RELATING COMPLEXITY TO PRACTICAL PERFORMANCE IN PARSING WITH WIDE-COVERAGE UNIFICATION GRAMMARS

John Carroll

University of Cambridge, Computer Laboratory
Pembroke Street, Cambridge CB2 3QG, UK
jac@cl.cam.ac.uk

Abstract

The paper demonstrates that exponential complexities with respect to grammar size and input length have little impact on the performance of three unification-based parsing algorithms, using a wide-coverage grammar. The results imply that the study and optimisation of unification-based parsing must rely on empirical data until complexity theory can more accurately predict the practical behaviour of such parsers[1].

1. INTRODUCTION

General-purpose natural language (NL) analysis systems have recently started to use declarative unification-based sentence grammar formalisms; systems of this type include SRI's CLARE system (Alshawi *et al.*, 1992) and the Alvey NL Tools (ANLT; Briscoe *et al.*, 1987a). Using a declarative formalism helps ease the task of developing and maintaining the grammar (Kaplan, 1987). In addition to syntactic processing, the systems incorporate lexical, morphological, and semantic processing, and have been applied successfully to the analysis of naturally-occurring texts (e.g. Alshawi *et al.*, 1992; Briscoe & Carroll, 1993).

Evaluations of the grammars in these particular systems have shown them to have wide coverage (Alshawi *et al.*, 1992; Taylor, Grover & Briscoe, 1989)[2]. However, although the practical throughput of parsers with such realistic grammars is important, for example when process-

ing large amounts of text or in interactive applications, there is little published research that compares the performance of different parsing algorithms using wide-coverage unification-based grammars. Previous comparisons have either focussed on context-free (CF) or augmented CF parsing (Tomita, 1987; Billot & Lang, 1989), or have used relatively small, limited-coverage unification grammars and lexicons (Shann, 1989; Bouma & van Noord, 1993; Maxwell & Kaplan, 1993). It is not clear that these results scale up to reflect accurately the behaviour of parsers using realistic, complex unification-based grammars: in particular, with grammars admitting less ambiguity parse time will tend to increase more slowly with increasing input length, and also with smaller grammars rule application can be constrained tightly with relatively simple predictive techniques. Also, since none of these studies relate observed performance to that of other comparable parsing systems, implementational oversights may not be apparent and so be a confounding factor in any general conclusions made.

Other research directed towards improving the throughput of unification-based parsing systems has been concerned with the unification operation itself, which can consume up to 90% of parse time (e.g. Tomabechi, 1991) in systems using lexicalist grammar formalisms (e.g. HPSG; Pollard & Sag, 1987). However, parsing algorithms assume more importance for grammars having more substantial phrase structure components, such as CLARE (which although employing some HPSG-like analyses still contains several tens of rules) and the ANLT (which uses a formalism derived from GPSG; Gazdar *et al.*, 1985), since the more specific rule set can be used to control which unifications are performed.

In NL analysis, the syntactic information associated with lexical items makes top-down parsing less attractive than bottom-up (e.g. CKY; Kasami, 1965; Younger, 1967), although the latter is often augmented with top-down predic-

---

[2] For example, Taylor *et al.* demonstrate that the ANLT grammar is in principle able to analyse 96.8% of a corpus of 10,000 noun phrases taken from a variety of corpora.

tion to improve performance (e.g. Earley, 1970; Lang, 1974; Pratt, 1975). Section 2 describes three unification-based parsers which are related to polynomial-complexity bottom-up CF parsing algorithms. Although incorporating unification increases their complexity to exponential on grammar size and input length (section 3), this appears to have little impact on practical performance (section 4). Sections 5 and 6 discuss these findings and present conclusions.

2. THE PARSERS

The three parsers in this study are: a bottom-up left-corner parser, a (non-deterministic) LR parser, and an LR-like parser based on an algorithm devised by Schabes (1991). All three parsers accept grammars written in the ANLT formalism (Briscoe et al., 1987a), and the first two are distributed as part of the ANLT package. The parsers create parse forests (Tomita, 1987) that incorporate subtree sharing (in which identical sub-analyses are shared between differing superordinate analyses) and node packing (where sub-analyses covering the same portion of input whose root categories are in a subsumption relationship are merged into a single node).

THE BOTTOM-UP LEFT-CORNER PARSER

The bottom-up left-corner (BU-LC) parser operates left-to-right and breadth-first, storing partial (active) constituents in a chart; Carroll (1993) gives a full description. Although pure bottom-up parsing is not usually thought of as providing high performance, the actual implementation achieves very good throughput (see section 4) due to a number of significant optimisations, amongst which are:

- Efficient rule invocation from cheap (static) rule indexing, using discrimination trees keyed on the feature values in each rule's first daughter to interleave rule access with unification and also to share unification results across groups of rules.

- Dynamic indexing of partial and complete constituents on category types to avoid attempting unification or subsumption operations which static analysis shows will always fail.

- Dynamic storage minimisation, deferring structure copying—e.g. required by the unification operation or by constituent creation—until absolutely necessary (e.g. unification success or parse success, respectively).

The optimisations improve throughput by a factor of more than three.

THE NON-DETERMINISTIC LR PARSER

Briscoe & Carroll (1993) describe a methodology for constructing an LR parser for a unification-based grammar, in which a CF 'backbone' grammar is automatically constructed from the unification grammar, a parse table is constructed from the backbone grammar, and a parser is driven by the table and further controlled by unification of the 'residue' of features in the unification grammar that are not encoded in the backbone. In this parser, the LALR(1) technique (Aho, Sethi & Ullman, 1986) is used, in conjunction with a graph-structured stack (Tomita, 1987), adapting for unification-based parsing Kipps' (1989) Tomita-like recogniser that achieves polynomial complexity on input length through caching.

On each reduction the parser performs the unifications specified by the unification grammar version of the CF backbone rule being applied. This constitutes an on-line parsing algorithm. In the general case, the off-line variant (in which all unifications are deferred until the complete CF parse forest has been constructed) is not guaranteed to terminate; indeed, it usually does not do so with the ANLT grammar. However, a drawback to the on-line algorithm is that a variant of Kipps' caching cannot be used, since the cache must necessarily assume that all reductions at a given vertex with all rules with the same number of daughters build exactly the same constituent every time; in general this is not the case when the daughters are unification categories. A weaker kind of cache on partial analyses (and thus unification results) was found to be necessary in the implementation, though, to avoid duplication of unifications; this sped the parser up by a factor of about three, at little space cost.

THE COMPILED-EARLEY PARSER

The Compiled-Earley (CE) parser is based on a predictive chart-based CF parsing algorithm devised by Schabes (1991) which is driven by a table compiling out the predictive component of Earley's (1970) parser. The size of the table is related linearly to the size of the grammar (unlike the LR technique). Schabes demonstrates that this parser always takes fewer steps than Earley's, although its time complexity is the same: $O(n^3)$. The space complexity is also cubic, since the parser uses Earley's representation of parse forests.

The incorporation of unification into the CE parser follows the methodology developed for unification-based LR parsing described in the previous section: a table is computed from a CF 'backbone', and a parser, augmented with on-line unification and feature-based subsumption opera-

tions, is driven by the table. To allow meaningful comparison with the LR parser, the CE parser uses a one-word lookahead version of the table, constructed using a modified LALR technique (Carroll, 1993)[3].

To achieve the cubic time bound, the parser must be able to retrieve in unit time all items in the chart having a given state, and start and end position in the input string. However, the obvious array implementation, for say a ten word sentence with the ANLT grammar, would contain almost 500000 elements. For this reason, the implementation employs a sparse representation for the array, since only a small proportion of the elements are ever filled. In this parser, the same sort of duplication of unifications occurs as in the LR parser, so lists of partial analyses are cached in the same way.

3. COMPLEXITIES OF THE PARSERS

The two variables that determine a parser's computational complexity are the grammar and the input string (Barton, Berwick & Ristad, 1987). These are considered separately in the next two sections.

GRAMMAR-DEPENDENT COMPLEXITY

The term dependent on the grammar in the time complexity of the BU-LC unification-based parser described above is $O(|C|^2|R|^3)$, where $|C|$ is the number of categories implicit in the grammar, and $|R|$, the number of rules. The space complexity is dominated by the size of the parse forest, $O(|C|)$ (these results are proved by Carroll, 1993). For the ANLT grammar, in which features are nested to a maximum depth of two, $|C|$ is finite but nevertheless extremely large (Briscoe et al., 1987b)[4].

The grammar-dependent complexity of the LR parser makes it also appear intractable: Johnson (1989) shows that the number of LR(0) states for certain (pathological) grammars is exponentially related to the size of the grammar, and that there are some inputs which force an LR parser to visit all of these states in the course of a parse.

Thus the total number of operations performed, and also space consumed (by the vertices in the graph-structured stack), is an exponential function of the size of the grammar.

To avoid this complexity, the CE parser employs a table construction method which ensures that the number of states in the parse table is linearly related to the size of the grammar, resulting in the number of operations performed by the parser being at worst a polynomial function of grammar size.

INPUT-DEPENDENT COMPLEXITY

Although the complexity of returning all parses for a string is always related exponentially to its length (since the number of parses is exponential, and they must all at least be enumerated), the complexity of a parser is usually measured for the computation of a parse forest (unless extracting a single analysis from the forest is worse than linear)[5].

If one of the features of the ANLT grammar formalism, the kleene operator (allowing indefinite repetition of rule daughters), is disallowed, then the complexity of the BU-LC parser with respect to the length of the input string is $O(n^{p+1})$, where p is the maximum number of daughters in a rule (Carroll, 1993). The inclusion of the operator increases the complexity to exponential. To retain the polynomial time bound, new rules can be introduced to produce recursive tree structures instead of an iterated flat tree structure. However, when this technique is applied to the ANLT grammar the increased overheads in rule invocation and structure building actually slow the parser down.

Although the time and space complexities of CF versions of the LR and CE parsers are $O(n^3)$, the unification versions of these parsers both turn out to have time bounds that are greater than cubic, in the general case. The CF versions implicitly pack identical *sequences* of sub-analyses, and in all reductions at a given point with rules with the same number of daughters, the packed sequences can be formed into higher-level constituents as they stand without further processing. However, in the unification versions, on each reduce action the daughters of the rule involved have to be unified with every possible alternative sequence of the sub-analyses that are being consumed by the rule

---

[3]Schabes describes a table with no lookahead; the successful application of this technique supports Schabes' (1991:109) assertion that "several other methods (such as LR(k)-like and SLR(k)-like) can also be used for constructing the parsing tables [...]"

[4]Barton, Berwick & Ristad (1987:221) calculate that GPSG, also with a maximum nesting depth of two, licences more than 10^{775} distinct syntactic categories. The number of categories is actually infinite in grammars that use a fully recursive feature system.

[5]This complexity measure does correspond to real world usage of a parser, since practical systems can usually afford to extract only a small number of parses from the frequently very large number encoded in a forest; this is often done on the basis of preference-based or probabilistic factors (e.g. Carroll & Briscoe, 1992).

(in effect expanding and flattening out the packed sequences), leading to a bound of n^{p+1} on the total number of unifications.

4. PRACTICAL RESULTS

To assess the practical performance of the three unification-based parsers described above, a series of experiments were conducted using the ANLT grammar (Grover, Carroll & Briscoe, 1993), a wide-coverage grammar of English. The grammar is defined in metagrammatical formalism which is compiled into a unification-based 'object grammar'—a syntactic variant of the Definite Clause Grammar formalism (Pereira & Warren, 1980)—containing 84 features and 782 phrase structure rules. Parsing uses fixed-arity term unification. The grammar provides full coverage of the following constructions: declarative sentences, imperatives and questions (yes/no, tag and wh-questions); all unbounded dependency types (topicalisation, relativisation, wh-questions); a relatively exhaustive treatment of verb and adjective complement types; phrasal and prepositional verbs of many complement types; passivisation; verb phrase extraposition; sentence and verb phrase modification; noun phrase complements and pre- and post-modification; partitives; coordination of all major category types; and nominal and adjectival comparatives.

Although the grammar is linked to a lexicon containing definitions for 40000 base forms of words, the experiments draw on a much smaller lexicon of 600 words (consisting of closed class vocabulary and, for open-class vocabulary, definitions of just a sample of words which taken together exhibit the full range of possible complementation patterns), since issues of lexical coverage are of no concern here.

COMPARING THE PARSERS

In the first experiment, the ANLT grammar was loaded and a set of sentences was input to each of the three parsers. In order to provide an independent basis for comparison, the same sentences were also input to the SRI Core Language Engine (CLE) parser (Moore & Alshawi, 1992) with the CLARE2.5 grammar (Alshawi et al., 1992), a state-of-the-art system accessible to the author.

The sentences were taken from an initial sample of 175 representative sentences extracted from a corpus of approximately 1500 that form part of the ANLT package. This corpus, implicitly defining the types of construction the grammar is intended to cover, was written by the linguist who developed the ANLT grammar and is used to check for any adverse effects on coverage when the grammar is modified during grammar development. Of

Parser	Grammar	CPU time	Storage allocated
BU-LC	ANLT	75.5	47.0
LR	ANLT	48.9	33.6
CE	ANLT	98.4	38.5
CLE	CLARE2.5	277.7	–

Table 1: Parse times (in CPU seconds on a Sun Sparc ELC workstation) and storage allocated (in megabytes) while parsing the 129 test sentences (1–12 words in length).

the initial 175 sentences, the CLARE2.5 grammar failed to parse 42 (in several cases because punctuation is strictly required but is missing from the corpus). The ANLT grammar also failed to parse three of these, plus an additional four. These sentences were removed from the sample, leaving 129 (mean length 6.7 words) of which 47 were declarative sentences, 38 wh-questions and other sentences with gaps, 20 passives, and 24 sentences containing co-ordination.

Table 1 shows the total parse times and storage allocated for the BU-LC parser, the LR parser, and the CE parser, all with ANLT grammar and lexicon. All three parsers have been implemented by the author to a similar high standard: similar implementation techniques are used in all the parsers, the parsers share the same unification module, run in the same Lisp environment, have been compiled with the same optimisation settings, and have all been profiled with the same tools and hand-optimised to a similar extent. (Thus any difference in performance of more than around 15% is likely to stem from algorithmic rather than implementational reasons). Both of the predictive parsers employ one symbol of lookahead, incorporated into the parsing tables by the LALR technique. Table 1 also shows the results for the CLE parser with the CLARE2.5 grammar and lexicon. The figures include garbage collection time, and phrasal (where appropriate) processing, but not parse forest unpacking. Both grammars give a total of around 280 analyses at a similar level of detail.

The results show that the LR parser is approximately 35% faster than the BU-LC parser, and allocates about 30% less storage. The magnitude of the speed-up is less than might be expected, given the enthusiastic advocation of non-deterministic CF LR parsing for NL by some researchers (e.g. Tomita, 1987; Wright, Wrigley & Sharman, 1991), and in the light of improvements observed for predictive over pure bottom-up parsing (e.g. Moore & Dowding, 1991). However, on the assumption that incorrect prediction of gaps is

the main avoidable source of performance degradation (c.f. Moore & Dowding), further investigation shows that the speed-up is near the maximum that is possible with the ANLT grammar (around 50%).

The throughput of the CE parser is half that of the LR parser, and also less than that of the BU-LC parser. However, it is intermediate between the two in terms of storage allocated. Part of the difference in performance between it and the LR parser is due to the fact that it performs around 15% more unifications. This might be expected since the corresponding finite state automaton is not determinised—to avoid theoretical exponential time complexity on grammar size—thus paying a price at run time. Additional reasons for the relatively poor performance of the CE parser are the overheads involved in maintaining a sparse representation of the chart, and the fact that with the ANLT grammar it generates less "densely packed" parse forests, since its parse table, with 14% more states (though fewer actions) than the LALR(1) table, encodes more contextual distinctions (Billot & Lang, 1989:146).

Given that the ANLT and CLARE2.5 grammars have broadly similar (wide) coverage and return very similar numbers of syntactic analyses for the same inputs, the significantly better throughput of the three parsers described in this paper over the CLE parser[6] indicates that they do not contain any significant implementational deficiencies which would bias the results[7].

SWAPPING THE GRAMMARS OVER

A second experiment was carried out with the CLE parser, in which the built-in grammar and lexicon were replaced by versions of the ANLT object grammar and lexical entries translated (automatically) into the CLE formalism. (The reverse of this configuration, in which the CLARE2.5 grammar is translated into the ANLT formalism, is not possible since some central rules contain sequences of daughters specified by a single 'list' variable, which has no counterpart in the ANLT and cannot directly be simulated). The through-

put of this configuration was only one fiftieth of that of the BU-LC parser. The ANLT grammar contains more than five times as many rules than does the sentence-level portion of the CLARE2.5 grammar, and Alshawi (personal communication) points out that the CLE parser had not previously been run with a grammar containing such a large number of rules, in contrast to the ANLT parsers.

THE EFFECT OF SENTENCE LENGTH

Although the mean sentence length in the first two experiments is much shorter than the 20–30 word length (depending on genre etc.) that is common in real texts, the test sentences cover a wide range of syntactic constructions and exhibit less constructional bias than would a set of sentences extracted at random from a single corpus. However, to investigate performance on longer sentences and the relationship between sentence length and parse time, a further set of 100 sentences with lengths distributed uniformly between 13 and 30 words was created by hand by the author and added to the previous test data. Table 2 shows the relationship between sentence length and mean parse time with the BU-LC and LR parsers.

In contrast to the results from the first experiment, the throughput of the LR parser is only 4% better than that of the BU-LC parser for sentences of 13–27 words in length. The former parses many sentences up to twice as fast, but a small proportion of the others are parsed almost twice as slowly. As well as their wide variability with respect to the BU-LC parser, the absolute variability of the LR parse times is high (reflected in large standard deviations—σ—see Table 2). Most of the sentences for which LR performance is worse contain more than one occurrence of the passive construction: due to their length this is particularly the case for the group of sentences of 28–30 words with which the LR parser performed particularly badly. However, it is likely that if the constraining power of the parse table were improved in this area the difference in throughput between LR and BU-LC would revert to nearer the 35% figure seen in the first experiment.

The standard deviations for numbers of parses are also relatively large. The maximum number of parses was 2736 for one 29-word sentence, but on the other hand some of even the longest sentences had fewer than ten parses. (But note that since the time taken for parse forest unpacking is not included in parse times, the latter do not vary by such a large magnitude).

The results of this experiment are displayed graphically in Figure 1, together with a quadratic function. Comparison with the function suggests

---

[6]Although the ANLT parser is implemented in Common Lisp and the CLE parser in Prolog, comparing parse times is a valid exercise since current compiler and run-time support technologies for both languages are quite well-developed, and in fact the CLE parser takes advantage of Prolog's built-in unification operation which will have been very tightly coded.

[7]The ANLT's speed advantage over CLARE is less pronounced if the time for morphological analysis and creation of logical forms is taken into account, probably because the systems use different processing techniques in these modules.

Sentence length (words)	BU-LC Parse time		LR Parse time		Number of parses	
	Mean	σ	Mean	σ	Mean	σ
1–3	0.11	0.06	0.05	0.02	1.3	0.7
4–6	0.23	0.18	0.15	0.11	1.4	0.8
7–9	0.42	0.24	0.28	0.17	1.8	1.3
10–12	1.17	0.92	0.76	0.52	3.8	2.4
13–15	0.97	0.28	0.86	0.38	10.0	13.7
16–18	1.92	0.75	1.89	1.00	14.3	17.5
19–21	3.54	1.42	3.74	2.46	60.1	117.3
22–24	3.87	1.62	3.61	3.07	143.8	200.1
25–27	5.45	1.98	5.05	3.59	168.8	303.1
28–30	7.86	2.37	12.89	5.65	343.5	693.7

Table 2: Mean and standard deviation parse times (in CPU seconds on an HP9000/710 workstation), and numbers of parses for the 229 test sentences (1–30 words in length) with the BU-LC and LR parsers.

that, at least for the BU-LC parser, parse time is related roughly quadratically to input length.

In previous work with the ANLT (Briscoe & Carroll, 1993), throughput with raw corpus data was worse than that observed in these experiments, though probably only by a constant factor. This could be due to the fact that the vocabulary of the corpus concerned exhibits significantly higher lexical ambiguity; however, for sentences taken from a specific corpus, constructional bias observed in a training phase could be exploited to improve performance (e.g. Samuelsson & Rayner, 1991).

5. DISCUSSION

All three of the parsers have theoretical worst-case complexities that are either exponential, or polynomial on grammar size but with an extremely large multiplier. Despite this, in the practical experiments reported in the previous section the parsers achieve relatively good throughput with a general-purpose wide-coverage grammar of a natural language. It therefore seems likely that grammars of the type considered in this paper (i.e. with relatively detailed phrase structure components, but comparatively simple from a unification perspective), although realistic, do not bring the parsing algorithms involved anywhere near the worst-case complexity.

In the experiments, the CE technique results in a parser with worse performance than the normal LR technique. Indeed, for the ANLT grammar, the number of states—the term that the CE technique reduces from exponential to linear on the grammar size—is actually smaller in the standard LALR(1) table. This suggests that, when considering the complexity of parsers, the issue of parse table size is of minor importance for realistic NL grammars (as long as an implementation represents the table compactly), and that improvements to complexity results with respect to grammar size, although interesting from a theoretical standpoint, may have little practical relevance for the processing of natural language.

Although Schabes (1991:107) claims that the problem of exponential grammar complexity "is particularly acute for natural language processing since in this context the input length is typically small (10–20 words) and the grammar size very large (hundreds or thousands of rules and symbols)", the experiments indicate that, with a wide-coverage NL grammar, inputs of this length can be parsed quite quickly; however, longer inputs (of more than about 30 words in length)—which occur relatively frequently in written text—are a problem. Unless grammar size takes on proportionately much more significance for such longer inputs, which seems implausible, it appears that in fact the major problems do not lie in the area of grammar size, but in input length.

All three parsers have worst-case complexities that are exponential on input length. This theoretical bound might suggest that parsing performance would be severely degraded on long sentences; however, the relationship between length of sentence and parse time with the ANLT grammar and the sentences tested appears to be approximately only quadratic. There are probably many reasons why performance is much better than the complexity results suggest, but the most important may be that:

- kleene star is used only in a very limited context (for the analysis of coordination),

- more than 90% of the rules in the grammar have no more than two daughters, and

- very few rules license both left and right recursion (for instance of the sort that is typically used to analyse noun compounding, i.e.

Figure 1: Mean parse times (in CPU seconds on an HP9000/710 workstation) for the test sentences with the BU-LC and LR parsers. A quadratic function is also displayed.

N --> N N).

Despite little apparent theoretical difference between the CLE and ANLT grammar formalisms, and the fact that no explicit or formal process of 'tuning' parsers and grammars to perform well with each other has been carried out in either of the ANLT or CLARE systems, the results of the experiment comparing the performance of the respective parsers using the ANLT grammar suggests that the parallel development of the software and grammars that has occurred nevertheless appears to have caused this to happen automatically. It therefore seems likely that implementational decisions and optimisations based on subtle properties of specific grammars can, and may very often be, more important than worst-case complexity when considering the practical performance of parsing algorithms.

6. CONCLUSIONS

The research reported is in a similar vein to that of, for example, Moore & Dowding (1991), Samuelsson & Rayner (1991), and Maxwell & Kaplan (1993), in that it relies on empirical results for the study and optimisation of parsing algorithms rather than on traditional techniques of complexity analysis. The paper demonstrates that research in this area will have to rely on empirical data until complexity theory is developed to a point where it is sufficiently fine-grained and ac-

curate to predict how the properties of individual unification-based grammars will interact with particular parsing algorithms to determine practical performance.

REFERENCES

Aho, A., R. Sethi & J. Ullman (1986) *Compilers: principles, techniques and tools.* Reading, MA: Addison-Wesley.

Alshawi, H., D. Carter, R. Crouch, S. Pulman, M. Rayner & A. Smith (1992) *CLARE: a contextual reasoning and cooperative response framework for the Core Language Engine.* SRI International, Cambridge, UK.

Barton, G., R. Berwick & E. Ristad (1987) *Computational complexity and natural language.* Cambridge, MA: MIT Press.

Billot, S. & B. Lang (1989) "The structure of shared forests in ambiguous parsing." In *Proceedings of the 27th Meeting of the Association for Computational Linguistics.* 143–151.

Bouma, G. & G. van Noord (1993) "Head-driven parsing for lexicalist grammars: experimental results." In *Proceedings of the 6th Conference of the European Chapter of the Association for Computational Linguistics.* 101–105.

Briscoe, E., C. Grover, B. Boguraev & J. Carroll (1987a) "A formalism and environment for the development of a large grammar of English." In *Proceedings of the 10th International Joint Conference on Artificial Intelligence.* 703–708.

293

Briscoe, E., C. Grover, B. Boguraev & J. Carroll (1987b) "Feature defaults, propagation and reentrancy." In *Categories, Polymorphism and Unification*, edited by E. Klein & J. van Benthem, Centre for Cognitive Science, Edinburgh University, UK. 19–34.

Briscoe, E. & J. Carroll (1993) "Generalised probabilistic LR parsing of natural language (corpora) with unification-based grammars." *Computational Linguistics*, 19(1): 25–59.

Carroll, J. (1993) *Practical unification-based parsing of natural language.* Computer Laboratory, Cambridge University, UK, Technical Report 314.

Carroll, J. & E. Briscoe (1992) "Probabilistic normalisation and unpacking of packed parse forests for unification-based grammars." In *Proceedings of the AAAI Fall Symposium on Probabilistic Approaches to Natural Language.* 33–38.

Earley, J. (1970) "An efficient context-free parsing algorithm." *Communications of the ACM*, 13.2: 94–102.

Gazdar, G., E. Klein, G. Pullum & I. Sag (1985) *Generalized phrase structure grammar.* Oxford, UK: Blackwell.

Grover, C., J. Carroll & E. Briscoe (1993) *The Alvey natural language tools grammar (4th release).* Computer Laboratory, Cambridge University, UK, Technical Report 284.

Johnson, M. (1989) "The computational complexity of Tomita's algorithm." In *Proceedings of the 1st International Workshop on Parsing Technologies.* 203–208.

Kaplan, R. (1987) "Three seductions of computational psycholinguistics." In *Linguistic Theory and Computer Applications*, edited by P. Whitelock *et al.*, New York: Academic Press. 149–188.

Kasami, J. (1965) *An efficient recognition and syntax analysis algorithm for context-free languages.* Air Force Cambridge Research Laboratory, Bedford, MA, Report AFCRL-65-758.

Kipps, J. (1989) "Analysis of Tomita's algorithm for general context-free parsing." In *Proceedings of the 1st International Workshop on Parsing Technologies.* 193–202.

Lang, B. (1974) "Deterministic techniques for efficient non-deterministic parsers." In *Automata, Languages and Programming, Lecture Notes in Computer Science 14*, edited by J. Loeckx, Berlin, Germany: Springer-Verlag. 255–269.

Maxwell, J. III & R. Kaplan (1993) "The interface between phrasal and functional constraints." *Computational Linguistics*, 19(4): 571–590.

Moore, R. & H. Alshawi (1992) "Syntactic and semantic processing." In *The Core Language Engine*, edited by H. Alshawi, Cambridge, MA: MIT Press. 129–148.

Moore, R. & J. Dowding (1991) "Efficient bottom-up parsing." In *Proceedings of the DARPA Speech and Natural Language Workshop.* 200–203.

Pereira, F. & D. Warren (1980) "Definite clause grammars for language analysis—a survey of the formalism and a comparison with augmented transition networks." *Artificial Intelligence*, 13(3): 231–278.

Pollard, C. & I. Sag (1987) *Information-based syntax and semantics: volume 1–fundamentals.* Chicago, IL: University of Chicago Press.

Pratt, V. (1975) "LINGOL – a progress report." In *Proceedings of the 5th International Joint Conference on Artificial Intelligence.* 422–428.

Samuelsson, C. & M. Rayner (1991) "Quantitative evaluation of explanation-based learning as an optimization tool for a large-scale natural language system." In *Proceedings of the 12th International Joint Conference on Artificial Intelligence.* 609–615.

Schabes, Y. (1991) "Polynomial time and space shift-reduce parsing of arbitrary context-free grammars." In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics.* 106–113.

Taylor, L., C. Grover & E. Briscoe (1989) "The syntactic regularity of English noun phrases." In *Proceedings of the 4th European Meeting of the Association for Computational Linguistics.* 256–263.

Tomabechi, H. (1991) "Quasi-destructive graph unification." In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics.* 315–322.

Tomita, M. (1987) "An efficient augmented-context-free parsing algorithm." *Computational Linguistics*, 13(1): 31–46.

Shann, P. (1989) "The selection of a parsing strategy for an on-line machine translation system in a sublanguage domain. A new practical comparison." In *Proceedings of the 1st International Workshop on Parsing Technologies.* 264–276.

Wright, J., E. Wrigley & R. Sharman (1991) "Adaptive probabilistic generalized LR parsing." In *Proceedings of the 2nd International Workshop on Parsing Technologies.* 154–163.

Younger, D. (1967) "Recognition and parsing of context-free languages in time n^3." *Information and Control*, 10(2): 189–208.