High Precision Extraction of Grammatical Relations

John Carroll

Cognitive and Computing Sciences
University of Sussex
Falmer, Brighton
BN1 9QH, UK

Ted Briscoe

Computer Laboratory
University of Cambridge
JJ Thomson Avenue
Cambridge CB3 0FD, UK

Abstract

A parsing system returning analyses in the form of sets of grammatical relations can obtain high precision if it hypothesises a particular relation only when it is certain that the relation is correct. We operationalise this technique—in a statistical parser using a manually-developed wide-coverage grammar of English—by only returning relations that form part of all analyses licensed by the grammar. We observe an increase in precision from 75% to over 90% (at the cost of a reduction in recall) on a test corpus of naturally-occurring text.

1 Introduction¹

Head-dependent relationships (possibly labelled with a relation type) have been advocated as a useful level of representation for grammatical structure in a number of different large-scale language-processing tasks. For instance, in recent work on statistical treebank grammar parsing (e.g. Collins, 1999) high levels of accuracy have been reached using lexicalised probabilistic models over head-dependent tuples. Bouma, van Noord and Malouf (2001) create dependency treebanks semi-automatically in order to induce dependency-based statistical models for parse selection. Lin (1998), Srinivas (2000) and others have evaluated the accuracy of both phrase structure-based and dependency parsers by matching head-dependent relations against ‘gold standard’ relations, rather than matching (labelled) phrase structure bracketings. Research on unsupervised acquisition of lexical information from corpora, such as argument structure of predicates (Briscoe and Carroll, 1997; McCarthy, 2000), word classes for disambiguation (Clark and Weir, 2001), and collocations (Lin 1999), has used grammatical relation/head/dependent tuples. Such

tuples also constitute a convenient intermediate representation in applications such as information extraction (Palmer *et al.*, 1993; Yeh, 2000), and document retrieval on the Web (Grefenstette, 1997).

A variety of different approaches have been taken for robust extraction of relation/head/dependent tuples, or *grammatical relations*, from unrestricted text. Dependency parsing is a natural technique to use, and there has been some work in that area on robust analysis and disambiguation (e.g. Lafferty, Sleator and Temperley, 1992; Srinivas, 2000). Finite-state approaches (e.g. Karlsson *et al.*, 1995; Ait-Mokhtar and Chanod, 1997; Grefenstette, 1998) have used hand-coded transducers to recognise linear configurations of words and part of speech labels associated with, for example, subject/object-verb relationships. An intermediate step may be to mark nominal, verbal etc. ‘chunks’ in the text and to identify the head word of each of the chunks. Statistical finite-state approaches have also been used: Brants, Skut and Krenn (1997) train a cascade of Hidden Markov Models to tag words with their grammatical functions. Approaches based on memory based learning have also used chunking as a first stage, before assigning grammatical relation labels to heads of chunks (Argamon, Dagan and Krymolowski, 1998; Buchholz, Veenstra and Daelemans, 1999). Blaheta and Charniak (2000) assume a richer input representation consisting of labelled trees produced by a treebank grammar parser, and use the treebank again to train a further procedure that assigns grammatical function tags to syntactic constituents in the trees. Alternatively, a hand-written grammar can be used that produces ‘shallow’ and perhaps partial phrase structure analyses from which grammatical relations are extracted (e.g. Carroll, Minnen and Briscoe, 1998; Lin, 1998).

Recently, Schmid and Rooth (2001) have described an algorithm for computing expected *governor labels* for terminal words in labelled headed

¹A previous version of this paper was presented at IWPT’01; this version contains new experiments and results.

parse trees produced by a probabilistic context-free grammar. A governor label (implicitly) encodes a grammatical relation type (such as subject or object) and a governing lexical head. The labels are *expected* in the sense that each is weighted by the sum of the probabilities of the trees giving rise to it, and are computed efficiently by processing the entire parse forest rather than individual trees. The set of terminal/relation/governing-head tuples will not typically constitute a globally coherent analysis, but may be useful for interfacing to applications that primarily accumulate fragments of grammatical information from text (such as for instance information extraction, or systems that acquire lexical data from corpora). The approach is not so suitable for applications that need to interpret complete and consistent sentence structures (such as the analysis phase of transfer-based machine translation). Schmid and Rooth have implemented the algorithm for parsing with a lexicalised probabilistic context-free grammar of English and applied it in an open domain question answering system, but they do not give any practical results or an evaluation.

In the paper we investigate empirically Schmid and Rooth’s proposals, using a wide-coverage parsing system applied to a test corpus of naturally-occurring text, extend it with various thresholding techniques, and observe the trade-off between precision and recall in grammatical relations returned. Using the most conservative threshold results in a parser that returns only grammatical relations that form part of all analyses licensed by the grammar. In this case, precision rises to over 90%, as compared with a baseline of 75%.

2 The Analysis System

In this investigation we extend a statistical shallow parsing system for English developed originally by Carroll, Minnen and Briscoe (1998). Briefly, the system works as follows: input text is labelled with part-of-speech (PoS) tags by a tagger, and these are parsed using a wide-coverage unification-based ‘phrasal’ grammar of English PoS tags and punctuation. For disambiguation, the parser uses a probabilistic LR model derived from parse tree structures in a treebank, augmented with a set of lexical entries for verbs, acquired automatically from a 10 million word sample of the British National Corpus (Leech, 1992), each entry containing subcategorisation frame information and an associated probability. The parser is therefore ‘semi-lexicalised’ in

that verbal argument structure is disambiguated lexically, but the rest of the disambiguation is purely structural.

The coverage of the grammar—the proportion of sentences for which at least one complete spanning analysis is found—is around 80% when applied to the SUSANNE corpus (Sampson, 1995). In addition, the system is able to perform parse failure recovery, finding the highest scoring sequence of phrasal fragments (following the approach of Kiefer *et al.*, 1999), and the system has produced at least partial analyses for over 98% of the sentences in the written part of the British National Corpus.

The parsing system reads off grammatical relation tuples (GRs) from the constituent structure tree that is returned from the disambiguation phase. Information is used about which grammar rules introduce subjects, complements, and modifiers, and which daughter(s) is/are the head(s), and which the dependents. In Carroll *et al.*’s evaluation the system achieves GR accuracy that is comparable to published results for other systems: extraction of non-clausal subject relations with 83% precision, compared with Grefenstette’s (1998) figure of 80%; and overall F-score² of unlabelled head-dependent pairs of 80%, as opposed to Lin’s (1998) 83%³ and Srinivas’s (2000) 84% (this with respect only to binary relations, and omitting the analysis of control relationships). Blaheta and Charniak (2000) report an F-score of 87% for assigning grammatical function tags to constituents, but the task, and therefore the scoring method, is rather different.

For the work reported in this paper we have extended Carroll *et al.*’s basic system, implementing a version of Schmid and Rooth’s expected governor technique (see section 1 above) but adapted for unification-based grammar and GR-based analyses. Each sentence is analysed as a set of weighted GRs where the weight associated with each grammatical relation is computed as the sum of the probabilities of the parses that relation was derived from, divided by the sum of the probabilities of all parses. So, if we assume that Schmid and Rooth’s example sentence *Peter reads every paper on markup* has 2 parses, one where *on markup* attaches to the preceding noun having overall probability 0.007 and the other where it has verbal attachment with probability 0.003, then some of the weighted GRs would be

²We use the F₁ measure defined as $2 \times \text{precision} \times \text{recall} / (\text{precision} + \text{recall})$.

³Our calculation, based on table 2 of Lin (1998).

1.0 ncsbj(reads, Peter, _)
0.7 ncmod(on, paper, markup)
0.3 ncmod(on, reads, markup)

Figure 1 contains a more extended example of a weighted GR analysis for a short sentence from the SUSANNE corpus, and also gives a flavour of the relation types that the system returns. The GR scheme is described in detail by Carroll, Briscoe and Sanfilippo (1998).

3 Empirical Results

3.1 Weight Thresholding

Our first experiment compared the accuracy of the parser when extracting GRs from the highest ranked analysis (the standard probabilistic parsing setup) against extracting weighted GRs from all parses in the forest. To measure accuracy we use the precision, recall and F-score measures of parser GRs against ‘gold standard’ GR annotations in a 10,000-word test corpus of in-coverage sentences derived from the SUSANNE corpus and covering a range of written genres⁴. GRs are in general compared using an equality test, except that in a specific, limited number of cases (described by Carroll, Minnen and Briscoe, 1998) the parser is allowed to return more generic relation types.

When a parser GR has a weight of less than one, we proportionally discount its contribution to the precision and recall scores. Thus, given a set T of GRs with associated weights produced by the parser, i.e.

$$T = \{(w_i, t_i) \mid w_i \text{ is the weight associated with GR } t_i, \text{ where } 0 < w_i \leq 1\}$$

and a set S of gold-standard (unweighted) GRs, we compute the weighted match between S and the elements of T as

$$m = \sum_{(w_i, t_i) \in T} w_i \delta(t_i \in S)$$

where $\delta(x) = 1$ if x is true and 0 otherwise. The weighted precision and recall are then

$$\frac{m}{\sum_{(w_i, t_i) \in T} w_i} \quad \text{and} \quad \frac{m}{|S|}$$

respectively, expressed as percentages. We are not aware of any previous published work using

Table 1: GR accuracy comparing extraction from just the highest-ranked parse compared to weighted GR extraction from all parses.

	Precision (%)	Recall (%)	F-score
Best parse	76.25	76.77	76.51
All parses	74.63	75.33	74.98

weighted precision and recall measures, although there is an option for associating weights with complete parses in the distributed software implementing the PARSEVAL scheme (Harrison *et al.*, 1991) for evaluating parser accuracy with respect to phrase structure bracketings. The weighted measures make sense for application tasks that can deal with sets of mutually-inconsistent GRs.

In this initial experiment, precision and recall when extracting weighted GRs from all parses were both one and a half percentage points lower than when GRs were extracted from just the highest ranked analysis (see table 1)⁵. This decrease in accuracy might be expected, though, given that a true positive GR may be returned with weight less than one, and so will not receive full credit from the weighted precision and recall measures.

However, these results only tell part of the story. An application using grammatical relation analyses might be interested only in GRs that the parser is fairly confident of being correct. For instance, in unsupervised acquisition of lexical information (such as subcategorisation frames for verbs) from text, the usual methodology is to (partially) analyse the text, retaining only reliable hypotheses which are then filtered based on the amount of evidence for them over the corpus as a whole. Thus, Brent (1993) only creates hypotheses on the basis of instances of verb frames that are reliably and unambiguously cued by closed class items (such as pronouns) so there can be no other attachment possibilities. In recent work on unsupervised learning of prepositional phrase disambiguation, Pantel and Lin (2000) derive training instances only from relevant data appearing in syntactic contexts that are guaranteed to be unambiguous. In our system, the weights on GRs indicate how certain the parser is of the associated relations being correct. We therefore investigated whether more highly weighted GRs are in fact more likely

⁵Ignoring the weights on GRs, standard (unweighted) evaluation results for all parses are: precision 36.65%, recall 89.42% and F-score 51.99.

⁴At <http://www.cogs.susx.ac.uk/lab/nlp/carroll/greval.html>.

1.0	aux(., continue, will)	0.4490	iobj(on, place, tax-payers)
1.0	detmod(., burden, a)	0.3276	ncmod(on, burden, tax-payers)
1.0	doobj(do, this, .)	0.2138	ncmod(on, place, tax-payers)
1.0	doobj(place, burden, .)	0.0250	xmod(to, continue, place)
1.0	ncmod(., burden, disproportionate)	0.0242	ncmod(., Fulton, tax-payers)
1.0	ncsubj(continue, Failure, .)	0.0086	obj2(place, tax-payers)
1.0	ncsubj(place, Failure, .)	0.0086	ncmod(on, burden, Fulton)
1.0	xcomp(to, Failure, do)	0.0020	mod(., continue, place)
0.9730	clausal(continue, place)	0.0010	ncmod(on, continue, tax-payers)
0.9673	ncmod(., tax-payers, Fulton)		

Figure 1: Weighted GRs for the sentence *Failure to do this will continue to place a disproportionate burden on Fulton taxpayers.*

Figure 2: Weighted GR accuracy as the threshold is varied.

to be correct than ones with lower weights. We did this by setting a *threshold* on the output, such that any GR with weight lower than the threshold is discarded.

Figure 2 plots weighted recall and precision as the threshold is varied between zero and one. The results are intriguing. Precision increases monotonically from 74.6% at a threshold of zero (the situation as in the previous experiment where all GRs extracted from all parses in the forest are returned) to 90.4% at a threshold of one. (The latter threshold has the effect of allowing only those GRs that form part of every single analysis to be returned). The influence of the threshold on recall is equally dramatic, although since we have not escaped the usual trade-off with precision the results are somewhat less positive. Recall decreases from 75.3% to 45.2%, initially rising slightly, then falling at a

gradually increasing rate. Between thresholds 0.99 and 1.0 there is only a two percentage point difference in precision, but recall differs by almost fourteen percentage points⁶. Over the whole range, as the threshold is increased from zero, precision rises faster than recall falls until the threshold reaches 0.65; here the F-score attains its overall maximum of 77.

It turns out that the eventual figure of over 90% precision is not due to ‘easier’ relation types (such as the dependency between a determiner and a noun) being returned and more difficult ones (for example clausal complements) being ignored. The majority of relation types are produced with frequency consistent with the overall 45% recall figure. Exceptions are *arg_mod* (encoding the English passive ‘by-phrase’) and *iobj* (indirect object), for which no GRs at all are produced. The reason for this is that both types of relation originate from an occurrence of a prepositional phrase in contexts where it could be either a modifier or a complement of a predicate. This pervasive ambiguity means that there will always be disagreement between analyses over the relation type (but not necessarily over the identity of the head and dependent themselves).

3.2 Parse Unpacking

Schmid and Rooth’s algorithm computes expected governors efficiently by using dynamic programming and processing the entire parse forest rather than individual trees. In contrast, we unpack the whole parse forest and then extract weighted GRs from each tree individually. Our implementation is certainly less elegant, but in practical terms for

⁶Roughly, each percentage point increase or decrease in precision and recall is statistically significant at the 95% level. In this and all significance tests in this paper we use a one-tailed paired *t-test* (with 499 degrees of freedom).

sentences where there are relatively small numbers of parses the speed is still acceptable. However, throughput goes down linearly with the number of parses, and when there are many thousands of parses—and particularly also when the sentence is long and so each tree is large—the parsing system becomes unacceptably slow.

One possibility to improve the situation would be to extract GRs directly from forests. At first glance this looks a possibility: although our parse forests are produced by a probabilistic LR parser using a unification-based grammar, they are similar in content to those computed by a probabilistic context-free grammar, as assumed by Schmid and Rooth’s algorithm. However, there are problems. If the test for being able to pack local ambiguities in the unification grammar parse forest is feature structure subsumption, unpacking a parse apparently encoded in the forest can fail due to non-local inconsistency in feature values (Oepen and Carroll, 2000)⁷, so every governor tuple hypothesis would have to be checked to ensure that the parse it came from was globally valid. It is likely that this verification step would cancel out the efficiency gained from using an algorithm based on dynamic programming. This problem could be side-stepped (but at the cost of less compact parse forests) by instead testing for feature structure equivalence rather than subsumption. A second, more serious problem is that some of our relation types encode more information than is present in a single governor tuple (the non-clausal subject relation, for instance, encoding whether the surface subject is the ‘deep’ object in a passive construction); this information can again be less local and violate the conditions required for the dynamic programming approach.

Another possibility is to compute only the n highest ranked parses and extract weighted GRs from just those. The basic case where $n = 1$ is equivalent to the standard approach of computing GRs from the highest probability parse. Table 2 shows the effect on accuracy as n is increased in stages to 1000, using a threshold for GR extraction of 1; also shown is the previous setup (labelled ‘unlimited’) in which all parses in the forest are considered.⁸ (All differences in precision in the table are significant to at least the 95% level, except between 1000 parses and

Table 2: Weighted GR accuracy using a threshold of 1, with respect to the maximum number of ranked parses considered.

Maximum Parses	Precision (%)	Recall (%)	F-score
1	76.25	76.77	76.51
2	80.15	73.30	76.57
5	84.94	67.03	74.93
10	86.73	62.47	72.63
100	89.59	51.45	65.36
1000	90.24	46.08	61.00
unlimited	90.40	45.21	60.27

an unlimited number). The results demonstrate that limiting processing to a relatively small, fixed number of parses—even as low as 100—comes within a small margin of the accuracy achieved using the full parse forest. These results are striking, in view of the fact that the grammar assigns more than 300 parses to over a third of the sentences in the test corpus, and more than a thousand parses to a fifth of them. Another interesting observation is that the relationship between precision and recall is very close to that seen when the threshold is varied (as in the previous section); there appears to be no loss in recall at a given level of precision. We therefore feel confident in unpacking a limited number of parses from the forest and extracting weighted GRs from them, rather than trying to process all parses. We have tentatively set the limit to be 1000, as a reasonable compromise in our system between throughput and accuracy.

3.3 Parse Weighting

The way in which the GR weighting is carried out does not matter when the weight threshold is equal to 1 (since then only GRs that are part of every analysis are returned, each with a weight of one). However, we wanted to see whether the precise method for assigning weights to GRs has an effect on accuracy, and if so, to what extent. We therefore tried an alternative approach where each GR receives a contribution of 1 from every parse, no matter what the probability of the parse is, normalising in this case by the number of parses considered. This tends to increase the numbers of GRs returned for any given threshold, so when comparing the two methods we found thresholds such that each method obtained the same precision figure (of roughly 83.38%). We then compared the recall figures (see table 3). The recall

⁷The forest therefore also ‘leaks’ probability mass since it contains derivations that are in fact not legal.

⁸At $n = 1000$ parses, the (unlabelled) weighted precision of head-dependent pairs is 91.0%.

Table 3: Accuracy at the same level of precision using different weighting methods, with a 1000-parse tree limit.

Weighting Method	Precision (%)	Recall (%)	F-score
Probabilistic (at threshold 0.99)	88.38	59.19	70.90
Equally (at threshold 0.768)	88.39	55.17	67.94

for the probabilistic weighting scheme is 4% higher (statistically significant at the 99.95% level).

3.4 Maximal Consistent Relation Sets

It is interesting to see what happens if we compute for each sentence the maximal consistent set of weighted GRs. (We might want to do this if we want complete and coherent sentence analyses, interpreting the weights as confidence measures over sub-analysis segments). We use a ‘greedy’ algorithm to compute consistent relation sets, taking GRs sorted in order of decreasing weight and adding a GR to the set if and only if there is not already a GR in the set with the same dependent. (But note that the correct analysis may in fact contain more than one GR with the same dependent, such as the *nc-subj ... Failure* GRs in Figure 1, and in these cases this method will introduce errors). The weighted precision, recall and F-score at threshold zero are 79.31%, 73.56% and 76.33 respectively. Precision and F-score are significantly better (at the 95.95% level) than the baseline.

3.5 Parser Bootstrapping

One of our primary research goals is to explore unsupervised acquisition of lexical knowledge. The parser we use in this work is ‘semi-lexicalised’, using subcategorisation probabilities for verbs acquired automatically from (unlexicalised) parses. In the future we intend to acquire other types of lexico-statistical information (for example on PP attachment) which we will feed back into the parser’s disambiguation procedure, bootstrapping successively more accurate versions of the parsing system. There is still plenty of scope for improvement in accuracy, since compared with the number of correct GRs in top-ranked parses there are roughly a further 20% that are correct but present only in lower-ranked parses. There appears to be less room for improvement with argument relations (*ncsubj, dobj*

etc.) than with modifier relations (*ncmod* and similar). This indicates that our next efforts should be directed to collecting information on modification.

4 Discussion and Further Work

We have extended a shallow parsing system for English that returns analyses in the form of sets of grammatical relations, presenting an investigation into the extraction of *weighted* relations from probabilistic parses. We observed that setting a threshold on the output such that any relation with weight lower than the threshold is discarded allows a trade-off to be made between recall and precision, and found that by setting the threshold at 1 the precision of the system was boosted dramatically, from a baseline of 75% to over 90%. With this setting, the system returns only relations that form part of all analyses licensed by the grammar: the system can have no greater certainty that these relations are correct, given the knowledge that is available to it.

Although we believe this technique to be well suited to probabilistic parsers, it could also potentially benefit any parsing system that can represent ambiguity and return analyses that are composed of a collection of elementary units. Such a system need not necessarily be statistical, since parse probabilities make no difference when checking that a given sub-analysis segment forms part of all possible global analyses. Moreover, a non-statistical parsing system could use the technique to construct a reliable annotated corpus automatically, which it could then be trained on.

Acknowledgements

We are grateful to Mats Rooth for early discussions about his expected governor label work. This research was supported by UK EPSRC projects GR/N36462/93 ‘Robust Accurate Statistical Parsing (RASP)’ and by EU FP5 project IST-2001-34460 ‘MEANING: Developing Multilingual Web-scale Language Technologies’.

References

- Ait-Mokhtar, S. and J-P. Chanod (1997) Subject and object dependency extraction using finite-state transducers. In *Proceedings of the ACL/EACL’97 Workshop on Automatic Information Extraction and Building of Lexical Semantic Resources*, 71–77. Madrid, Spain.
- Argamon, S., I. Dagan and Y. Krymolowski (1998) A memory-based approach to learning shallow natural language patterns. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics*, 67–73. Montreal.

- Blaheta, D. and E. Charniak (2000) Assigning function tags to parsed text. In *Proceedings of the 1st Conference of the North American Chapter of the ACL*, 234–240. Seattle, WA.
- Bouma, G., G. van Noord and R. Malouf (2001) Alpino: wide-coverage computational analysis of Dutch. *Computational Linguistics in the Netherlands 2000. Selected Papers from the 11th CLIN Meeting*.
- Brants, T., W. Skut and B. Krenn (1997) Tagging grammatical functions. In *Proceedings of the 2nd Conference on Empirical Methods in Natural Language Processing*, 64–74. Providence, RI.
- Brent, M. (1993) From grammar to lexicon: unsupervised learning of lexical syntax. *Computational Linguistics*, 19(3), 243–262.
- Briscoe, E. and J. Carroll (1997) Automatic extraction of subcategorization from corpora. In *Proceedings of the 5th ACL Conference on Applied Natural Language Processing*, 356–363. Washington, DC.
- Buchholz, S., J. Veenstra and W. Daelemans (1999) Cascaded grammatical relation assignment. In *Proceedings of the Joint SIGDAT Conference on Empirical Methods in Natural Language Processing and Very Large Corpora*, College Park, MD. 239–246.
- Carroll, J., E. Briscoe and A. Sanfilippo (1998) Parser evaluation: a survey and a new proposal. In *Proceedings of the 1st International Conference on Language Resources and Evaluation*, 447–454. Granada, Spain.
- Carroll, J., G. Minnen and E. Briscoe (1998) Can subcategorisation probabilities help a statistical parser?. In *Proceedings of the 6th ACL/SIGDAT Workshop on Very Large Corpora*. Montreal, Canada.
- Clark, S. and D. Weir (2001) Class-based probability estimation using a semantic hierarchy. In *Proceedings of the 2nd Conference of the North American Chapter of the ACL*. Pittsburgh, PA.
- Collins, M. (1999) *Head-driven statistical models for natural language parsing*. PhD thesis, University of Pennsylvania.
- Grefenstette, G. (1997) SQLET: Short query linguistic expansion techniques, palliating one-word queries by providing intermediate structure to text. In *Proceedings of the RIAO'97*, 500–509. Montreal, Canada.
- Grefenstette, G. (1998) Light parsing as finite-state filtering. In A. Kornai (Eds.), *Extended Finite State Models of Language*. Cambridge University Press.
- Harrison, P., S. Abney, E. Black, D. Flickinger, C. Gdaniec, R. Grishman, D. Hindle, B. Ingria, M. Marcus, B. Santorini, & T. Strzalkowski (1991) Evaluating syntax performance of parser/grammars of English. In *Proceedings of the ACL'91 Workshop on Evaluating Natural Language Processing Systems*, 71–78. Berkeley, CA.
- Karlsson, F., A. Voutilainen, J. Heikkilä and A. Anttila (1995) *Constraint Grammar: a Language-Independent System for Parsing Unrestricted Text*. Berlin, Germany: de Gruyter.
- Kiefer, B., H-U. Krieger, J. Carroll and R. Malouf (1999) A bag of useful techniques for efficient and robust parsing. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, 473–480. University of Maryland.
- Lafferty, J., D. Sleator and D. Temperley (1992) Grammatical trigrams: A probabilistic model of link grammar. In *Proceedings of the AAAI Fall Symposium on Probabilistic Approaches to Natural Language*, 89–97. Cambridge, MA.
- Leech, G. (1992) 100 million words of English: the British National Corpus. *Language Research*, 28(1), 1–13.
- Lin, D. (1998) Dependency-based evaluation of MINIPAR. In *Proceedings of the The Evaluation of Parsing Systems: Workshop at the 1st International Conference on Language Resources and Evaluation*. Granada, Spain (also available as University of Sussex technical report CSRP-489).
- Lin, D. (1999) Automatic identification of non-compositional phrases. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, 317–324. College Park, MD.
- McCarthy, D. (2000) Using semantic preferences to identify verbal participation in role switching alternations. In *Proceedings of the 1st Conference of the North American Chapter of the ACL*, 256–263. Seattle, WA.
- Oepen, S. and J. Carroll (2000) Ambiguity packing in constraint-based parsing — practical results. In *Proceedings of the 1st Conference of the North American Chapter of the ACL*, 162–169. Seattle, WA.
- Palmer, M., R. Passonneau, C. Weir and T. Finin (1993) The KERNEL text understanding system. *Artificial Intelligence*, 63, 17–68.
- Pantel, P. and D. Lin (2000) An unsupervised approach to prepositional phrase attachment using contextually similar words. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*, 101–108. Hong Kong.
- Sampson, G. (1995) *English for the Computer*. Oxford University Press.
- Schmid, H. and M. Rooth (2001) Parse forest computation of expected governors. In *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics*, 458–465. Toulouse, France.
- Srinivas, B. (2000) A lightweight dependency analyzer for partial parsing. *Natural Language Engineering*, 6(2), 113–138.
- Yeh, A. (2000) Using existing systems to supplement small amounts of annotated grammatical relations training data. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*, 126–132. Hong Kong.