

Lab 3: Recursive Descent Parser

Use the [html version](#)¹ of this lab if you want to easily access the links or copy and paste commands.
Or use the [pdf version](#)² if you want nicer formatting or a printable sheet.

Goals and motivation of this lab, and what to submit

Parsing a sentence consists of finding the correct syntactic structure for that sentence using a given grammar, where the grammar contains the structural constraints of the language. Parsing is one of the major tasks which helps in processing and analyzing natural language.

In this lab we explore *recursive descent parsing* using some simple Phrase Structure Grammars (CFGs), again with some help from NLTK. We aim to give you a sense of how much computation is *potentially* involved in parsing sentences, and thus why cleverer parsing algorithms are needed. We also aim to give you some practice writing grammars, to see how the choices you make can interact with the parsing algorithm, sometimes in undesirable ways.

The main part of this lab is based on a graphical demo app that runs inside NLTK and shows you a recursive descent parser in action. There are also some additional optional sections at the end that you can do if you want to look at the grammar rules and parsed sentences in a more realistic broad-coverage grammar (from a subset of the Penn Treebank).

At the end of today's lab session, please email me a short message including the following information (If you are working with a partner, only one email is needed for both of you):

- your name and your partner's name;
- your answer to Q3
- your grammar for Q4

I will post solutions to each lab but will not provide individual feedback, I just want to make sure that you are actually making an effort.

Preliminaries

First, download the files [lab3.py](#)³ and [lab3_fix.py](#)⁴ into your `labs` folder (or folder of your choice).

Open `lab3.py` in your Python editor and start up a Python interpreter.

Check that `nltk` is working by downloading the corpus for this lab (actually only used in the optional sections at the end):

```
import nltk
nltk.download('treebank')
```

Running the code

First, in your iPython window, type:

```
cd <foldername>
```

¹<http://homepages.inf.ed.ac.uk/sgwater/teaching/lisa2015/labs/lab3.html>

²<http://homepages.inf.ed.ac.uk/sgwater/teaching/lisa2015/labs/lab3.pdf>

³<http://homepages.inf.ed.ac.uk/sgwater/teaching/lisa2015/labs/lab3.py>

⁴http://homepages.inf.ed.ac.uk/sgwater/teaching/lisa2015/labs/lab3_fix.py

⁵<http://homepages.inf.ed.ac.uk/sgwater/teaching/lisa2015/labs/lab2.html>

where `<foldername>` is replaced with the name of the folder where you put `lab3.py` and `lab3_fix.py`. Then run this week's lab as follows:

```
%run lab3.py
```

You may also want to check [last week's instructions](#)⁵ on how to get figures to appear in a separate window (or not), according to your preference.

Recursive Descent Parser (RDP)

We'll start with the following toy grammar, implemented as `grammar1` in the lab code:

```
# Grammatical productions.
S -> NP VP
NP -> Pro | Det N | N
Det -> Art
VP -> V | V NP | V NP PP
PP -> Prep NP
# Lexical productions.
Pro -> "i" | "we" | "you" | "he" | "she" | "him" | "her"
Art -> "a" | "an" | "the"
Prep -> "with" | "in"
N -> "salad" | "fork" | "mushrooms"
V -> "eat" | "eats" | "ate" | "see" | "saw" | "prefer" | "sneezed"
Vi -> "sneezed" | "ran"
Vt -> "eat" | "eats" | "ate" | "see" | "saw" | "prefer"
Vp -> "eat" | "eats" | "ate" | "see" | "saw" | "prefer"
```

Our code reads in this grammar and turns it into an `nltk.grammar.ContextFreeGrammar` object using `parse_grammar` (defined in `lab3_fix.py`, which you need not examine).

Run the command:

```
app(grammar1, sentence1)
```

This shows an app. The five coloured buttons at the bottom of the resulting app window control the parser. Try 'Step' a few times. Note how the display corresponds to the description of the recursive descent algorithm in Monday's lecture, in particular the way the current subgoal and possible expansions are highlighted.

Now 'Autostep'.

Watch progress in the tree window, and individual actions in the 'Last Operation' display. Hit 'Autostep' again to stop the parser. Try 'Step' again a few times, and see how the app indicates which options at a choice point have already been tried, and which remain.

Try some other sentences (using `Edit -> Edit Text` in the menu bar) and watch how the parser runs. You can speed things up by selecting `Animate -> Fast Animation`. The parser stops when it gets the first valid parse of a sentence. You can start it again to see if there are more valid parses, or you can reset it under the `File` menu.

Try parsing:

```
i ate the salad with a fork
```

Slow the animation back down, and start with `Autostep`, but stop that when it gets to work on the `Det N` version of `NP` as applied to "the salad". Step through the next bit one at a time, to see just how much backtracking (and wasted effort) is involved in getting back to the *correct* expansion of `VP`, followed by a complete (and lengthy) recapitulation of the parsing of "the salad" as an `NP`.

Adding productions to the grammar

Q1. Run the parser on the sentence `He ate salad` (note the capital 'H'). Can you parse this sentence using the toy grammar provided? What production should be added to handle this sentence?

Look at the parse trees for the following sentences:

```
he ate salad with a fork
he ate salad with mushrooms
```

Q2: Is there a problem with either of the parse trees? (*Hint*: Observe the attachments of prepositional phrases ('with a fork', 'with mushrooms') in both the parse trees).

Using the demo app, add a production $NP \rightarrow NP PP$, *after* the other NP rules. (Click on Edit->Edit Grammar to do this). Re-run the parser on one of the above sentences and take note of the parse tree. Then, change the order of the rules $NP \rightarrow N$ and $NP \rightarrow NP PP$, so that the $NP PP$ expansion is used by the parser first. Run the parser on the above sentence.

Q3: What is the problem with this new ordering and how does the parser's behaviour differ from the other ordering? Explain how this behaviour depends on the particular way this RD parser chooses which rule to expand when there are multiple options.

Remove the offending rule before going on.

Ungrammatical sentences

Run the parser on following sentences:

```
he sneezed
he sneezed the fork
```

"sneeze" is an intransitive verb which doesn't allow an object. Though the second sentence is ungrammatical, it is parsed by our grammar.

Q4: Modify the grammar to handle such cases. What rules did you add/change? (*Hint*: Required *lexical productions* are already provided in the grammar)

Number agreement (optional)

If you're more interested in looking at real treebank rules, feel free to skip this and go to the next section.

Our toy grammar ignores the concept of number agreement. Change the grammar to handle number agreement by creating subcategories of the appropriate categories and use your new grammar to parse/rule out the following sentences, or others of your choosing:

```
he eats salad
he eat salad
he ate a mushrooms
```

Try an ambiguous sentence, such as `the sheep ran`. You can force the parser to look for additional parses after it finds a complete parse by hitting Autostep again. Now try `the sheep sneeze` and compare.

Exploring a treebank grammar (optional)

The previous parts of the lab dealt with a toy grammar. The remaining parts give you some idea of what a broad-coverage treebank and treebank-derived grammar looks like. We consider a small part of the Penn Phrase Structure Treebank. This data consists of around 3900 sentences, where each sentence is annotated with its phrase structure tree. Our code uses nltk libraries to load this data and extract parsed sentences.

Extracting and viewing parsed sentences

Make sure you've downloaded the corpus (see Preliminaries) and uncomment the following two lines in the lab code:

```
psents = treebank.parsed_sents()
print_parse_info(psents[0])
```

Then save and re-run the file.

It should print the first parsed sentence (at index 0), and the grammatical and lexical productions used in the first sentence.

The first line of these two extracts the list of parsed sentences from the treebank (which is imported at the top of the file), and `psents[0]` gives the first parsed sentence. When you `print` it (this happens inside `print_parse_info`), it shows up as a sort of left-to-right sideways tree, with parentheses around each node label and its children. Words are at the leaves, with each word dominated by a single pre-terminal label (POS category). These POS-word pairs are then grouped together into constituents labelled with nonterminal labels (phrasal categories).

You can also look at the parses and productions for other sentences by changing which sentence is passed in to `print_parse_info`. Verify this by looking at the parse and productions for the second sentence in the corpus.

What type of object is the parsed sentence object? (*Hint*: use the `type` function)

Check the methods available for this object using the command `help(psents[0])`. You can ignore the ones beginning with `'_'`, they are Python-internal ones.

Try using the `draw()` method on some of the sentences. Remember, `draw()` is a *method*, not a *function*, so the syntax is, e.g., `psents[0].draw()`.

Extract the list of words and the list of (word, pos-tag) tuples from `psents[0]` using some of the other available methods.

Distribution of Productions

Construct a frequency distribution of productions by completing the code in the `production_distribution` function, which returns two dictionaries, one for lexical productions and the other for grammatical productions.

Do you expect there to be more lexical or non-lexical (grammatical) productions in the grammar? Roughly how many non-lexical productions do you think there might be? Now check how many productions of each type there actually are. Do the numbers surprise you?

What are the 10 most frequent and least frequent lexical and grammatical productions? (*Hint*: you'll need to sort the frequency distributions as you did in the previous lab.)

By reusing code from a previous lab, plot histograms of the top 50-100 productions of each type. Do these distributions look familiar?

Subcategorization

Does the treebank grammar enforce subcategorization constraints such as number agreement or verb argument structure (e.g., transitive/intransitive/ditransitive)? If so, find some examples of rules that enforce these constraints. If not, find some examples of parses where you could replace one of the words in the parse with another word that is grammatical under the treebank grammar, but not in real English.