

Unsupervised Lexical Learning with Categorical Grammars using the LLL Corpus

Stephen Watkinson and Suresh Manandhar,
Department of Computer Science,
University of York,
York YO10 5DD,
UK.

June 4, 1999

Abstract

In this paper we report on an unsupervised approach to learning Categorical Grammar (CG) lexicons. The learner is provided with a set of possible lexical CG categories, the forward and backward application rules of CG and unmarked positive only corpora. Using the categories and rules, the sentences from the corpus are probabilistically parsed. The parses and the history of previously parsed sentences are used to build a lexicon and annotate the corpus. We report the results from experiments on two generated corpora and also on the more complicated LLL corpus, that contain examples from subsets of the English language. These show that the system is able to generate reasonable lexicons and provide accurately parsed corpora in the process. We also discuss ways in which the approach can be scaled up to deal with larger and more diverse corpora.

1 Introduction

In this paper we discuss a potential solution to two problems in Natural Language Processing (NLP), using a combination of statistical and symbolic machine learning techniques. The first problem is learning the syntactic roles, or categories, of words of a language *i.e.* learning a lexicon. Secondly, we discuss a method of annotating a corpus with parses.

The aim is to learn Categorical Grammar (CG) lexicons, starting from a set of lexical categories, the functional application rules of CG and an unannotated corpus of positive examples. The CG formalism (discussed in Section 2) is chosen because it assigns distinct categories to words of different types, and the categories describe the exact syntactic role each word can play in a sentence.

This problem is similar to the unsupervised part of speech tagging work of, for example, Brill [3] and Kupiec [8]. In Brill’s work a lexicon containing the parts of speech available to each word is provided and a simple tagger attaches a complex tag to each word in the corpus, which represents all the possible tags that word can have. Transformation rules are then learned which use the context of a word to determine which simple tag it should be assigned. The results are good, generally achieving around 95% accuracy on large corpora such as the Penn Treebank.

Kupiec [8] uses an unsupervised version of the Baum-Welch algorithm, which is a way of using examples to iteratively estimate the probabilities of a Hidden Markov Model for part of speech tagging. Instead of supplying a lexicon, he places the words in equivalence classes. Words in the same equivalence class must take one of a specific set of parts of speech. This improves the accuracy of this algorithm to about the same level as Brill’s approach.

In both cases, the learner is provided with a large amount of background knowledge – either a *complete* lexicon or set of equivalence classes. In the approach presented here, the most that is provided is a *small* partial lexicon. In fact the system learns the lexicon.

The second problem – annotating the corpus – is solved because of the approach we use to learn the lexicon. The system uses parsing to determine which are the correct lexical entries for a word, thus annotating the corpus with the parse derivations (also providing less probable parses if desired). An example of another approach to doing this is the Fidditch parser of Hindle [5] (based on the deterministic parser of Marcus [9]), which was used to annotate the Penn Treebank [10]. However, instead of learning the lexicon, a complete grammar and lexicon must be supplied to the Fidditch parser.

Our work also relates to CG induction, which has been attempted before. Osborne [11] has an algorithm that learns a grammar for sequences of *part-of-speech tags* from a *tagged* corpora, using the Minimum Description Length (MDL) principle – a well-defined form of compression. While this is a supervised setting of the problem, the use of the more formal approach to compression is of interest for future work. Also, results of 97% coverage are impressive, even though the problem is rather simpler, as it uses tagged corpora and learns from sequences of tags. Kanazawa [6] and Buszkowski [4] use a unification based approach with a corpus annotated with semantic structure, which in CG is a strong indicator of the syntactic structure. Unfortunately, they do not present results of experiments on natural language corpora and again the approach is essentially supervised.

Two unsupervised approaches to learning CGs are presented by Adriaans [1] and Solomon [13]. Adriaans, describes a purely symbolic method that uses the context of words to define their category. An oracle is required for the learner to test its hypotheses, thus providing negative evidence. This would seem to be awkward from an engineering view point *i.e.* how one could provide an oracle to achieve this, and implausible from a psychological point of view, as humans do not seem to receive such evidence [12]. Unfortunately, again no results on natural language corpora seem to be available.

Solomon’s approach [13] uses unannotated corpora, to build lexicons for simple CG. He uses a simple corpora of sentences from children’s books, with a slightly *ad hoc* and non-incremental, heuristic approach to developing categories for words. The results show

that a wide range of categories can be learned, but the current algorithm, as the author admits, is probably too naive to scale up to working on full corpora. No results on the coverage of the CGs learned are provided.

In Section 3 we discuss our learner. In Section 4 we describe experiments on three corpora containing examples of a subset of English and Section 5 contains the results, which are encouraging with respect to both problems. Finally, in Section 6, we compare the results with the systems mentioned above and discuss ways the system can be expanded and larger scale experiments may be carried out. Next, however, we describe Categorical Grammar.

2 Categorical Grammar

Categorical Grammar (CG) [17, 14] provides a functional approach to lexicalised grammar, and so, can be thought of as defining a syntactic *calculus*. Below we describe the basic (AB) CG, although in future it will be necessary to pursue a more flexible version of the formalism.

There is a set of *atomic* categories in CG, which are usually nouns (n), noun phrases (np) and sentences (s). It is then possible to build up *complex* categories using the two slash operators “/” and “\”. If A and B are categories then A/B is a category and A\B is a category. With basic CG there are just two rules for combining categories: the forward (FA) and backward (BA) *functional application* rules. Following Steedman’s notation [14] these are:

$$\begin{array}{l} X/Y Y \Rightarrow X \quad (FA) \\ Y X\backslash Y \Rightarrow X \quad (BA) \end{array}$$

Figure 1 shows a wide range of examples of the kinds of categories that can be used in CG. In Figure 1 the parse derivation for “John ate the apple” is presented.




Figure 1: A Example Parse in Pure CG

The CG described above has been shown to be weakly equivalent to context-free phrase structure grammars [2]. While such expressive power covers a large amount of natural language structure, it has been suggested that a more flexible and expressive formalism may capture natural language more accurately [17, 14]. This has led to some distinct branches of research into usefully extending CG.

CG has at least the following advantages for our task.

Syntactic Role	CG Category	Example
Sentence	s	<i>the dog ran</i>
Noun	n	<i>dog</i>
Noun Phrase	np	<i>the dog</i>
Intransitive Verb	s\np	<i>ran</i>
Transitive Verb	(s\np)/np	<i>kicked</i>
Ditransitive Verb	((s\np)/np)/np	<i>gave</i>
Sentential Complement Verb	(s\np)/s	<i>believe</i>
Determiner	np/n	<i>the</i>
Adjective	n/n	<i>hungry</i>
Auxiliary Verb	(s\np)/(s\np)	<i>does</i>
Preposition	(n\n)/np	<i>in</i>
	((s\np)\(s\np))/np	

Table 1: The categories available to the learner

- Learning the lexicon and the grammar is one task.
- The syntax directly corresponds to the semantics.

The first of these is vital for the work presented here. Because the syntactic structure is defined by the complex categories assigned to the words, it is not necessary to have separate learning procedures for the lexicon and for the grammar rules. Instead, it is just one procedure for learning the lexical assignments to words.

Secondly, the syntactic structure in CG parallels the semantic structure, which allows an elegant interaction between the two. While this feature of CG is not used in the current system, it could be used in the future to add semantic background knowledge to aid the learner (*e.g.* Buszkowski’s discovery procedures [4]).

3 The Learner

The system we have developed for learning lexicons and assigning parses to unannotated sentences is shown diagrammatically in Figure 2. In the following sections we explain the learning setting and the learning procedure respectively.

3.1 The Learning Setting

The input to the learning setting has five parts: the corpus, the lexicon, the CG rules, the set of legal categories and a probabilistic parser, which are discussed below.

The Corpus The corpus is a set of unannotated positive examples represented in Prolog as facts containing a list of words *e.g.*

```
ex([mary,loved,a,computer]).
```




Figure 2: A Diagram of the Structure of the Learner

The Lexicon The lexicon is at most partially available at the start, as this is what the learner induces. It is stored by the learner as a set of Prolog facts of the form:

`lex(Word, Category, Frequency).`

Where `Word` is a word, `Category` is a Prolog representation of the CG category assigned to that word and `Frequency` is the number of times this category has been assigned to this word up to the current point in the learning process.

The Rules The CG functional application rules (see Section 2) are supplied to the learner. Extra rules may be added in future for fuller grammatical coverage.

The Categories The learner has a complete set of the categories that can be assigned to a word in the lexicon. The complete set is shown in Table 1.

The Parser The system employs a probabilistic chart parser, which calculates the N most probable parses, where N is the beam set by the user. The probability of a word being assigned a category is based on the relative frequency, which is calculated from the current lexicon. This probability is smoothed (for words that have not been given fixed categories prior to execution) to allow the possibility that the word may appear as other categories. For all categories for which the word has not appeared, it is given a frequency of one. This is particularly useful for new words, as it ensures the category of a word is determined by its context.

Each non-lexical edge in the chart has a probability calculated by multiplying the probabilities of the two edges that are combined to form it. Edges between two vertices are not added if there are N edges labelled with the same category and a higher probability,

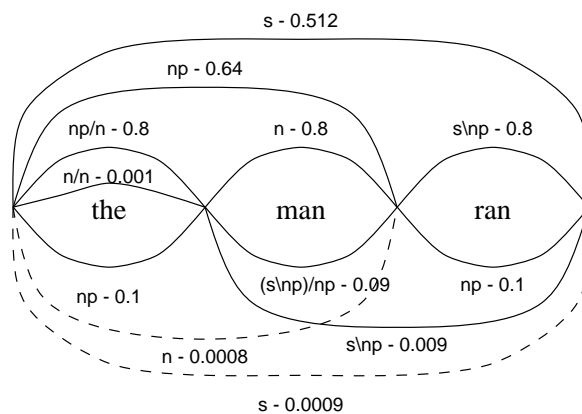


Figure 3: Example chart showing edge pruning

between the same two vertices (if one has a lower probability it is replaced). Also, for efficiency, edges are not added between vertices if there is an edge already in place with a much higher probability. The chart in Figure 3 shows examples of edges that would not be added. The top half of the chart shows one parse and the bottom half another. If N was set to 1 then the dashed edge spanning all the vertices would not be added, as it has a lower probability than the other s edge covering the same vertices. Similarly, the dashed edge between the first and third vertices would not be added, as the probability of the n is so much lower than the probability of the np .

It is important that the parser is efficient, as it is used on every example and each word in an example may be assigned any category. As will be seen it is also used extensively in selecting the best parses. In future we hope to investigate the possibility of using more restricted parsing techniques, *e.g.* deterministic parsing technology such as that described by Marcus [9], to increase efficiency and allow larger scale experiments.

3.2 The Learning Procedure

Having described the various components with which the learner is provided, we now describe how they are used in the learning procedure.

Parsing the Examples Examples are taken from the corpus one at a time and parsed. Each example is stored with the group of parses generated for it, so they can be efficiently accessed in future. The parse that is selected (see below) as the current correct parse is maintained at the head of this group. The head parse contributes information to the lexicon and annotates the corpus. The parses are also used extensively for the efficiency of the parse selection module, as will be described below. When the parser fails to find an analysis of an example, either because it is ungrammatical, or because of the incompleteness of the coverage of the grammar, the system skips to the next example.

The Parse Selector Once an example has been parsed, the N most probable parses are considered in turn to determine which can be used to make the most compressive lexicon (by a given measure), following the compression as learning approach of, for example, Wolff [16]. The current size measure for the lexicon is the sum of the sizes of the categories for each lexical entry. The size of a category is the number of atomic categories within it.

It is not enough to look at what a parse would add to the lexicon. Changing the lexicon may change the results given by the parser for previous examples. Changes in the frequency of assignments can cause the probabilities of previous parses to change. This can correct mistakes made earlier when the evidence from the lexicon was too weak to assign the correct parse. Such correction is achieved by reparsing previous examples that may be affected by the changed lexicon. Not reparsing those examples that will not be affected, saves a great deal of time. In this way a new lexicon is built from the reparsed examples for each hypothesised parse of the current example. The parse leading to the most compressive of these is chosen. The amount of reparsing is also reduced by using stored parse information.

This may appear an expensive way of determining which parse to select, but it enables the system to calculate the most compressive lexicon and keep an up-to-date annotation for the corpus. Also, the chart parser works in polynomial time and it is possible to do significant pruning, as outlined, so few sentences need to be reparsed each time. However, in the future we will look at ways of determining which parse to select that do not require complete reparsing.

Lexicon Modification The final stage takes the current lexicon and replaces it with the lexicon built with the selected parse. The whole process is repeated until all the examples have been parsed. The final lexicon is left after the final modification. The most probable annotation of the corpus is the set of top-most parses after the final parse selection.

4 Experiments

Experiments were performed on three different corpora all containing only positive examples. Experiments were performed with and without a partial lexicon of closed-class words (words of categories with a finite number of members) with fixed categories and probabilities, *e.g.* determiners and prepositions. All experiments were carried out on a SGI Origin 2000.

Experiments on Corpus 1 The first corpus was built from a context-free grammar (CFG), using a simple random generation algorithm. The CFG (shown in Figure 4) covers a range of simple declarative sentences with intransitive, transitive and ditransitive verbs and with adjectives. The lexicon of the CFG contained 39 words with an example of noun-verb ambiguity. The corpus consisted of 500 such sentences (Figure 5 shows examples). As the size of the lexicon was small and there was only a small amount of ambiguity, it was

unnecessary to supply the partial lexicon, but the experiment was carried out for comparison. We also performed an experiment on 100 unseen examples to see how accurately they were parsed with the learned lexicon. The results were manually verified to determine how many sentences were parsed correctly.

$S \rightarrow NP VP$	$VP \rightarrow Vbar$
$Vbar \rightarrow IV$	$Vbar \rightarrow TV NP$
$Vbar \rightarrow DV NP NP$	$NP \rightarrow PN$
$NP \rightarrow Nbar$	$Nbar \rightarrow Det N$
$N \rightarrow Adj N$	
$PN \rightarrow john$	$Det \rightarrow the$
$N \rightarrow boy$	$Adj \rightarrow small$
$IV \rightarrow ran$	$TV \rightarrow timed$
$DV \rightarrow gave$	

Figure 4: The CFG used to generate Corpus 1 with example lexical entries

```
ex([mary, ran]).
ex([john, gave, john, a, boy]).
ex([a, dog, called, the, fish, a, small, ugly, desk]).
```

Figure 5: Examples from Corpus 1

Experiments on Corpus 2 The second corpus was generated in the same way, but using extra rules (see Figure 6) to include prepositions, thus making the fragment of English more complicated. The lexicon used for generating the corpus was larger – 44 words in total. Again 500 examples were generated (see Figure 7 for examples) and experiments were carried out both with and without the partial lexicon. Again we performed an experiment on 100 unseen examples to see how accurately they are parsed.

$NP \rightarrow Nbar PP$	$VP \rightarrow Vbar PP$
$PP \rightarrow P NP$	
$P \rightarrow on$	

Figure 6: The extra rules required for generating Corpus 2 with example lexical entries

Experiments on Corpus 3 We also performed experiments using the LLL corpus [7]. This is a corpus of generated sentences for a substantial fragment of English. It is annotated with a certain amount of semantic information, which was ignored. The corpus contains


```
ex([the, fish, with, a, elephant, gave, banks, a, dog, with, a, bigger,
statue]).
ex([a, elephant, with, jim, walked, on, a, desk]).
ex([the, girl, kissed, the, computer, on, a, fish]).
```

Figure 7: Examples from Corpus 2

554 sentences, including a lot of movement (*e.g.* topicalized and question sentences). Examples are shown in Figure 8. While our CG rules can handle a reasonable variety of declarative sentences it is by no means complete, not allowing any movement. This is a limitation in the current approach. Also, this corpus is small and sparse, making learning difficult. It was determined to experiment to see how well the system performed under these conditions. Experiments were performed with closed-class words included only, as previous experiments on parts of the corpora [15] suggested that it would be necessary due to the sparseness of the corpus. We also present results for parsing unseen examples.

```
ex([which, rough, reports, above, hilary, wrote, hilary, in, sandy, beside,
which, telephone]).
ex([inside, no, secretary, wont, all, small, machines, stop]).
ex([which, old, report, disappears]).
```

Figure 8: Examples from Corpus 3

All experiments were performed with the minimum number of categories needed to cover the corpus, so for example, in the experiments on Corpus 1 the categories for prepositions were not available to the parser. This will obviously affect the speed with which the learner performs. Also, the parser was restricted to two possible parses in each case.

5 Results

In Table 2 we report the results of these experiments. The CCW Preset column indicates whether the closed-class words were provided or not. The lexicon accuracy column is a measure, calculated by manual analysis, of the percentage of lexical entries *i.e.* entries that have word-category pairs that can plausibly be accepted as existing in English. This should be taken together with the parse accuracy, which is the percentage of correctly parsed examples *i.e.* a linguistically correct syntactic analysis. The results for the first two corpora are extremely encouraging with 100% accuracy in both measures. While these experiments are on relatively simple corpora, these results strongly suggest the approach can be effective. Note that any experiment on corpus 2 without the closed-class words being set did not terminate, as the sentences in that corpus are significantly longer and each word may be a large number of categories. It is therefore clear, that setting the closed-class words greatly increases speed and that we need to consider methods of relieving the strain on the parser if the approach is to be useful on more complex corpora.

Corpus	CCW Preset	Lexicon Acc. (%)	Parse Acc. (%)	Exec. Time (s)
1	×	100	100	5297
1	✓	100	100	625
2	✓	100	100	10524
2	×	×	×	×
3	✓	73.2	28.5	182284

Table 2: Accuracies and timings for the different learning experiments

The results with the LLL corpus are also encouraging in part. A lexical accuracy of 73.2% is a good result especially considering that the grammar is not designed to cover sentences containing movement and there are many examples of these in this corpus. This is due to the learner getting the majority of the phrases within examples correct, with incorrect analysis of the parts that the grammars goes not cover *e.g.* the question word at the start of a sentence. Clearly a grammar with greater coverage could provide much improved results.

The results for parse accuracy, both in the training and test sets, do not suggest that the system is very successful in providing parses. However, they actually do not show the full picture. Practically all the sentences are mostly parsed correctly. Very few of the errors could actually be handled correctly by the grammar with which the system was supplied. In terms of less stringent measures of parse accuracy *e.g.* bracket crossing, the performance would be much higher. Obviously, in the longer term, provided a grammar with more coverage would possibly provide the desired results for this more stringent measure.

Table 3 shows predictably good results for parsing the test sets with the learned lexicons.

Corpus	Closed-Class	Parse Accuracy (%)
1	×	100
1	✓	100
2	×	100
2	✓	100
3	✓	37.2

Table 3: Unseen example parsing accuracy

6 Conclusions

We have presented an unsupervised learner that is able to both learn CG lexicons and annotate natural language corpora, with less background knowledge than other systems in the literature. Results from preliminary experiments are encouraging with respect to both

problems, particularly as the system appears to be reasonably effective on small, sparse corpora. It is encouraging that where errors arose this was often due only to incomplete background knowledge.

The results presented are encouraging with respect to the work that has already been mentioned - 100% can clearly not be improved upon and compares very favourable with the systems mentioned in Section 1. However, it is also clear that this was achieved on unrealistically simple corpora and when the system was used on the more diverse LLL corpus it did not fair quite as well. However, given the fact that the problem setting discussed here is somewhat harder than that attempted by other systems and the lack of linguistic background knowledge supplied, it is hoped that it will be possible to use the approach with some simple changes to greatly improve the performance on the LLL corpus and other larger coverage corpora. In particular, it will obviously be necessary to include coverage of the movement phenomena in natural language in the CG supplied.

The use of CGs to solve the problem provides an elegant way of using syntactic information to constrain the learning problem and provides the opportunity for expansion to a full grammar learning system in the future by the development of a category hypothesiser. It is hoped that this will be part of future work.

We also hope to carry out experiments on larger and more diverse corpora, as the corpora used thus far are too small to be an exacting test for the approach. We need to expand the grammar to cover more linguistic phenomena to achieve this, as well as considering other measures for compressing the lexicon (*e.g.* using an MDL-based approach). Larger experiments will lead to a need for increased efficiency in the parsing and reparsing processes. This could be done by considering deterministic parsing approaches [9], or perhaps shallower syntactic analysis.

While many extensions may be considered for this work, the evidence thus far suggests that the approach outlined in this paper is effective and efficient for these natural language learning tasks.

References

- [1] Pieter Willem Adriaans. *Language Learning from a Categorical Perspective*. PhD thesis, Universiteit van Amsterdam, 1992.
- [2] Y. Bar-Hillel, C. Gaifman, and E. Shamir. On categorial and phrase structure grammars. In *Language and Information*, pages 99 – 115. Addison-Wesley, 1964. First appeared in *The Bulletin of the Research Council of Israel*, vol. 9F, pp. 1-16, 1960.
- [3] Eric Brill. Unsupervised learning of disambiguation rules for part of speech tagging. In *Natural Language Processing Using Very Large Corpora*. Kluwer Academic Press, 1997.
- [4] Wojciech Buszkowski. Discovery procedures for categorial grammars. In Ewan Klein and Johan van Benthem, editors, *Categories, Polymorphism and Unification*, pages 35

- 64. Centre for Cognitive Science, University of Edinburgh and Institute for Language, Logic and Information, University of Amsterdam, 1987.
- [5] Donald Hindle. Deterministic parsing of syntactic non-fluencies. In Mitch Marcus, editor, *Proceedings of the 21st Annual Meeting of the Association for Computational Linguistics*, pages 123 – 128. Association for Computational Linguistics, 1983.
- [6] Makoto Kanazawa. *Learnable Classes of Categorical Grammars*. PhD thesis, Institute for Logic, Language and Computation, University of Amsterdam, 1994.
- [7] Dimitar Kazakov, Stephen Pulman, and Stephen Muggleton. The FraCas dataset and the LLL challenge. Technical report, SRI International, 1998.
- [8] Julian Kupiec. Robust part-of-speech tagging using a hidden markov model. *Computer Speech and Language*, 6:225–242, 1992.
- [9] Mitchell P. Marcus. *A Theory of Syntactic Recognition*. The MIT Press Series in Artificial Intelligence. The MIT Press, 1980.
- [10] Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. Building a large annotated corpus of english: The penn treebank. Technical Report IRCS-93-47, Institution for Research in Cognitive Science, 1993.
- [11] Miles Osborne. Minimisation, indifference and statistical language learning. In *Workshop on Empirical Learning of Natural Language Processing Tasks, ECML’97*, pages 113 – 124, 1997.
- [12] Steven Pinker. Language acquisition. In Daniel N. Osherson and Howard Lasnik, editors, *An Invitation to Cognitive Science: Language*, volume 1, pages 199–241. The MIT Press, 1990.
- [13] W. Daniel Solomon. Learning a grammar. Technical Report UMCS-AI-91-2-1, Department of Computer Science, Artificial Intelligence Group, University of Manchester, 1991.
- [14] Mark Steedman. Categorical grammar. *Lingua*, 90:221 – 258, 1993.
- [15] S.P. Watkinson and Suresh Manandhar. Unsupervised lexical learning with categorical grammars. To be published in ULNLP’99.
- [16] J.G. Wolff. Cognitive development as optimisation. In Leonard Bolc, editor, *Computational Models of Learning*, Symbolic computation-artificial intelligence. Springer Verlag, 1987.
- [17] Mary McGee Wood. *Categorical Grammars*. Linguistic Theory Guides. Routledge, 1993. General Editor Richard Hudson.