b) grammatical than (17a).

17. a. The teachers taught by Berlitz method passed the test.
    b. The children taught by Berlitz method passed the test.

This result (and the others reported by Crain & Steedman, 1985; Kurtzman, Note 9; and Altman, Note 1) might reflect either the initial analysis constructed or, alternatively, the ease of reanalysis. In short, this finding simply cannot choose between the possibility that principles like a priori plausibility govern an initial choice of syntactic analysis and the possibility that plausibility influences the probability and ease of reanalysis.

Ferreira and Clifton (1986) try to distinguish these possibilities by measuring reading times for ambiguous sentence structures such as (18) in disambiguating contexts, e.g. (19) as a context for (18a) and (18b).

18. a. The editor played the tape agreed the story was big.
    b. The editor played the tape and agreed the story was big.
    c. Sam loaded the boxes on the cart onto the van.
    d. Sam loaded the boxes on the cart before his coffee break.
19. John worked as a reporter for a newspaper. He knew a major story was brewing over the mayor scandal. He went to his editors with a tape and some photos because he needed their approval to 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 big. The other editor urged John to be cautious.

Even in a strongly biased context such as (19), there was clear evidence of a temporary misanalysis in the nonminimal attachment sentence forms (18b) and (18d). These results do not argue against the existence of the plausibility, parsimony, or referential success principles as principles which govern some stage of processing; but they do argue strongly against the use of these principles to select an initial syntactic analysis of a sentence in any immediate or nearly word-by-word fashion.

Issues of Reanalysis

In any depth-first model of processing (whether analysis is guided by structural, lexical, or discourse principles) there will be circumstances when it will be necessary to abandon the current analysis of the sentence. The particular circumstances requiring reanalysis will of course depend on the particular principles assumed; but the mere fact that the interpretation of words and phrases may depend not just on a priori context but also on subsequent context entails that every depth-first on-line processing model will need some theory of reanalysis.

Gorrell (1985) proposes one interesting account of reanalysis. In essence, he suggests that syntactic analysis is not completely depth-first or serial. Rather, it is "staggered serial." On this account, more than one analysis of the input is computed. The simplest analysis is initially adopted, but if it proves to be incorrect, the alternative will still be available; e.g. the sentential complement analysis will still be available when "was" is encountered in (20a) (Gorrell, 1985, p. 188).

20. a. John knew the old woman on the train was ill.
    b. They motioned to the man that they couldn't hear.
    c. *John left is surprising.
    d. That apparently healthy sheep die is disturbing.
    e. The horse raced past the barn fell.

The basic claim of the staggered serial account is that "easy" reanalyses—those which do not involve conscious effort—are just those for which the alternative analysis is still available at the point of reanalysis.

Two factors will determine whether an alternative analysis will be available when needed: whether the analysis is computed in the first place and the "distance" between the initially adopted analysis and the alternative, which Gorrell suggests is a function of how long the parser has been committed to the first analysis. According to Gorrell (1985, p. 194): "parallel processing only occurs in the environment of an overt marker which is ambiguous, e.g., an ambiguous verb or lexical item such as 'that'." Thus, reanalysis of (20a) is easy because the alternative to the direct object reading is computed and is still available when the parser encounters the error signal "was." Likewise, in (20b) reanalysis will be simple due to the presence of "that" which permits both analyses to be computed; (20c) is harder, because no overt element is present to mark the alternative analysis. And (20d) is claimed to be hard because of the distance between the error (the analysis of "that" as a demonstrative rather than as a complementiser) and the error signal ("is"). Finally, according to Gorrell, the reason why (20e) involves conscious reanalysis is because no alternative to the simple main-clause analysis is computed in the first place.

Unfortunately, the staggered serial account will not do, even for the limited data in (20) (which provided the basic motivation for this view). Though I cannot go into all the issues in depth here, this approach suffers from several shortcomings. First, counter to Gorrell's discussion, the principles he gives actually predict that (20e) should be *easier* to revise than (20a): Since the verb ("raced") is ambiguous, two anlayses of (20e) should initially be computed; and, the amount of time the parser is committed to the initial analysis is actually less in (20e) (where the misanalysis only spans four words of the input) than in (20a) (where it spans six words of the input). Clearly

Gorrell intends for only a single analysis to be computed in (20e). But it is entirely unclear what principled definition of ambiguous verb would classify an item with multiple subcategorisation frames as ambiguous (e.g. "know") but would not classify verbs with passive participles which are homophonous with their simple past (e.g. "raced") as unambiguous. Thus, it seems the staggered serial model does not really have an account of the circumstances under which an alternative analysis is computed.

Gorrell notes that the failure to find robust reanalysis effects for sentences with very short (one or two word) ambiguous phrases, as in the closure sentences in Frazier and Rayner (1982) and Kennedy and Murray (1984), may indicate that readers may maintain alternative structural descriptions for a small number of words—for two-word phrases, but not four- or five-word phrases in the studies cited. This assumption, however, is itself incompatible with the discussion of the examples in (20); e.g. the statement that reanalysis of (20a) is easy because the alternative is still available after a *six*-word phrase. I emphasise this point because the apparent impossibility of giving any consistent explicit account of the duration of an alternative structural analysis of a sentence is the major empirical obstacle to every version of parallel processing I've encountered (e.g. consider the parallelism implicit in discourse driven parsing models such as the one discussed earlier). It is of course to Gorrell's credit that he has proposed a sufficiently explicit account of limited parallelism that one may argue against it.

A further empirical difficulty arises for the staggered serial account. The account predicts (see discussion in Gorrell, 1985, p. 195) that processing verbs with multiple subcategorisation frames (e.g. optionally transitive verbs) should be harder than processing verbs with a single frame (e.g. obligatory transitive or obligatory intransitives). This prediction is falsified in Clifton et al. (1984) and in Mitchell (this volume; discussed earlier). Finally, it should be noted that in its present form at least, the staggered serial account of reanalysis is ad hoc. Gorrell (1985, p. 194) correctly notes: "Presumably, there is a cost to constructing parallel representations and the parser will not so do without sufficient reason." But it is not clear why the parser should engage in parallelism just for the particular subset of analyses proposed. Nor is it clear why, if it is bothering to construct parallel analyses at all, it should initially adopt just one analysis, or why it should choose the syntactically simpler one. In short, though the staggered serial account is probably the most extensive, explicit and interesting account of reanalysis to date, it can be faulted on both descriptive and explanatory grounds.

Frazier and Rayner (1982) argued that the complexity of revising an analysis is a function of the "clarity" of the error signal, i.e. whether the signal that the current analysis is ill-formed or inappropriate indicates the location and nature of the error. Like Gorrell, they suggest that the longer the parser is committed to an analysis, the harder it tends to be to revise. In

some cases this might be due to the difficulty of locating the original error if the error signal is delayed, as in (20e) and (24). However, in many other cases, minor differences in the length of an ambiguous phrase results in substantial complexity differences. These may be attributed to the greater cost of revising a syntactic error once it has already been semantically interpreted (see discussion in Frazier & Rayner, 1982).

Take (21), for example. The minimal attachment analysis of the words "the men" (shown in [21a]) seems perfectly easy to revise: Indeed, once the preposition is encountered, there is no alternative to the analysis shown in (21b).

21.

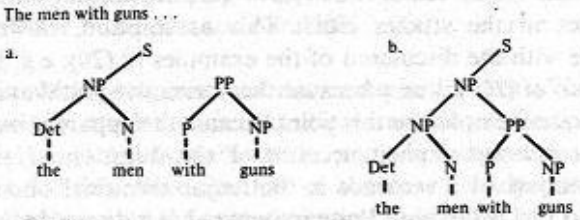

In other words, the revision shown in (21b) will be unambiguously warranted by the occurrence of the prepositional phrase (since the grammar of English will allow no alternative attachment of this phrase). Thus the revision in (21) should be trivial. However, given a discourse context which already contains a potential referent for the simple NP (as in [22] where "the men" in the second sentence might initially be taken to refer to the men introduced in the first sentence), it is possible that the simple NP analysis in (21a) leads to an incorrect semantic interpretation. It is only in this latter case that one is aware of any difficulty in processing the sentence.

22. There were men standing on the corner who were unarmed.
The men with guns were on the roof.

Sentence (23a) is an interesting case of revision. Intuitively it is more difficult than the closely related sentence in (20a) but less difficult than (20e). In (23a) the error signal "hid" is informative with respect to the nature of the error; "hid" will be missing a subject, assuming *the girl* is initially minimally attached as the direct object of *see*. Further, since the misanalysed NP (the girl) is the immediately preceding phrase, the location of the error (as well as its nature) is flagged by the error signal. Thus, (23) constrasts with (20e) where only the nature of the error ("fell" is again missing a subject) but not its location is marked by the error signal. However, unlike (20a), perception verbs permit a predication analysis (cf. [23b] and [23c]). Hence, the fact that *hid* requires a subject does not by itself require the processor to analyse *the*

girl as the subject of a sentential complement. Rather, in this case, it is only the presence of the tense marker which excludes the predication analysis of (23a) (cf. Williams, 1980), requiring the direct object analysis to be revised.

23. a. I saw [[the girl] [hid in the forest]].
b. I saw [the girl] [hidden in the forest].
c. I saw [the girl] [hide in the forest].

The idea that it is the informativeness of the error signal which determines the complexity of reanalysis explains another observation about the difficulty of making revisions, namely, that errors at one level of analysis tend to be difficult if they are detected only at some later level, especially if they appear to have been confirmed at this or some other subsequent level of analysis. For example, in (24), due to Marcus (1980), *cotton* is lexically misanaysed (as a derivative adj. or as part of a compound) but the error is only later signalled by the syntactic illformedness which ensues as a result of this lexical analysis.

24. The cotton clothing is made from grows in Mississippi.

Eventually a theory of reanalysis must account explicitly for all differences in the ease of making revisions. It must also account for why and how the processor manages to identify an overlooked analysis of the sentence if the intially chosen analysis does not break down. For example, the first analysis may be completely well formed, but simply not quite as plausible on semantic and pragmatic grounds as some alternative analysis. In response to this problem, Rayner et al. (1983) propose the existence of two independent processing subsystems: a syntactic processor and a thematic processor. The two subsystems operate in parallel, each carrying out its own idiosyncratic task on the basis of its own operating principles and its own characteristic information sources. By hypothesis, the thematic processor evaluates the relative plausibility of all thematic frames associated with the head of a phrase using discourse context and world knowledge. It may thus choose some frame which is not compatible with the initially selected syntactic analysis of the phrase.

We might assume that a thematic frame for a verb will always contain one external argument (i.e. one argument appearing outside the verb phrase) and possibly one or more internal arguments. A thematic role and syntactic category label will be supplied for each argument, as illustrated for "see" in (25). By convention, the external argument (the experiencer in [25]) is underlined; thus, thematic frames permit the external argument to be distinguished from any internal arguments. Assuming all internal arguments must be sisters to the head, a frame like (25a) will be consistent only with a

syntactic analysis in which "see" has exactly one sister phrase and it is a noun phrase. Consequently, if the syntactic processor builds the (minimal attachment) structure in (26) for the phrase "saw a cop with a revolver" and the thematic processor selected the frame in (25a), there will be a clear conflict: In (26) "see" has two sisters but (25a) claims it has only one.

25. see
    a. [Experiencer Theme]
       NP NP
    b. [Experiencer Theme—Instrument]
       NP NP PP
    etc.
26. The spy saw the cop with a revolver . . .

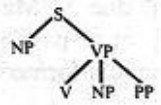

This conflict can act as an error signal, indicating to the syntactic processor that there is a locally more plausible analysis of the sentence. Rayner et al. present eye-movement data in support of this account.

The general idea that thematic structure actively participates in sentence analysis by confirming a syntactic analysis or initiating reanalysis attempts receives further support by a (makes-sense judgement) study conducted by Mike Tanenhaus (personal communication).

27. a. A gambler visited his cousin tonight. (Transitive expectation)
    b. A gambler visited tonight.
28. a. A gambler cheated his cousin tonight. (Intransitive expectation)
    b. A gambler cheated tonight.
29. a. The girl was walking on the grass. (Locative expectation)
    b. The girl was walking.
30. a. The girl was reading on the grass. (No locative expectation)
    b. The girl was reading.

If a sentence contains a direct object noun phrase or a locative prepositional phrase (as in the a-forms above), the processor has no alternative other than to analyse the phrase syntactically and assign a thematic role to it, whether the occurrence of the phrase is lexically predicted or not (i.e. whether the phrase conforms to the most frequent usage of the verb or not). However, in the simple intransitive sentence forms without a locative phrase (the b-forms above), the processor will expect a particular thematic role that won't arrive if the verb is one that usually occurs with two arguments (e.g. *visit* usually occurs with both an agent and a theme, *walk* usually occurs with both an agent and a locative phrase). Thus we might expect the facilitation from a match between the actual sentence form and the preferred verb form to be greater in the "short" b-forms than in the a-forms. This is precisely what Tanenhaus found, i.e. preferred transitives (locatives) took longer than preferred intransitives (nonlocative) sentences, with a crossover interaction—preferred monadic verbs did not take longer in the dyadic (b-) sentence forms. Of course, if thematic predictions governed the selection of an initial analysis (rather than playing a confirmatory role or a role in reanalysis) we would not expect this crossover (see also Carlson & Tanenhaus, Note 3; Stowe, Note 11).

Summary

We have reviewed the evidence supporting the claim that the human sentence processor initially computes just one constituent structure analysis of a sentence. We have also examined the principles which determine the particular analysis the processor tries first, including proposed structural principles, lexical principles and discourse principles. It was emphasised that quite different theories of sentence processing result depending on whether the proposed nonstructural (lexical and discourse) principles guide initial syntactic analysis or alternatively reflect the fast operation of a theoretically distinct (nonsyntactic) processing subsystem. The lexical and discourse principles discussed here may be viewed either way: They may provide an alternative to structural principles and thus guide which particular syntactic analysis the processor tries first, or they may be viewed as principles which guide subsequent analysis, and thus only modulate syntactic preferences, by determining the eventual fate of the initially selected syntactic analysis. Though most experimental evidence is ambiguous and may be interpreted either way, in my (clearly partisan) opinion the current evidence which is not open to either interpretation (see discussion earlier of Mitchell, this volume; Ferreira & Clifton, 1986) favours the view where lexical and especially discourse principles do not directly influence syntactic analysis by determining which analysis the processor will initially construct. Finally it was emphasised that any depth-first theory of sentence processing must ultimately offer a detailed account of reanalysis. At present only the barest outlines of such a theory have been sketched in.

PROCESSING LONG-DISTANCE DEPENDENCIES

The previous section was concerned with the recovery of constituent structure relations. Here we will briefly consider long-distance grammatical dependencies, such as the so-called "filler-gap" dependencies, e.g. the relation between "what" and the empty object following "eat" in (31).

31. What did John eat ____?

We will refer to any empty position in the constituent structure representation as a "gap" and refer to any phase which controls the interpretation of a gap as a "filler" (regardless of the grammatical nature of the relation involved).

Traditionally (Fodor, 1978; 1979; Jackendoff & Culicover, 1971; Wanner & Maratsos, 1978) it has been assumed that gaps are only identified if the processor predicts the occurence of a phrase of a certain type, and then the input lexical string fails to contain an item of that type in the appropriate position. This gap-identification strategy would correctly account for the fact that (32a) is the preferred interpretation of (32), where the gap is taken to correspond to the second of two postverbal noun phrases (though see Woolford, 1986 for an alternative account of these facts and Frazier, Note 6, for a discussion of problems with this gap-identification strategy).

32. Which patient did the nurse bring the doctor?
    a. Which patient did the nurse bring the doctor ____?
    b. Which patient did the nurse bring ____ the doctor?

It is clear that the postulation of gaps is influenced by lexical preferences (cf. Clifton et al., 1984; Fodor, 1978; Tanenhaus et al., 1985). However, just as in the case of nonempty phrases, there remain questions about whether the expected form or use of a lexical item influences initial-gap hypotheses or just the subsequent evaluation of gap-hypotheses based on nonlexical information (see discussion earlier).

It has usually been assumed that the assignment of a filler to a gap is accomplished immediately, as soon as the gap is postulated. Several indirect sources of evidence support this assumption (e.g. Crain & Fodor, 1985; Frazier et al., 1983; Stowe, 1986; Tanenhaus et al., 1985; Wanner & Maratsos, 1978). Recent evidence from auditory (Clifton & Frazier, in press; Swinney, personal communication) and visual (Clifton & Frazier, in press) priming studies present direct evidence in support of this assumption.

With respect to the assignment of fillers to gaps, it is quite clear that perceivers of English will assign the more recent of two grammatically permissable fillers in cases of ambiguity (cf. Crain & Fodor, 1985; Fodor, 1978; Frazier, Clifton, & Randall, 1983).[4] There is also intuitive evidence that this "recent filler" strategy applies in other languages (see Engdahl, Note 4, for results of an extremely interesting study in Swedish). What is not clear at present is whether the identification of an obligatory filler induces a special mode of processing where the filler, if you will, actively searches for a gap, as suggested by Frazier (Note 6).

The processing classification of gap types is just being worked out now and is far from securely established, even if we restrict our attention to a single language such as English. Most studies of filler-gap processing (e.g. all those cited here) have restricted their attention to gaps of a very limited type (those considered to be wh-trace in Chomsky, 1981, obligatorily controlled null pronominals and most recently NP-trace, cf. Bever & McElree, Note 2). It is quite clear that certain gap types (presumably those which can receive case vs. those which cannot) are distinguished even in very initial stages of processing. For example, Frazier and Clifton (in press) show that "thematically controlled' null pronominals (e.g. the PRO in [33]) are not subject to the recent filler strategy.

33. a. John lent the book to Mary PRO to read ____ on vacation (PRO = null pronominal)
    b. John borrowed the book from Mary PRO to read ____ on vacation.

Another gap type which is not well understood is the parasitic gap (cf. Chomsky, 1982; Engdahl, 1983). The processing of parasitic gaps is only beginning to be studied empirically. Unlike ordinary gaps, parasitic gaps may only occur in sentences containing some independent filler-gap relation, as illustrated in (34b), where the parasitic gap is labelled "pg."

34. a. *Which film did you discuss the screenplay without seeing ____?
    b. Which film did you discuss ____ without seeing ____$_{pg}$?
    c. Which film did you discuss ____ without seeing it?

Parasitic-gap sentences exhibit a huge range of acceptability, both within and

[4] In the case of multiple filler-gap relations, the recent filler strategy automatically results in nested rather than intersecting dependencies, as illustrated

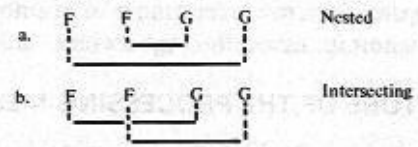

With noun phrase filler-gap relations, nested dependencies do seem to be easier across a variety of languages. However, when it comes to verbal dependencies in "verb raising" structures, it appears that intersecting dependencies are easier than nested ones (cf. Bach, Brown, & Marslen-Wilson, in press). This strongly suggests that perceivers deal differently with structures missing an argument—such as the filler-gap relations discussed in the text—than with structures in which the argument structure is determined by the item that has been displaced.

across speakers. For some speakers (like myself), they range from being completely acceptable to being completely unacceptable, depending on the precise details of the structure, lexical times, length, and meaning of the sentence. For other speakers, even the best examples of parasitic-gap sentences seem to be only marginally acceptable. Nevertheless, in a recent visual grammaticality judgement study by Seely (Note 10), it was found that parasitic gaps in clauses containing null pronominal subjects (e.g. the "without" clause in [34]) were considered just as acceptable as pronouns, as indicated by the error rate, despite the fact that huge acceptability differences between pronoun and parasitic gaps were obtained for sentences containing overt subjects.

In short, though many theoretically crucial questions remain concerning the gap-filling routines employed in English, at least certain basic generalisations seem to have been established. Gaps are recognised and assigned a filler "on-line" (i.e. the processor does not delay such decisions until the ends of clauses or sentences but rather makes them at or immediately following the position of the gap). In cases of temporary ambiguity, the more recent of two potential (grammatically permissible) fillers is assigned to a gap. Certain gap types are distinguished from each other even in initial parsing, i.e. all gaps are not treated equal even in the earliest stages of parsing. Thus, as in the case of constituent structure parsing, the identification and assignment of fillers and gaps appear to occur essentially immediately, following a depth-first (one analysis at a time) strategy. To the extent that the recent filler strategy results because recently encountered phrases may be retrieved from memory more quickly than distant phrases (see Frazier et al., 1983), there may be another similarity between structure building and gap filling. The decision principle of adopting the first available analysis may turn out to guide the analysis of both constituent structure relations and long-distance (filler-gap) relations.

In addition to long-distance relations involving arguments, anaphoric relations involving verbs and verb phrases have recently been investigated. The focus of attention has been on the nature of the antedecent for verbal gaps (see Black, Coltheart, & Byng, this volume) and for various types of verb-phrase anaphors (see Tanenhaus & Carlson, 1986 and references therein).

THE STRUCTURE OF THE PROCESSING MECHANISM

Modularity of the Language Processor

In preceding sections, the emphasis has been on the processing principles implicated in the syntactic processing of sentences. Little has been said about either the structure of the processing mechanism or the principles governing

other aspects of language processing. Fodor (1983) has argued for a modular conception of mind, in which the language system, like the various perceptual systems (e.g. the visual system), is viewed as an informationally encapsulated input system, governed by biological principles. On this conception, the early structuring of an input from some natural domain (e.g. language) is accomplished very quickly by mandatory application of specialised routines which deal exclusively with inputs from that domain, producing a "shallow output" which is then open to further interpretation and evaluation by a general cognitive system with full access to general knowledge and beliefs about the world.

The proposals discussed here are consistent in spirit but not in detail with Fodor's modular conception of mind. In particular, the thematic processor (of Rayner et al., 1983, see earlier) implies the existence of an interface system which operates in tandem with the language input system proper. This permits knowledge of the general cognitive variety to influence an input system indirectly, once it has been translated into the specialised vocabulary of the input system. Specifically, in the case at hand, pragmatic plausibility considerations influence thematic assignments, through thematic frame selection. Given our assumption that thematic frames are prestored and contain syntactic category labels, the selection of a thematic frame may influence ongoing syntactic analysis (though not the selection of an initial syntactic analysis).

I have emphasised the thematic processor hypothesis not because there is overwhelming evidence in favour of the hypothesis, nor even a detailed theory of the mechanisms implicated by the hypothesis: Rather, the reason for focusing on this hypothesis is conceptual; anyone favouring more extensive interaction between the language processing system (or more accurately, the grammatical processing system) and nonlinguistic processing system(s) must address the problem of how this interaction could take place. One may not simply assume communication between a nonlinguistic and linguistic system if there is no shared vocabulary. In brief, one must say how any fact I may know about the expected properties of real-world objects or likely events in the world could influence my grammatical decisions (like where to attach a phrase of a particular syntactic category), if one assumes that nongrammatical knowledge influences grammatical processing decisions.

Modularity Within the Language Processor

If the language processing system were one monolithic system, rather than a series of largely independent subsystems, we might expect to be able to identify the information source(s) and decision principles or mode of operation (multiple analysis vs. single analysis, immediate commitment vs.

delay) of the language processor once and for all. In other words, we would be able to make statements like: The language processor operates according to a depth-first strategy, making immediate commitments, after checking all grammatical information sources. But this is not what we find. Instead, it seems that we must restrict such statements, giving one specification for the lexical access subsystem, another for the constituent structure subsystem, and yet another for the thematic subsystem. In short, it appears that the information sources and the decision principles or mode of operation of the processor differs from one subsystem to the next. For example, considerable evidence suggests that lexical access involves multiple analysis (i.e. the processor considers all common representations of an input), using only information about the phonological or orthographic form of an input (i.e. ignoring, at this stage, information about syntactic or semantic compatibility with context) (cf. Seidenberg, Tanenhaus, Leiman, & Bienkowski, 1982; Swinney, 1979; Tanenhaus & Lucas, in press). Clearly this subsystem differs from the constituent structure system and from the gap-filling routines, which clearly do have access to syntactic information but do not follow a multiple analysis strategy.

While it is clear that the lexical-access system differs from the syntactic system, what is not completely clear at present is whether the syntactic processing system responsible for constituent structure analysis also accomplishes all other aspects of syntactic processing, e.g. the identification and assignment of filler-gap dependencies. Several investigators (Berwick & Weinberg, 1984; Clifton & Frazier, in press; Frazier, et al., 1983; Freedman & Forster, 1985) have proposed a distinction between a syntactic subsystem concerned with structure building (i.e. constituent structure) and a subsystem concerned with the identification or evaluation of (the binding and control) relations between phrases in the constituent structure representation. By contrast, working within a generalised phrase structure grammar framework (where filler-gap dependencies are captured by rules cast in the same format as those traditionally assumed for phrase structure dependencies that span only a single "layer" of the tree), Crain and Fodor (1985) propose that there is just a single syntactic processing system responsible for all syntactic decisions.

Whatever the outcome of this particular debate, one thing is reasonably clear. The nature of the debate in psycholinguistics has changed. The issue is no longer whether there is structure in the comprehension system, but rather the precise nature of the structure and the explanation or motivation for its existence.

CONCLUSIONS

At present, we don't have anything like a complete theory of language comprehension, even for a single language such as English. However, there are certain areas where psycholinguists have been able to ask very detailed questions about processing and to develop quite rich theories permitting a large range of disparate facts to be reduced to just one or two general principles. What is striking is how very simple the processing principles are. Once we pay careful attention to the representations involved, to the precise functional decomposition of the overall comprehension task, and to the particular information sources available for accomplishing some subtask, it seems we only need to appeal to rather trivial principles like "don't change current assignments without evidence" (implicit in the discussion in the second section) and "structure the input as soon as possible" (cf. minimal attachment, late closure, recent filler). To my mind, the simplicity of these principles is the most impressive evidence possible for the current modular approach to language and mind. After all, stated in its most general form, a modular theory is simply one which explains some extremely complicated phenomena in terms of the interaction of several simple subsystems.

ACKNOWLEDGEMENTS

I would like to thank Maria Black, Max Coltheart, and one anonymous reviewer for helpful comments on this chapter. The work reported here was supported by NSF Grant BNS 85 10177 to Rayner and Frazier and NIH Grant HD-18788 to Clifton and Frazier.

REFERENCES

Bach, E., Brown, C., & Marslen-Wilson, W. (in press). Crossed and nested dependencies in German and Dutch: A psycholinguistics study. *Language and Cognitive Processes.*

Berwick, R. C., & Weinberg, A. (1984). *The grammatical basis of linguistic performance.* Cambridge, Mass.: M.I.T. Press.

Bever, T. G. (1970). The cognitive basis for linguistic structures. In J. R. Hayes (Ed.), *Cognition and the development of language.* New York: John Wiley & Sons.

Carlson, G. (1983). Logical form: Type of evidence. *Linguistic and Philosophy, 6,* 295–318.

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

Chomsky, N. (1982). Some concepts and consequences of the theory of government and binding *Linguistic Inquiry Monograph 6.*

Clark, H. & Sengul, C. J. (1979). In search of referents for nouns and pronouns. *Memory and Cognition, 7,* 35–41.

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

Clifton, C. & Frazier, L. (in press). The use of syntactic information in filling gaps. *Journal of Psycholinguistic Research.*

Crain, S. & Fodor, J. D. (1985). How can grammars help parsers? In D. Dowty, L. Kartunnen, & A. Zwicky (Eds.), *Theoretical perspectives on natural language parsing*. Cambridge: Cambridge University Press.

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, & H. Zwicky (Eds.), *Natural language parsing*. Cambridge: Cambridge University Press.

Dell. G., McKoon, G., & Ratcliff, R. (1983). The activation of antecedent information during the processing of anaphoric reference. *Journal of Verbal Learning and Verbal Behaviour, 22*, 121–133.

Ehrlich, K. & Rayner, K. (1983). Pronoun assignment and semantic integration during reading: Eye movements and immediacy of processing. *Journal of Verbal Learning and Verbal Behaviour, 22*, 75–87.

Engdahl, E. (1983). Parasitic gaps. *Linguistics and Philosophy, 6*, 5–34.

Ferreira, F. & Clifton, C. (1986). The independence of syntactic processing. *Journal of Memory and Language, 25*, 348–368.

Flores d'Arcais, F. B. (1982). Automatic syntactic computation in sentence comprehension. *Psychological Research, 44*, 231–242.

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

Fodor, J. D. (1978). Parsing strategies and constraints on transformations. *Linguistic Inquiry, 9*, 427–474.

Fodor, J. D. (1979). Superstrategy. In W. Cooper & E. C. T. Walker (Eds.), *Sentence processing*. Hillsdale, N.J.: Lawrence Erlbaum Associates Inc.

Ford, M., Bresnan, J., & Kaplan, R. (1983). A competence-based theory of syntactic closure. In J. Bresnan (Ed.), *The mental representation of grammatical relations*. Cambridge, Mass.: M.I.T. Press, 727–796.

Forster, K. (1979). Levels of processing and the structure of the language processor. In W. E. Cooper & E. C. T. Walker (Eds.), *Sentence processing*. Hillsdale, N.J.: Lawrence Erlbaum Associates Inc.

Frazier, L. (1985). Syntactic complexity. In D. Dowty, L. Karttunen, & H. Zwicky (Eds.), *Natural language parsing*. Cambridge: Cambridge University Press.

Frazier, L. & Clifton, C. (in press). Thematic relations in parsing. *University of Massachusetts Occasional Papers in Linguistics*.

Frazier, L., Clifton, C., & Randall, J. (1983). Filling gaps: Decision principles and structure in sentence comprehension. *Cognition, 13*, 187–222.

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 analysis of structurally ambiguous sentences. *Cognitive Psychology, 14*, 178–210.

Frazier, L. & Rayner, K. (in press). Parameterising the language processing system: Left- vs. right-branching within and across languages. In J. Hawkins (Ed.), *Explaining linguistic universals*. Oxford: Basil Blackwell Publishing Co.

Freedman, S. & Forster, K. (1985). The psychological status of overgenerated sentences. *Cognition, 19*, 101–132.

Garrod, S. & Stanford, A. J. (1985). On the real-time character of interpretation during reading. *Language and Cognitive Processes, 1*, 43–59.

Gazdar, G. (1981). Unbounded dependencies and co-ordinate structure. *Linguistic Inquiry, 12*, 155–184.

Gorrell, P. (1985). Natural language parsing and reanalysis. *Proceedings of NELS, 16*, 186–196.

Grosz, B. T., Joshi, A., & Weinstein, S. (1983). Providing a unified account of definite noun phrases in discourse. *Proceedings of the ACL, MIT*, 44–50.

Holmes, V. M., Kennedy, A., & Murray, W. (in press). Syntactic structure and the garden path. *Quarterly Journal of Experimental Psychology*.

Jackendoff, R. & Culicover, P. (1971). A reconsideration of dative movement. *Foundations of Language, 7*, 397–412.

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

Joshi, A. & Weinstein, S. (1981). Control of inference: Role of some aspects of discourse centring. *Proceeding of the International Joint Conference on Artificial Intelligence*. August, Vancouver, B.C.

Kamp, H. (1981). A theory of truth and semantic representation. In Groenendijk (Eds.), *Formal methods in the study of language, Vol. I*. Amsterdam: Mathematische Centrum.

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

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

Kintsch, W. & van Dijk, T. A. (1978). Toward a model of text comprehension and production. *Psychological Review, 85*, 363–394.

Marcus, M. (1980). *A theory of syntactic recognition for natural language*. Cambridge, Mass.: M.I.T. 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.

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 Behaviour, 22*, 358–374.

Rayner, K. & Frazier, L. (in press). Parsing temporarily ambiguous complements. *Quarterly Journal of Experimental Psychology*.

Riesbeck, C. & Schank, R. (1978). Comprehension by computer: Expectation-based analysis of sentences in context. In W. I. M. Levelt & G. B. Flores d'Arcais (Eds.), *Studies in the perception of language*. Chichester: Wiley.

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

Stowe, L. (1986). Parsing wh-constructions: Evidence for on-line gap location. *Language and Cognitive Processes, 1*, 27–245.

Swinney, D. (1979). Lexical access during sentence comprehension: (Re) Considerations of context effects. *Journal of Verbal Learning and Verbal Behaviour, 18*, 645–659.

Tanenhaus, M., Stowe, L., & Carlson, G. (1985). The interaction of lexical expectation and pragmatic information. *Proceedings of the Seventh Annual Cognitive Science Society Meetings*.

Tanenhaus, M. & Lucas, M. (in press). Context effects in lexical processing. In U. Frauenfelder & L. Tyler (Eds.), special issue of *Cognition*.

Wanner, E. & Maratsos, M. (1978). An ATN approach to comprehension. In M. Halle, J. Bresnan, & G. A. Miller (Eds.), *Linguistic theory and psychological reality*. Cambridge, Mass.: M.I.T. Press.

Williams, E. (1980). Predication. *Linguistics Inquiry, 11*, 203–238.

Woolford, E. (1986). The distribution of empty nodes in Navajo: A mapping approach. *Linguistic Inquiry, 17*, 301–330.

REFERENCE NOTES

1. Altman, G. (1985). *Modularity and interaction in sentence processing*. Paper presented at Cognitive Science Meeting, Amherst, Mass., June 1985.
2. Bever, T. G. & McElree, B. (1986). *Empty categories access their antecedents during comprehension*. University of Rochester manuscript.
3. Carlson, G. & Tanenhaus, M. (1986). *Thematic roles and language comprehension*. University of Iowa manuscript.
4. Engdahl, E. (1981). *Interpreting sentences with multiple filler-gap dependencies*. Max-Planck-Institut for Psycholinguistik, Nijmegen, manuscript.
5. Fodor, J. D. *Freedom of expression*. University of Connecticut manuscript.
6. Frazier, L. Syntactic processing: Evidence from Dutch. Max-Planck Institute, Nijmegen, Holland, manuscript. (To appear in *Natural Language and Linguistic Theory*).
7. Frazier, L. (1979). *On comprehending sentences: Syntactic parsing strategies*. University of Connecticut doctoral dissertation.
8. Hudson, S., Tanenhaus, M. K., & Dell, G. S. (1986). *The effect of the discourse centre on the local coherence of a discourse*. Paper presented at Cognitive Science Society conference, Amherst, August 15–17.
9. Kurtzman, H. (1985). *Studies in syntactic ambiguity resolution*. M.I.T. doctoral dissertation.
10. Seely, T. Daniel. (in preparation). *Processing parasitic gaps: The dependence hypothesis*.
11. Stowe, L. *Thematic structure and sentence comprehension*. Manuscript.
12. Tanenhaus, M. & CArlson, G. (1986). *Processing verb phrase anaphors*. Paper presented at Annual Meeting of Cognitive Science Society, Amherst, Massachusets, August 1986.
13. Ueda, M. (1984). *Notes on parsing in Japanese*. University of Massachusetts manuscript.