

Matthew J. Traxler, Martin J. Pickering, Charles Clifton Jr., and
Roger van Gompel

ARCHITECTURES AND MECHANISMS FOR SENTENCE PROCESSING: IS SYNTACTIC PARSING A FORM OF LEXICAL AMBIGUITY RESOLUTION?

Some accounts of syntactic parsing propose that readers and listeners determine a sentence's syntactic structure in much the same way that they determine the meaning of a word (MacDonald, Pearlmutter, & Seidenberg, 1994). Namely, they recover information from a mental lexicon that describes what syntactic structures are permissible and preferred. MacDonald et al. note (p. 692), "...recent types of theorizing eliminate the strong distinction between accessing a meaning and constructing a syntactic representation, an idea which was central to previous accounts." Other accounts of syntactic parsing propose that readers and listeners consult grammatical principles and guide their syntactic structure-building decisions by determining (at least) the lexical categories of the words in the sentence (e.g., Frazier, 1979, 1987; Frazier & Clifton, 1996). In this chapter, we will contrast these different accounts by describing three dimensions on which they differ and by examining their compatibility with the empirical record.

THREE PARSERS

Because natural language is full of ambiguity, readers and listeners are constantly compelled either to delay analyzing some aspects of the input stream or entertain syntactic analyses that may turn out to be wrong. Empirical evidence suggests that

syntactic processing is rarely delayed (e.g., Altmann & Steedman, 1988; Frazier & Rayner, 1982; Trueswell, Tanenhaus, & Garnsey, 1994). Thus, two general classes of models are viable: those that propose that only a single analysis is ever considered at any given time and those that propose that multiple analyses are constructed and ranked. These two kinds of model differ as to how serial versus how parallel processing is presumed to be, as to whether the parser is limited to particular sources of information during an initial stage of processing, and as to the criteria the parser uses to rank analyses (or decide which single analysis to construct).

The *garden-path* model (Frazier, 1979, 1987) falls at the serial end of the serial-parallel dimension, as it presumes that readers and listeners construct only a single syntactic analysis at a time. Further, the garden-path parser ignores some potentially useful information during an initial stage of parsing. For example, referential context may influence the plausibility of alternative syntactic analyses (see, e.g., Altmann & Steedman, 1988), but according to the garden-path model, the parser does not use this information until after it has built an initial syntactic structure. The parser bases its preferences on *minimal attachment* and *late closure* strategies. Minimal attachment posits that readers and listeners minimize analysis time by building the phrase structure tree with the fewest nodes that is consistent with the input. Late closure posits that readers and listeners conserve processing resources by attaching incoming material to the most recently constructed phrase or clause rather than moving to a new or previously constructed phrase or clause.

The *construal* hypothesis (Frazier & Clifton, 1996) differs from the garden path model in that it resolves some (but not all) syntactic ambiguities in parallel. Construal proposes that text can be divided into *primary relations* and *non-primary* relations.

Frazier and Clifton (p. 41) propose, "Primary phrases and relations include a. the subject and predicate of any (+ or -) finite clause; b. complements and obligatory constituents of primary phrases." For primary relations, the construal parser acts just like the garden path parser. Text that instantiates primary relations is attached deterministically in a phrase-structure tree and is subject to minimal attachment and late closure. Text that instantiates non-primary relations is associated to the currently active theta domain (the entire extended maximal projection of the most recent theta assigner), rather than attached to a specific node in the phrase-structure tree. Once the non-primary phrase has been associated to the active theta domain, all attachment sites within the domain are assessed as hosts for the non-primary phrase simultaneously. Non-structural factors like semantic plausibility and referential context can inform the assessment of alternative attachments. Thus, construal proposes that the parser is a partially parallel mechanism and is only sometimes prevented from using non-syntactic information to construct an initial syntactic analysis. Its preferences for text instantiating primary phrases are the same as for the garden-path parser (i.e., it still prefers to construct the simplest phrase-structure tree), but simplicity does not determine its preferences for text instantiating non-primary phrases.

Constraint-based parsers differ from garden-path and construal parsers in that they consider and rank multiple syntactic analyses for all ambiguous fragments (e.g., Boland, 1997; Boland, Tanenhaus, Garnsey, & Carlson, 1995; MacDonald et al., 1994; Spivey-Knowlton & Sedivy, 1995; Trueswell, Tanenhaus, & Kello, 1993). They draw no substantive distinctions between initial and later stages of processing (or between initial analysis and re-analysis), they do not ignore potentially useful sources of information (e.g.,

referential context), and they prefer likely analyses (as opposed to simple analyses). Finally, constraint-based theories propose that ambiguity resolution results from competition between alternative syntactic analyses. Spivey-Knowlton and Sedivy propose (p. 260), "Near equal activation of the two alternatives will result in lengthy competition, hence greatly slowed reading times at the point of ambiguity."

A constraint-based parser determines which analysis to foreground by accessing information about the syntactic structures that a particular lexical item permits and prefers. For example, the verb *realized* can take a simple noun-phrase object as its post-verbal argument (as in *The investor realized a profit.*), but it may also take a full sentence complement (as in *The investor realized that a profit would save the company.*), or a reduced complement (as in *The investor realized a profit would save the company*). Further, the verb *realized* most commonly takes a reduced complement as its post verbal argument (Pickering, Traxler, & Crocker, 1999), whereas a verb like *doubted* most commonly takes a noun-phrase object as its post-verbal argument. When parsing syntactically ambiguous sentence fragments like *The man realized a profit...* the parser may take into account the fact that the verb *realized* or some larger class of verbs including *realized* prefers to take a reduced complement and therefore assign more activation to the sentence-complement interpretation and less to the noun-phrase object interpretation. In contrast, when the parser encounters the fragment *The man doubted the explanation...(would satisfy the police.)* it may take into account the fact that the verb *doubted* prefers to take a noun-phrase object and assign more activation to the noun-phrase object interpretation and less to the sentence-complement interpretation (see, e.g., Ford, Bresnan, & Caplan, 1983). Further, the amount of activation assigned to the alternative syntactic analyses will depend on how strongly the

individual verb prefers one structure over another. If a verb greatly prefers sentence complements, information available from the sentence fragment and preceding context supports this interpretation, and if the actual sentence requires this interpretation, then processing should proceed smoothly. If the verb has no strong preference, and context does not strongly bias one way or the other, then alternative analyses will receive similar amounts of activation and processing will be slowed or disrupted by competition between alternative syntactic analyses.

EXPERIMENTAL EVIDENCE FOR PARALLELISM, LIKELIHOOD PREFERENCES, AND COMPETITION

Evidence for parallel consideration of alternative interpretations in language processing is most prevalent in experiments on lexical ambiguity resolution. Swinney's (1979) classic study demonstrated that readers and listeners activate multiple meaning for ambiguous words before selecting a contextually appropriate meaning. Although some experimental work suggests that strongly biasing discourse context may cause readers and listeners to activate only the contextually appropriate meaning (e.g., Seidenberg, Tanenhaus, Leiman, & Bienkowski, 1982), consideration of alternative interpretations in parallel seems to occur commonly in lexical processing.

Lexical ambiguity resolution also offers some evidence for competitive resolution of ambiguity. In one eye-movement monitoring study (Rayner & Duffy, 1986), fixation times were collected for ambiguous words that had either one frequently occurring and one infrequently occurring meaning (biased words) or two equally frequent meanings (equibiased words). Equibiased words were defined as having no more than a 67% preference for the more frequent meaning, and biased words as having no less than 78% preference. Compared to unambiguous control words, readers spent more time fixating equibiased

words. By contrast, biased words produced fixation times equal to unambiguous control words. Rayner and Duffy proposed that the two meanings of equibiased words became available at the same time, thereby forcing the reader to do additional processing to select one meaning for integration into the discourse representation. A competition between alternative meanings is one possible description of this additional processing (MacDonald, Pearlmutter, & Seidenberg, 1994).

A further study investigated context effects on the processing of equibiased and biased ambiguous words (Duffy, Morris, & Rayner, 1988). Equibiased words were preceded by context that made one of two meanings appropriate. Biased words were preceded by context that made the subordinate meaning appropriate. Fixation times on equibiased words preceded by biasing context did not differ from unambiguous control words. By contrast, readers spent more time fixating biased ambiguous words than unambiguous words when context supported the less-frequent meaning of the ambiguous word. Duffy et al. provide two possible explanations for these findings. The first proposes that word meanings are delivered for integration serially. On this view, frequency and context interact to determine when different meanings become available. The other explanation proposes that all of the meanings of an ambiguous word are activated simultaneously, but with different amounts of evidence favoring the different meanings. According to Duffy et al. (p. 441), "Selection (of a single meaning for integration) would be relatively fast when the evidence clearly supported one meaning; it would be slow when equal amounts of evidence accrued for both meanings." Where discourse context supports the dominant meaning of an ambiguous word, lexical preferences and contextual information are in accord and the dominant meaning quickly out-competes the non-dominant meaning. When discourse context supports the non-dominant

meaning, lexical and contextual information conflict, and it takes longer for the contextually appropriate non-dominant meaning to out-compete the inappropriate dominant meaning. Thus, the data from this study may offer additional evidence that, in some circumstances, lexical ambiguity is resolved by a competitive process.

Evidence for parallel consideration of alternative syntactic structures comes from studies that manipulate the lexical preferences of individual verbs (Garnsey, Pearlmutter, Myers, & Lotocky, 1997; Spivey-Knowlton & Sedivy, 1995; Trueswell et al. 1993, 1994). Trueswell et al. (1993) monitored readers' eye-movements while they processed sentences like 1a - d:

(1a) The student forgot the solution was in the back of the book.

(1b) The student hoped the solution was in the back of the book.

(1c) The student forgot that the solution was in the back of the book.

(1d) The student hoped that the solution was in the back of the book.

The verb *forgot* prefers to take a noun-phrase object (as opposed to a reduced complement) as its post-verbal argument; whereas the verb *hoped* prefers to take a sentence-complement. In all of the test sentences, the correct interpretation was the sentence-complement analysis. Sometimes the post-verbal argument was a reduced complement (as in 1a & b), sometimes it was a full complement (as in 1c & d). Sentences with full complements provided baseline data, as the complementizer makes the direct-object interpretation impossible. If readers referred to detailed lexical information to rank alternative analyses, then they should have experienced greater disruption processing sentences like 1a, where the verb prefers the inappropriate noun-phrase object

interpretation, than sentences like 1b, where the verb prefers the sentence-complement interpretation that actually appears. In the event, readers showed greater disruption while processing the syntactically disambiguating material was in the in sentence 1a than in 1b, as predicted. Trueswell et al. also noted that the magnitude of the difference between the ambiguous forms (1a & b) and the unambiguous forms (1c & d) correlated with the degree to which the verbs preferred to take a full (vs. reduced) sentence complement. Thus, they concluded that the verbs activated alternative syntactic structures to differing degrees, and that resolution time was a function of how much evidence was available for alternative syntactic analyses. Finally, since difficulty at disambiguation was a function of whether the disambiguating material confirmed or disconfirmed the verbs' preferences, these studies also provide some evidence that the parser initially foregrounds (or adopts) the analysis that is most likely in the circumstances. A further eye-movement monitoring study produced compatible results using a semantic plausibility manipulation (Garnsey, Pearlmutter, Myers, & Lotocky, 1997; but see also Ferreira & Henderson, 1990; Pickering, Traxler, & Crocker, 1999).

RESEARCH ON MODIFIER ATTACHMENT IN 6 1/2 LANGUAGES

Processing of modifiers has been studied in Dutch, (American & British) English, French, German, Italian, and Spanish, which makes a discussion of this work especially appropriate for this volume. In this section, we will describe some of the work on modifier attachment in languages other than English and some of the work that has been done in the Eye-Tracking Lab at the University of Glasgow (Traxler, Pickering, & Clifton, 1998; van Gompel, Pickering, & Traxler, 1999). We will focus

