

DETERMINISTIC PARSING AND UNBOUNDED DEPENDENCIES

Ted Briscoe

Dept of Linguistics, Lancaster University
Bailrigg, Lancashire
LA1 4YT, UK

ABSTRACT

This paper assesses two new approaches to deterministic parsing with respect to the analysis of unbounded dependencies (UDs). UD in English are highly locally (and often globally) ambiguous. Several researchers have argued that the difficulty of UD undermines the programme of deterministic parsing. However, their conclusion is based on critiques of various versions of the Marcus parser which represents only one of many possible approaches to deterministic parsing. We examine the predictions made by a LR(1) deterministic parser and the Lexicat deterministic parser concerning the analysis of UD. The LR(1) technique is powerful enough to resolve the local ambiguities we examine. However, the Lexicat model provides a more psychologically plausible account of the parsing of UD, which also offers a unified account of the resolution of local and global ambiguities in these constructions.

INTRODUCTION

Church (1980:117) and Johnson-Laird (1983:313) have argued that the high degree of ambiguity in unbounded dependencies undermines the programme of deterministic parsing. Their conclusion is based on critiques of various versions of the Marcus parser (Marcus, 1980; Berwick & Weinberg, 1984). This parser represents only one of many possible approaches to deterministic parsing. Therefore, the conclusion that deterministic parsing, in general, is impractical or psychologically implausible may be premature.

In the next section, we outline the problems for the deterministic analysis of unbounded dependencies. In the succeeding sections, we present two alternative parsing techniques (and associated grammars) which make differing predictions concerning the onset and location of indeterminacy in the analysis of unbounded dependencies. We argue that the LR(1) parser is capable of deterministically resolving the local ambiguities which occur in these constructions, whilst the Lexicat parser is not. In the final section, we evaluate these predictions in the light of the Determinism Hypothesis (Marcus, 1980) and the Interactive Determinism Hypothesis (Briscoe & Boguraev, 1984; Briscoe, in press) and argue that the Lexicat parser in conjunction with the Interactive Determinism Hypothesis provides the most psychologically plausible and unified account of the parsing of unbounded dependencies.

UNBOUNDED DEPENDENCY AMBIGUITIES

Unbounded dependencies are found in English constituent questions, relative clauses and topicalised constructions. The dependency is between the preposed constituent and its point of attachment. For example, in (1) *Who* is preposed and functioning as direct object of the transitive verb *like*.

(1) Who do you like e

Most current theories of grammar represent the grammatical role of the preposed constituent by associating it with the normal position of a constituent having that grammatical role. In several theories, this position is occupied by a phonologically null category or trace which is grammatically linked to the preposed constituent. We will use e to mark this position because each of the grammars associated with the parsers we consider adopts a 'positional' account of the recovery of the grammatical role of the preposed constituent. However, we use e to mark an unambiguous point of attachment without any commitment to the presence of phonologically null categories or structure in the syntactic representation of unbounded dependencies.

The dependency between preposed constituent and point of attachment is unbounded because an unlimited amount of lexical material can occur between these two points in grammatical English constructions of this type. For example, it is possible to construct more and more embedded examples like those in (2) which exhibit the same dependency as (1).

(2)

Who do you think Kim likes e

Who do you expect that Kim hopes Sandy likes e

A parser for English capable of producing a syntactic representation adequate to guide semantic interpretation must recover the grammatical role of the preposed constituent. However, whenever these constructions contain verbs of ambiguous valency the correct point of attachment for the preposed constituent also becomes ambiguous. For example, in (3) there are two potential attachment points, or doubtful gaps (Fodor, 1979), written e?.

(3) Who do you want e? to succeed e?

The correct attachment of *Who* is ambiguous because both *want* and *succeed* can take, but do not require, NP objects.

The ambiguity in (3) is global with respect to the sentence; however, identical local ambiguities exist for a parser operating incrementally from left to right. For example, in (4) the attachment of *Who* as object of *want*, although correct, remains doubtful until the end of the sentence.

(4) Who do you want e? to succeed Bill

Thus, at the point when the parser reaches a potential but ambiguous attachment point in the left-to-right analysis of the input, it cannot be sure that this is the correct attachment point because there may be another further downstream in the input, as in (3). Moreover, the point of attachment further downstream may be unambiguous and obligatory, resolving the local ambiguity in the other direction, as in (5).

(5) Who do you want e? to replace e

To resolve the local ambiguities in unbounded dependencies the parser requires access to an unbounded amount of left and right context, measured in terms of lexical material. Firstly, when a potential attachment point is found, the parser must know whether or not a preposed constituent exists to be attached. This requires potentially unbounded access to the left context of the analysis since the preposed constituent could have occurred an unbounded distance back from its point of attachment. Secondly, when a potential but ambiguous attachment point is found, the parser must decide whether it is the correct point of attachment. However, since this decision cannot be made determinately when the potential attachment point occurs, the parser requires access to the right context of forthcoming material downstream from the current position. The examples in (6) illustrate that this would require unbounded lookahead.

(6) Who does Kim want e? to think that the boss will replace Sandy (with e)

Who does Kim want e? to think that the boss expects the directors to replace Sandy (with e)

In (6) the correct point of attachment cannot be determined until the end of the sentence which can be arbitrarily far away (in terms of lexical material) in the right context.

Berwick & Weinberg (1984:153f) argue that the Marcus parser can adequately represent an unbounded left context with finite resources if a successive cyclic and trace theoretic analysis (eg. Chomsky, 1981) of unbounded dependencies is adopted. However, both Church (1980) and Fodor (1985) demonstrate that the three cell lookahead buffer in the Marcus parser is not powerful enough to provide the required access to the right context in order to choose the correct point of attachment deterministically in many unbounded dependency constructions.

Marcus' (1980) Determinism Hypothesis claims that local ambiguities which are not resolvable in terms of the lookahead buffer are resolved by parsing strategy and that therefore, many unbounded dependency constructions

should be psychologically complex, 'garden paths' requiring extensive reanalysis. There is some evidence for syntactic preferences in the analysis of unbounded dependencies; the oddity of the examples in (7), which all require attachment of the preposed phrase in a doubtful position, suggests that the human parser prefers to ignore doubtful points of attachment and wait for a later one.

(7)

- a) Who did you want to give the present to Sue?
- b) I gave the boy who you wanted to give the books to three books.
- c) Sue wouldn't give the man who was reading the book

However, as Fodor (1985) points out, the great majority of unbounded dependency constructions are not complex. Furthermore, the Marcus parser predicts a sharp distinction between 'short' unbounded dependencies which fall within the buffer and the remainder. No such distinction seems to be supported by the psychological data. Finally, unbounded dependencies exhibit an identical kind of ambiguity which can be either local or global. Therefore, we would expect a unified account of their resolution, but the Determinism Hypothesis offers no account of the resolution of global ambiguities (see eg. Briscoe, in press).

ALTERNATIVE DETERMINISTIC PARSERS

There are a number of alternative deterministic parsing techniques many of which are in common use in compilers for high-level computer programming languages. Berwick (1985:313f) compares the Marcus parser to the most general of these techniques, LR(1) parsing (eg. Aho & Ullman, 1972), and argues that the Marcus parser can be seen as both an extension and restriction of the LR(1) technique. In fact, he argues that it is equivalent to a bounded context parser (eg. Aho & Ullman, 1972) which only allows literal access to grammatical symbols in the c-command domain in the left context and to two grammatical symbols in the right context.

To date, little attention has been given to alternative deterministic techniques as models of natural language parsing in their own right, though. One exception is the work of Shieber (1983) and Pereira (1985), who have proposed that a simple extension of the LALR(1) technique can be used to model human natural language parsing strategies. The LALR(1) technique is a more efficient variant of the LR(1) technique. Since our implementation of the Shieber/Pereira model uses the latter technique, we will refer throughout to LR(1). With the grammar discussed below, the behaviour of a parser using either technique should be identical (see Aho & Ullman, 1972). In addition, Briscoe & Boguraev (1984) and Briscoe (in press) propose that a bounded context, deterministic parser in conjunction with an extended categorial grammar will also model these strategies. Below these two alternative approaches are compared with respect to unbounded dependency constructions.

The Shieber/Pereira Parser

The LR(1) technique involves extensive preprocessing of the grammar used by the parser to compute all the distinct analysis 'paths' licensed by the grammar. This preprocessing results in a parse table which will deterministically specify the operations of a shift-reduce parser provided that the input grammar is an 'LR(1) grammar'; that is, provided that it is drawn from a subset of the unambiguous context-free grammars (see Aho & Ullman, 1972; Briscoe, in press). The parse table encodes the set of possible left contexts for an LR(1) grammar as a deterministic finite-state machine. In intuitive terms, the LR(1) technique is able to resolve deterministically a subset of the possible local ambiguities which can be represented in a context-free grammar, and none of the global ambiguities. If an LR(1) parsing table is constructed from a grammar covering a realistic, ambiguous fragment of English, the resulting non-deterministic parsing table will contain 'clashes' between shift and reduce operations and between different reduce operations. Shieber and Pereira demonstrate that if shift/reduce clashes are resolved in favour of shifting and reduce/reduce clashes in favour of reducing with the rule containing the most daughters, then the parser will model several psychologically plausible parsing strategies, such as right association (eg. Frazier, 1979).

Shieber (1983) and Pereira (1985) both provide grammars with a GPSG-style (Gazdar et al., 1985) SLASH feature analysis of unbounded dependencies. (8) presents a grammar fragment written in the same style to mimic the GPSG account of unbounded dependencies in a context-free notation.

(8)

Terminals

Det N Vtr V Aux want to wh \$

Non-terminals

SENT S VP VPinf VP/NP VPinf/NP NPwh NP

- 0) SENT → S \$
- 1) S → NP VP
- 2) NP → Det N
- 3) VP → Vtr NP
- 4) VP/NP → Vtr
- 5) NP → N
- 6) NPwh → wh
- 7) VP/NP → want VPinf/NP
- 8) VP/NP → want VPinf
- 9) VPinf → to VP
- 10) VP → V NP
- 11) VP → V
- 12) VPinf/NP → to VP/NP
- 13) VP/NP → V
- 14) S → NPwh Aux VP
- 15) S → NPwh Aux NP VP/NP
- 16) VP → want VPinf
- 17) VP → want NP VPinf
- 18) VP/NP → want NP VPinf/NP

The LR(1) technique applied to this grammar is very successful at resolving local ambiguities in these constructions; neither of the sentences in (9) result in any indeterminacy during parsing, despite the potential local ambiguity over the attachment of the preposed constituent.

(9)

- a) Who do you want Bill to succeed?
- b) Who do you want to succeed Bill?

That is, these local ambiguities fall within the subset of possible local ambiguities representable in a context-free grammar which are resolvable by this technique. On the other hand, parsing the globally ambiguous example in (3) using the same parse table derived from this grammar results in a reduce/reduce conflict, because the LR(1) technique cannot resolve global ambiguity (by definition). The conflict arises when the parser is in the configuration shown in Figure 1. At this point, the parser must choose between reducing *succeed* to VP or to VP/NP. When this indeterminacy arises the entire sentence has been shifted onto the parse stack.

Stack	Input Buffer
NPwh Aux NP want to V \$	
Who do you want to succeed	

Figure 1 - Configuration of LR(1) Parser

In general, because of the LR technique of preprocessing the grammar and despite the unbounded nature of the ambiguity, the decision point will always be at the end of the sentence. Therefore, local ambiguities involving the point of attachment of preposed constituents will not involve parsing indeterminacy using this technique. In this instance, the suspicion arises that the power of technique may be too great for a model of human parsing because examples such as those in (7) above do appear to be garden paths. However, normally such effects are only predicted when a parsing conflict is resolved incorrectly by the rules of resolution (eg. Shieber, 1983) and no conflict will arise parsing these examples with a grammar like that in (8).

At first sight it is surprising that these local ambiguities cause no problems since an LR(1) parser appears to have less access to the right context than the Marcus parser. However, the LR(1) parser makes greater use of the left context and also delays many syntactic decisions until most of the input is in the parse stack; in the configuration in Figure 1 no clause level attachments have been made, despite the fact that the complete sentence has been shifted into the parse stack.

The reduce/reduce conflict in the globally ambiguous case occurs much later than the position of the initial doubtful attachment point. Moreover, the conflict cannot be resolved using the Shieber/Pereira resolution rules as they stand, since both possible reductions (VP → V; VP/NP → V) only involve one daughter.

The Lexicat Parser

The LEXical-CATEGorical parser is a deterministic, shift-reduce parser developed for extended categorial grammars which include a rule of syntactic composition, as well as the more usual rule of application. An earlier version of the parser is briefly described in Briscoe & Boguraev (1984). Briscoe (in press) provides a complete description of Lexicat. Ades & Steedman (1982), Steedman (1985) and Briscoe (in press) discuss composition in further detail from the perspectives of syntax, semantics and parsing.

In a categorial grammar most syntactic information is located in the assignment of categories to lexical items. The rules of composition and application and a lexicon which suffices for the fragment under consideration are given in (10).

- (10)
- Function-Argument Application
 $X \ YIX \Rightarrow \ Y$
- Function-Function Composition
 $X'Y \ Y'Z \Rightarrow \ X'Z$
- Bill : NP to : VPinf/VP
 you : NP do : S/VP/NP
 who : NP succeed : VP!NP, VP
 want : VP/VPinf!NP, VP/VPinf

This grammar assigns the two analyses shown in Figure 2 to the ambiguous example (3).

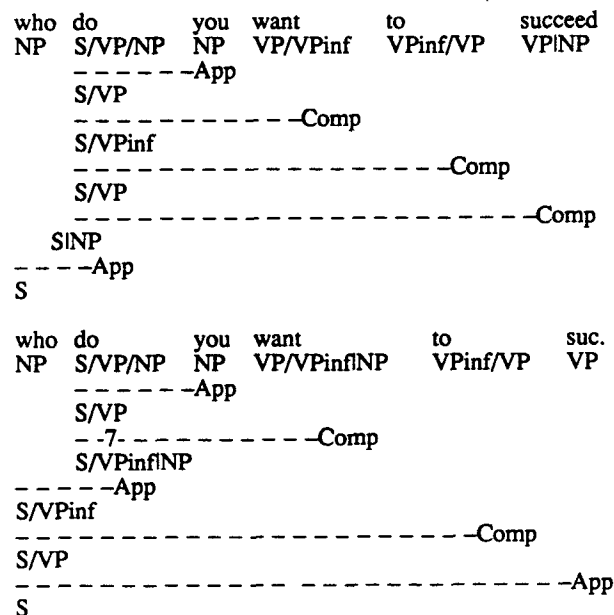


Figure 2 - Analysis of Unbounded Dependencies

The grammar represents the grammatical role of the preposed constituent by relating it directly to the verbal category. The material intervening between the preposed constituent and its point of attachment is composed into one (partial) constituent. Steedman (1985) provides linguistic motivation for a very similar analysis.

Lexicat employs one push down stack for analysis and storage of partially analysed material in the left and right context. Parsing proceeds on the basis of a three cell window into the stack. The item in the first cell (at the right hand end) of the stack represents a one word lookahead into the right context. This cell can only contain the lexical entry for the next word in the input. So, in common with LR(1) parsers but unlike the Marcus parser, Lexicat is restricted to lookahead of one lexical item. The second cell contains the syntactic category or categories associated with the current (partial) constituent under analysis. The third cell provides the left context for the parse and can contain the syntactic category or categories associated with the adjacent (partial) constituent to the left of the current constituent. Cells further down the stack contain (partial) constituents awaiting further integration into the analysis, but do not form part of the left context for parsing decisions.

Lexicat is a (1,1)-bounded context parser because it only has access to one set of grammatical symbols to the left and one set of grammatical symbols to the right of the current constituent (see Briscoe, in press). As such it is demonstrably less powerful than the LR(1) technique, which allows access to any aspect of the left context which can be represented as a regular expression, and the Marcus parser, which allows access to grammatical symbols in the c-command domain in the left context and two (not necessarily terminal) symbols in the right context (eg. Berwick, 1985:313f). However, it is unclear, at present, what precisely can be concluded on the basis of these differences in the parsing techniques because of the differing properties of the grammatical theories employed in each model.

The Lexicat parser does not employ a parse table or use parsing states to maintain information about the nature of the left context. Rules of application and composition (with various restrictions on the directionality of reduction and the range of categories to which each rule applies) are used directly by the parser to perform reductions. There are two stages to each step in the parsing algorithm; a checking phase and a reduction phase. A rule of forward application and a rule of forward composition are used to check for the possibility of reduction between the categories in Cell1 and Cell2. If reduction is possible, Lexicat shifts the next item from the input buffer onto the stack. If reduction is impossible, Lexicat moves to the reduction phase and attempts a reduction between Cell2 and Cell3 using rules of backward and forward application and a more constrained rule of forward composition. If this fails, then Lexicat shifts. This completes one step of the parsing algorithm, so Lexicat returns to the checking phase. This process continues until the parse is complete or the input is

exhausted, at which point the parser halts under the usual conditions.

Quite often, a shift/reduce conflict arises during the checking phase in which case Lexicat opts, by default, to shift. In most constructions this resolution rule results in analyses which conform to the parsing strategies of late closure and right association (eg. Frazier, 1979). However, in unbounded dependencies it results in late attachment of the preposed constituent. For example, if Lexicat is in the configuration shown in Figure 3, then a shift/reduce conflict will occur in the checking phase.

Stack	Input Buffer
3 2	1
NP S/VPinfNP S/VPinf	VPinf/VP succeed
Who (do you want)	to

Figure 3 - Configuration of Lexicat Parser

The ambiguity of the partial constituent *do you want* results from the ambiguous valency of *want*. Depending on which category in Cell2 is chosen, reduction by forward composition will or will not be possible. By default, Lexicat will shift in the face of this conflict; thus the potential for reduction by backwards application between Cell2 and Cell3 will not be considered during this step. In the next configuration, the preposed constituent will be in Cell4 outside the parser's 'window' into the stack. Therefore, the possibility of attaching the preposed constituent does not arise again until *do you want to succeed* has been composed into one (partial) constituent. At this point, the only possible attachment which remains is as object of *succeed*.

If the parser is analysing (9a), and *Bill* rather than *to* is the item in Cell1 in the configuration in Figure 3, then essentially the same situation obtains; there will be a shift/reduce conflict, shift will be chosen by default and the parser will go on to build the late attachment analysis. If, on the other hand, the parser is analysing (9b) and *Bill* is in the input buffer at the end of the sentence, the parse configuration at the moment of indeterminacy will still be as in Figure 3 and the same default analysis will be chosen since the parser has no access to the contents of the input buffer to guide its decisions. However, in this case the parse will fail because *Bill* will be attached as object of *succeed* and *Who* will be left dangling.

Unlike the LR(1) model, Lexicat faces parsing indeterminacy at the point when the first potential point of attachment occurs. The resolution rule in favour of shifting predicts that late attachment of preposed constituents is preferred and this prediction is compatible with the garden path data in (7). The Lexicat parser employs the grammar directly without preprocessing and therefore conforms to Berwick & Weinberg's (1984)

transparency condition.

INTERACTIVE DETERMINISM

Marcus' (1980) Determinism Hypothesis claims that local ambiguity is resolved autonomously either by lookahead or, if this fails, by parsing strategy. This predicts that strategy-violating local ambiguities which fall outside the span of the Marcus parser's lookahead buffer will be garden paths. The theory tells us little about the resolution of global ambiguities, but implies that the mechanism employed must be similar to that used to recover from garden paths, involving interaction with other components of the language comprehension system.

Using the Determinism Hypothesis, it is difficult to select between the two models outlined above (or indeed to conclusively rule out the Marcus parser) because the differing predictions concerning the onset of indeterminacy in the face of identical ambiguities are unimportant. The Determinism Hypothesis concerns only judgements of psychological complexity. These judgements marginally favour the Lexicat parser, but the data relating to garden paths with unbounded dependencies is hardly overwhelming. Moreover, the Determinism Hypothesis seems highly suspect in the light of the unbounded dependency examples because it predicts that local and global ambiguities are resolved using completely different mechanisms. At the very least, this approach is unparsimonious and leaves the resolution of global ambiguity largely unexplained. In addition, the extreme similarity between local and global ambiguities in unbounded dependency constructions suggests that one mechanism for the resolution of local and global ambiguity is quite feasible.

Briscoe & Boguraev (1984) and Briscoe (in press) propose a different account of the relationship between parsing and other components of comprehension than that entailed by the Determinism Hypothesis, dubbed the Interactive Determinism Hypothesis (IDH). Under this account of deterministic parsing, the parser proceeds autonomously until it is faced with an indeterminacy and then requests help from other components of the comprehension system. By default, the parser will apply a resolution rule or parsing strategy at such a point, but this can be overruled by specific non-syntactic information at the onset of the indeterminacy. The IDH implies that both local and global ambiguity is resolved at its onset (relative to some parsing technique) either by strategy or interactive blocking of the strategy. The IDH predicts, in addition, that garden paths will arise when a strategy-violating, local ambiguity is not resolved interactively, as a result of the absence or removal of the relevant non-syntactic information.

Under the IDH, the differing predictions concerning the onset of indeterminacy in ambiguous unbounded dependency constructions become crucial in any comparison of the two parsing models outlined above.

The Lexicat parser makes far stronger predictions because indeterminacy occurs much earlier in the analysis when less of the input is available in the left context. Since Lexicat prefers late attachment by default, it predicts that when a doubtful point of attachment is reached, which is the correct point of attachment, non-syntactic information in the available left context should block the preference to shift and force early reduction with the preposed constituent. By contrast, the Shieber/Pereira parser does not meet an indeterminacy except in globally ambiguous cases and then not until all the input is in the left context. It therefore predicts in conjunction with the IDH that there should be no garden paths involving unbounded dependencies and that there should be some non-syntactic information in the entire input which resolves the global ambiguity. The former prediction appears to be wrong in the light of the examples in (7) and the latter is so weak as to be trivial.

It turns out that there is some evidence supporting the far stronger predictions of the Lexicat model in conjunction with the IDH. This evidence comes from the distribution of prosodic boundaries in relation to the onset of strategy-violating syntactic ambiguities. For example in (11)

(11) Without her, contributions to the fund would be inadequate.

the comma (an orthographic counterpart to certain types of prosodic boundary) marks the location of an intonational or major tone group boundary which would normally occur in the spoken version of this sentence. The prosodic boundary prevents the potential misinterpretation in which *her contributions* is reduced as one NP. In unbounded dependency constructions, Danly (reported in Cooper & Paccia-Cooper, 1980:159f) has demonstrated that the final syllable of the verb preceding a correct attachment point is lengthened relative to an environment without a potential attachment point, or with a potential but incorrect one. Syllable lengthening is indicative of a phrasal or minor tone group boundary. Paul Warren and the author have since tested Danly's result by acoustically analysing ten readers' productions of four examples containing doubtful but correct early points of attachment and four similar examples with doubtful and incorrect early attachment points. The results tend to confirm the original finding since lengthening was found consistently (although the measurements did not achieve statistical significance; see Briscoe, in press). A final piece of evidence that lengthening occurs before a correct point of attachment comes from the acceptability of contraction in (12a), but not in (12b).

- (12)
- a) Who do you wanna succeed
 - b) *Who do you wanna succeed Bill

Contraction forces late attachment of *Who* in a), but b) is unacceptable because the only possible interpretation involves attachment 'into' the contracted form. Fodor (1979:277n17) notes that it is only the occurrence of contraction which appears to provide determinate information about the correct analysis and that, since

contraction is optional, this information cannot be relied on. However, metrical phonologists (eg. Nespor & Vogel, 1986) argue that such rules are not blocked syntactically by the presence of the trace/gap, but by an intervening prosodic boundary and that this explains the coincidence of other phonetic effects, such as lengthening, at points where contraction is blocked (Cooper & Paccia-Cooper, 1980:Ch10). In other words, contraction is the tip of a far more systematic prosodic iceberg which does reliably cue the presence of a correct attachment point.

When Lexicat reaches a potential point of attachment, it is faced with a shift/reduce ambiguity. By default, Lexicat prefers to shift, but this strategy can be blocked by a prosodic boundary intervening between the verb and item about to be shifted into the parse stack. Therefore, the parser opts for early attachment of the preposed constituent. In terms of Lexicat's operation, the prosodic boundary in the unbounded dependency construction plays the same role as the prosodic boundary in (11); they both block the shift operation. By contrast, in the Shieber/Pereira parser it is difficult to see how a prosodic boundary in unbounded dependencies could be used to select one of two possible reductions, whilst in an example like (11) it would need to force the parser to shift rather than reduce. In addition, the relevant non-syntactic information occurs at the onset of the indeterminacy for the Lexicat model but well before this point for the Shieber/Pereira model. This corroborates the far stronger prediction made by Lexicat, and also makes the mechanism of interaction for this model simpler (see Briscoe, in press).

Finally, we should note that in the garden paths in (7) it is intuitively clear that examples b) and c) would be spoken with prosodic boundaries at the correct attachment point, and probably written with commas otherwise. Example a) on the other hand, is more subtle, but the experimental results reported above suggest that *want* would be lengthened in this context signalling the early attachment point. Thus, the IDH's prediction that garden paths are the result of the removal or distortion of non-syntactic information which functions to prevent the parser's default analysis in the face of indeterminacy is corroborated.

CONCLUSION

The paper has presented two approaches to the deterministic analysis of unbounded dependencies. The LR(1) technique is capable of resolving the type of local ambiguities which appear to occur in these constructions, suggesting that Church and Johnson-Laird were wrong to reject deterministic parsing on the basis of this data. However, we have argued that the Lexicat parser provides a better psychological model of the parsing of unbounded dependencies because a) it predicts the garden path data and b), in conjunction with the IDH, it predicts the apparent distribution of prosodic boundaries in these constructions more successfully, and c) it provides a unified account of the resolution of local and global

ambiguities, and d) it is a simpler model of deterministic parsing which does not require preprocessing the grammar or maintaining state information concerning the left context.

REFERENCES

- Ades, A. & Steedman, M. (1982). On the order of words. *Linguistics & Philosophy*, 4, 517-558.
- Aho, A. & Ullman, J. (1972). *The Theory of Parsing, Translating and Compiling*. Vol. I, Englewood Cliffs, NJ: Prentice-Hall.
- Berwick, R. (1985). *The Acquisition of Syntactic Knowledge*. Cambridge, Mass.: MIT Press.
- Berwick, R. & Weinberg, A. (1984). *The Grammatical Basis of Linguistic Performance*. Cambridge, Mass.: MIT Press.
- Briscoe, E. (In press). *Modelling Human Speech Comprehension; A Computational Approach*. Chichester, UK: Ellis Horwood.
- Briscoe, E. & Boguraev, B. (1984). Control structures and theories of interaction in speech understanding systems. In *Proc. of Coling84*, Stanford, Ca, pp. 259-266.
- Chomsky, N. (1981). *Lectures on Government and Binding*. Dordrecht, Holland: Foris.
- Church, K. (1980). *On Memory Limitations in Natural Language Processing*. Bloomington, Ind.: Indiana University Linguistics Club.
- Cooper, W. & Paccia-Cooper, J. (1980). *Syntax & Speech*. Cambridge, Mass.: Harvard University Press.
- Fodor, J.D. (1979). Superstrategy. In Walker, E. & Cooper, W. (eds.) *Sentence Processing*, Hillsdale, NJ: Lawrence Erlbaum.
- Fodor, J.D. (1985). Deterministic parsing and subadjacency. *Language and Cognitive Processes*, 1.1, 3-42.
- Frazier, L. (1979). *On Comprehending Sentences: Syntactic Parsing Strategies*. Bloomington, Ind.: Indiana University Linguistics Club.
- Gazdar, G., Klein, E., Pullum, G., & Sag, I. (1985). *Generalized Phrase Structure Grammar*. Oxford, UK: Blackwell.
- Johnson-Laird, P. (1983). *Mental Models*. Cambridge, UK: CUP.
- Marcus, M. (1980). *A Theory of Syntactic Recognition for Natural Language*. Cambridge, Mass.: MIT Press.
- Nespor, M. & Vogel, I. (1986). *Prosodic Phonology*. Dordrecht, Holland: Foris.
- Pereira, F. (1985). A new characterisation of attachment preferences. In Dowty, D., Karttunen, L., & Zwicky, A. (eds.) *Natural Language Parsing*. Cambridge, UK: CUP.
- Shieber, S. (1983). Sentence disambiguation by a shift-reduce parsing technique. In *Proc. of 21st ACL*, Cambridge, Mass., pp. 113-118.
- Steedman, M. (1985). Dependency and coordination in the grammar of Dutch and English. *Language* 55, 523-68.