

CS 288: Statistical NLP

Assignment 4: Parsing

Due 4/6/09

In this assignment, you will build an English treebank parser. You will consider both the problem of learning a grammar from a treebank and the problem of parsing with that grammar.

Setup: The data for this assignment is available on the web page as usual. It uses the same Penn Treebank data as the first and third assignments, this time with full parses.

The starting class for this assignment is

```
edu.berkeley.nlp.assignments.PCFGParserTester
```

Make sure you can access the source and data files.

Description: In this project, you will build a broad-coverage parser. You may either build an agenda-driven PCFG parser, or an array-based CKY parser. I will first go over the data flow, then describe the support classes that are provided. You are free to use these classes, or not, as you see fit.

Currently, files 200 to 2199 of the Treebank are read in as training data, as is standard for this data set. Depending on whether you run with `-validate` or `-test`, either files 2200 to 2299 are read in (validation), or 2300 to 2399 are read in (test). You can look in the data directory if you're curious about the native format of these files, but all I/O is taken care of by the provided code. You can always run on fewer training or test files to speed up your preliminary experiments, especially while debugging (where you might first want to train and test on the same, single file, or even just a single tree).

Once the trees are read, the code constructs a `BaselineParser`, which implements the `Parser` interface (with only one method: `getBestParse()`). The parser is then used to predict trees for the sentences in the test set. You can control the maximum length of training and testing sentences with the parameters `-maxTrainLength` and `-maxTestLength`. The standard setup is to train on sentences of all lengths and test on sentences of length less than or equal to 40 words. 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; a good research parser will parse a 20 word sentence in more like 0.1 seconds).

This baseline parser is quite terrible - it takes a sentence, tags each word with its most likely tag (i.e. runs a unigram tagger), then looks for that exact tag sequence in the training set. If it finds an exact match, it answers with a known parse for that tag sequence. 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. This baseline is just a crazy placeholder — you're going to provide a better solution.

You should familiarize yourself with these basic classes:

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

You will be using these classes no matter what kind of parser you build. If you choose to build an agenda-based parser, you will find `util.GeneralPriorityQueue` useful. If you choose to build an array-based CKY parser, you will instead find the `UnaryClosure` class (also in the harness file) useful. It computes the reflexive, transitive closure of the unary subset of a grammar, and also maps closure rules to their best backing paths.

I have also provided a basic lexicon (`Lexicon`). This lexicon is minimal, but handles rare and unknown words adequately for the present purposes. If you want to employ your tagger from assignment 2 instead of using this lexicon, that is perfectly acceptable.

At this point, you should manually scan through a few of the training trees to get a sense of the format and 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. As we discussed in class, most parsers (including yours) require grammars in which the rules are at most binary branching. You can binarize and unbinarize trees using the `TreeAnnotations` class. The default implementation binarizes the trees in a way that doesn't generalize the n-ary grammar at all (convince yourself of this). You should run some trees through the binarization process and look at the results to get an 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. This grammar is composed of binary and unary rules, each bearing its relative frequency estimated probability from the training trees you provide on construction. It therefore encodes a PCFG, and your main goal is to build a parser which parses novel sentences using that PCFG. Building this parser will be the bulk of the work for this assignment.

Extensions: Once you've got a parser which, given a test sentence, returns a parse of that sentence using the training grammar, you have several choices in how to proceed with this assignment. One choice is to focus on the grammar, using better annotation / refinement techniques like horizontal and vertical markovization to improve the accuracy of your parser. The current representation is equivalent to a 1st-order vertical process with an infinite-order horizontal process. Everyone should at least try out a 2nd-order / 2nd-order grammar, meaning using parent annotation (symbols like `NP^S` instead of `NP`) and forgetful binarization (symbols like `@VP->...NP_PP` which abstract the horizontal history, instead of `@VP->_VBD_RB_NP_PP` which record the entire history). If you choose to focus on this aspect, however, you should do more exploration of various kinds of annotation. A truly ambitious approach would be to explore automatic state-splitting.

Another direction is to focus on efficiency of the parser. You could, for example, implement a pruning technique like A* parsing or coarse-to-fine pruning and compare it to an exhaustive approach.

Still another direction 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 the 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 a substantial error analysis of the mistakes your parser makes.

To summarize my expectations in this assignment, I expect that the bulk of your time will go into understanding the setup and then building a basic PCFG parser (agenda-driven or CKY). Once that is working, I expect you to try out a 2nd-order / 2nd-order grammar, which should take little additional work; this grammar should work much better than your original grammar. You should also report at least some errors (with examples) that your parser seems to make often, though you need not do a detailed data analysis (unless you want to focus on that). Finally, I expect you to do some other kind of experiment that you find interesting. I do NOT require this last part of the assignment to be something that will double your time on this assignment. However, not all of the extensions I mentioned will be possible unless you do decide to invest substantial additional time.

Coding Tips: Whenever you run the java VM, you should invoke it with as much memory as you need (and in server mode for JVMs which don't default to server):

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

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 operations or `hashCode` / `equals` methods, you might be able to speed up your computation substantially by backing your sets, maps, and counters with `IdentityHashMaps` instead of `HashMaps`. This change requires the use of something like a `util.Interner` for canonicalization.