

# CS 224N / Ling 237

## Assignment 4: A Treebank Parser

Due Friday, 20 May 2005, 5pm

Thanks to Dan Klein for the original assignment.

## 1 Setup

The code for assignment 4 is under `/afs/ir/class/cs224n/hw4`. The precompiled classes are in `hw4/java/classes`. The data for this assignment is in `hw4/data`. Copy the starter code to your local directory, and make sure you can compile it. Spend some time looking through the principal source file for this assignment:

```
java/src/cs224n/assignments/PCFGParserTester.java
```

Make sure you can run the main method of the `PCFGParserTester` class. You can either run it from the data directory, or pass that directory in as the `-path` value. Use the `-data` value to specify which dataset to use. Running:

```
java -server -mx500m cs224n.assignments.PCFGParserTester  
-path /afs/usr/class/cs224n/hw4/data -data miniTest
```

will train and test your parser on a few sentences from a toy grammar. Running:

```
java -server -mx500m cs224n.assignments.PCFGParserTester  
-path /afs/usr/class/cs224n/hw4/data -data treebank
```

will train and test your parser on the WSJ section of the Penn Treebank dataset.

## 2 Building A PCFG Parser

In this project, you will build a broad-coverage parser. You're free to build a beam-decoded shift-reduce parser, or implement any other general (P)CFG parsing solution you find interesting and manageable. However, the bulk of the support code assumes that you will

be building a parser where the grammar rules have been pre-transformed into exclusively unary and binary grammar rules. The easiest thing to build would be a probabilistic CKY parser. Another good thing to attempt to build is an agenda-driven PCFG parser.

At the beginning of the main method in `PCFGParserTester` some training and test trees are read in. Currently, the training trees are used to construct a `BaselineParser` that implements the `Parser` interface (which only has one method: `getBestParse`). The parser is then used to predict trees for the sentences in the test set. For each test sentence the parse given by your parser is evaluated by comparing the constituents it generates with the constituents in the hand-parsed version. From this, precision, recall, and the F1 score are calculated.

This baseline parser is really quite poor – it takes a sentence, tags each word with its most likely tag (i.e., a unigram tagger), then looks for occurrences of the tag sequence in the training set. If it finds an exact match, it answers with the training parse of the matching training sentence. If no match is found, it constructs a right-branching tree, with nodes' labels chosen independently, conditioned only on the length of the span of a node. If this sounds like a strange (and terrible) way to parse, it should, and you're going to provide a better solution.

You should familiarize yourself with these basic classes:

<code>ling.Tree</code>	CFG tree structures, (pretty-print	with
<code>Lexicon</code>	<code>ling.Trees.PennTreeRenderer</code> )	
<code>Grammar</code> , <code>UnaryRule</code> , <code>BinaryRule</code>	Pre-terminal productions and probabilities	
	CFG rules and accessors	

`Tree` is a linguistic tree class that you will use no matter what kind of parser you implement. `Lexicon` is a minimal lexicon, but it handles rare and unknown words adequately for the present purposes. You can use it to determine the pre-terminal productions for your parser if you like. You might also want to consider using your POS-tagger to fix the tags for a sentence instead. `Grammar` is a class you can use to learn a PCFG from the training trees. Since the training set is hand-parsed this learning is very easy. We simply set

$$\hat{P}(N^j \rightarrow \zeta) = \frac{C(N^j \rightarrow \zeta)}{\sum_{\gamma} C(N^j \rightarrow \gamma)}$$

where  $C(N^j \rightarrow \gamma)$  is count observed for that production in the data set. While you could consider smoothing rule rewrite probabilities, it is sufficient for this assignment to just work with unsmoothed MLE probabilities for rules. (Doing anything else makes things rather more complex and slow, since every rewrite will have a nonzero probability, so you should definitely get things working with an unsmoothed grammar before considering adding smoothing.) `UnaryRule`, `BinaryRule` are simply the classes the grammar uses to store these learned productions. They each bear the frequency estimated probabilities from the training set.

(If you build an agenda-driven PCFG parser, you will also find `util.PriorityQueue` useful. It is a fast priority queue implementation, but it does not support a promotion /

increasePriority operation. If you want to increase the priority of an edge in the queue, you need to add that edge a second time, with the new higher priority, and be aware that whenever you pop an edge off of the queue, it might be a duplicate “dead” edge. We made this design decision because, in our experience, highly-tuned agenda-driven parsers spend a large fraction of their time dealing with the overhead that queues supporting promotion incur. But you are of course welcome to change the priority queue if you’d really like it to support promotion.)

Although it is not required, we strongly recommend that you get your parser working on the miniTest dataset before you attempt the treebank dataset. The miniTest data set consists of 3 training sentences and 1 test sentence from a toy grammar. There are just enough productions in the training set for the test sentence to have an ambiguity due to PP-attachment. There are unary, binary, and trinary grammar rules in in training sentences.

As we discussed in class, most parsers require grammars in which the rules are at most binary branching. You can binarize and unbinarize trees using the `TreeAnnotations` class. The implementation we give you binarizes the trees in a way that doesn’t generalize the grammar at all. You should run some trees through the binarization process to get the idea of what’s going on. If you annotate/binarize the training trees, you should be able to construct a Grammar out of them using the constructor provided.

Your first job is to build a parser using this grammar. For the miniTest dataset your parser should match the given parse of the test sentence exactly. Once you’ve got this working you can move on to the treebank.

Scan through a few of the training trees in the treebank dataset to get a sense of range of inputs. Something you’ll notice is that the grammar has relatively few non-terminal symbols (27 plus part-of-speech tags) but thousands of rules, many trinary-branching or longer. Currently, files 200 to 2199 of the Treebank are read in as training data, 2200 to 2299 are validation data, and 2300 to 2319 are test data (you can look in the data directory if you’re curious about the native format of these files). At the moment, only the training and test set are used, but you are welcome to use the validation set too if you see a use for it. The static integer `MAX_LENGTH` determines the maximum length of sentences to test on (it does not affect the training set). You can lower `MAX_LENGTH` for preliminary experiments, but your final parser should work on sentences of at least length 20 in a reasonable time (5 seconds per 20-word sentence should be achievable without too much optimization).

Once you have a parser working on the tree bank your next task is improve upon the supplied grammar by adding 2nd-order vertical markovization. This means using parent annotation symbols like `NP^S` to indicate a subject noun phrase instead of just `NP`. You can test your new grammar on the miniTest data set if you want though the results won’t be very interesting. When you test it on the treebank the results should be pretty impressive: a 3-4% improvement over the a parser using the original grammar.

At this point an F1 performance of 80% is probably achievable (but don’t worry too much about this exact figure, its just an estimate). With a little more work it may be possible to improve upon that by 2-5%.

## 2.1 Coding Tips

Whenever you run the java VM, you should invoke it with as much memory as you need, and in server mode:

```
java -server -mx500m package.ClassName
```

You'll get much faster performance than just running with no options.

If your parser is running very slowly, run the VM with the `-Xprof` command line option. This will result in a flat profile being output after your process completes. If you see that your program is spending a lot of time in hash map, hash code, or equals methods, you might be able to speed up your computation substantially by backing your sets, maps, and counters with `IdentityHashMap` instead of `HashMap`. This requires the use of something like an `Interner` for canonicalization ask around if you're not sure what that would entail, some people in the class already know this trick.

## 2.2 The Requirements

To summarize you are expected to implement at least the following:

- A PCFG parser, probably a CKY parser of an agenda-driven one, that can parse sentences of at least length 20 at a rate of approximately 5 sentences per second.
- An improved grammar with 2nd order vertical markovization.

Anything beyond that is optional. Some extensions are described in the next section.

## 3 Further Investigations

If you have a parser which, given a test sentence, returns a parse in a reasonable amount of time, then you have done enough to get full credit. If you are interested in extra-credit there are several ways in which you can embellish your parser:

- One choice is to focus on the grammar, using additional annotation and binarization techniques like horizontal and vertical markovization to improve the accuracy of your parser. The supplied grammar is equivalent to a 1st-order vertical process with an infinite-order horizontal process. You are required to implement 2nd-order vertical markovization, but you could also try an 3rd-order vertical markovization or horizontal markovization, meaning forgetful binarization (symbols like `@VP-> . . .NP_PP` which omit details of the history, instead of `@VP->_VBD_RB_NP_PP` which record the entire history).
- Another choice is to focus on efficiency of the parser. You could compare a beam method with the exact agenda-driven method, or implement a pruning technique like

an A\* heuristic or a figure-of-merit and compare it to the basic agenda method (this may require substantial effort). Or you could study the trade-offs of pre-tagging the sentence before parsing, rather than letting the parser do the tagging.

- A final choice is to investigate some other aspect of parsing that can be easily tested in your existing code. For example, you might investigate whether some compact subset of a grammar gives nearly the same accuracy. Or if you're interested in cognitive issues, you could study whether correct trees really do tend to have bounded stack depths in one direction or the other (and whether incorrect trees perhaps have different profiles). Or you could do an error analysis of the mistakes your parser makes.

Please note that while extensions are great, we are also aware that you need to get started on your final projects. If you are choosing between adding an extension for this homework or spending more time on your final project, we recommend that you focus on the final project.

## 4 Write Up

Your write up should include the following:

- A brief description of your implementation including any important design decisions you made.
- An evaluation of successes and failures of your parser. You don't have to do a detailed data analysis, but you should give a few examples to illustrate any errors your parser makes often and describe its performance in general.
- A discussion of further improvements you might make to deal with some of the observed errors.
- An explanation of any extensions you implemented and an analysis of whether or not they improved performance.

As for the last assignment you should turn in a hard copy of your write up. There is no required length but you can probably do it 4 pages or less. We only need one copy of the write up per group. If you finish your report over the weekend and cannot get into Gates to turn it in you can time stamp your write up and turn it in on Monday. It is your responsibility under the honor code to time stamp your write up with the time when you actually finished it.

## 5 Submitting the assignment

### 5.1 The program: electronic submission

Submission of the program code will be via a script. To submit your program, put all the needed files in one directory on a Sweet Hall machine (or, possibly, another machine with AFS on it and with you being correctly authenticated) and type:

```
/afs/ir/class/cs224n/bin/submit-hw4
```

If you need to resubmit it type

```
/afs/ir/class/cs224n/bin/submit-hw4 -replace
```

Make sure that you include all the source code for your programs. We will run programs on the Sweet Hall systems, so they should run without problems there. Please don't include large data files in your submission. Your code doesn't have to be beautiful but we should be able to scan it and figure out what you did without too much pain.