

NLP Lab Session Week 1

Word Frequencies from text in NLTK

September 2, 2015

Starting an NLTK Session

Python can be run on the command line to run Python programs (that have been written in a text editor or an IDE (Integrated Development Environment)) or as an interpreter, where you just type little pieces of Python on the interpreter line and it runs them for you. We will mostly be running Python in interpreter mode in an IDLE window.

For Windows:

If the Python IDLE icon is on your desktop, just double-click that to open the Idle window. Or to open the Python IDLE window, go to All Programs->Python 2.7 -> IDLE (Python GUI)

For Macs:

Start a terminal session. Type either “python” or “idle”.

It's ok if it's Python 2.6, but don't try to use any version of Python 3 with these lab notes.

Now download the file

LabWeek1examples.txt

and save it in a folder where you will keep materials for this class.

You will probably want to work by having the IDLE window open for testing NLTK and a browser window open with these instructions. You may also want to have a separate tab or window open to the NLTK book: http://www.nltk.org/book_1ed/, where these examples are taken from Chapter 1.

In the following, examples for you to try are given following the Python Idle prompt of >>>. You can copy and paste the Python example into the python interpreter, or you can type the example in.

Getting Started in Python and NLTK

Start by typing a couple of examples of **arithmetic** into the Python interpreter. For example, “1 + 2” is an arithmetic expression and Python will compute a result:

```
>>> 1 + 2
```

Note that if you want to type in a **string** of text, you surround the string with quotes.

```
>>> 'hello'
```

(In Python, you can usually also use double quotes. But note that whenever you copy/paste quotes from a Word or PDF document, you may get non-programming quotes, and you may have to type them directly from the keyboard.)

In programming, when we have a value of some type (like the number 3 or the string 'hello'), we can save that value by assigning it to a **variable**.

```
>>> num = 1 + 2
```

```
>>> num
```

In this example, the name of the variable is "num" and its value is 3.

Python and NLTK Resources

Python WikiBooks: http://en.wikibooks.org/wiki/Python_Programming

Python tutorial from python.org: <http://docs.python.org/2/tutorial/>

NLTK book: http://nltk.org/book_1ed/

After today's lab session, if you want to know more about Python variables, strings and lists, you can look at these topics in the resources.

Working with Text from the NLTK book module

Next, you use the Python "import" statement to load the data used in the book examples into the Python environment:

```
>>> from nltk.book import *
```

This command has loaded modules, functions and data from the nltk.book module. Among the loaded objects, it has loaded 9 of the text examples available from the corpora package (only a small number of them!). It has used the variable names text1 through text9 for these examples, and already assigned them values. If you type the variable name, you get a description of the text:

```
>>> text1
```

The variables sent1 through sent9 have been set to be a list of tokens of the first sentence of each text.

```
>>> sent1
```

Note that the first sentence of the book Moby Dick is "Call me Ishmael." and that this sentence has been already separated into tokens in the variable sent1. Compare the value of sent1 with the following variable:

```
>>> sent1text = 'Call me Ishmael.'
```

Searching Text

The text data structure has a number of functions to operate on text. One is called "concordance", and it will search for any word that you give to the function and show you the occurrences and some surrounding context.

```
>>> text1.concordance("monstrous")
```

Observe the use of the arrow keys with the enter key to select and modify previous lines in Python, and try a similar example.

```
>>> text2.concordance("affection")
```

Another function is "similar" which finds all the words that are used in the same context as the one given, where the context is the word before and the word after.

```
>>> text1.similar("monstrous")
```

We can use this to compare how the same word is used differently in other texts.

```
>>> text2.similar("monstrous")
```

Counting Vocabulary

Each text from the books was separated into a list of tokens, and this is one of the first NLP processing steps. The tokens usually consist of words and all the punctuation and other symbols occurring in the text. To further investigate text, we can count the occurrences of words.

We start by using the Python length function, "len" to tell us how many things are in a list. (Strictly speaking, each text variable is an object of type `nltk.text.Text`, which contains the text string and some other functions, but we're trying not to explain much programming here.) Even though the text variables are of type `nltk.text.Text`, the interface allows us to use the variables as if it were a list (this list is actually `text3.tokens`)

```
>>> len(text3)
```

```
>>> len(text4)
```

We can see that we don't want to print out all the tokens, but we can use the Python list "slice" notation to print the first 50 tokens:

```
>>> text3[:50]
```

Now this is the total number of tokens, and we might also want to find out how many unique words there are, not counting repetitions. The Python "set" function removes the repetitions, and we can apply the "sorted" function to that, returning the resulted sorted list of tokens. If we typed the following, lots of words would flash by on the screen.

```
XXX >>> sorted(set(text3))
```

Instead, we can just find the length of that list.

```
>>> len(sorted(set(text3)))
```

Or we can specify just to print the first 30 words in the list of sorted words:

```
>>> sorted(set(text3))[:30]
```

Now let's compute the ratio of the total number of tokens to the number of unique tokens and we'll get an average of how many repetitions there are for each word. First we get a

division operator that uses real arithmetic (aka floating point) instead of integer and then we divide to get the ratio.

```
>>> from __future__ import division
```

```
>>> len(text3) / len(set(text3))  
(On average, each word is used about 16 times.)
```

Now let's search for and count occurrences of particular words and compare that to the total number of words.

```
>>> text3.count("smote")
```

Compute the fraction of the number of occurrences of the word compared with the total number of words and then multiply by 100 to get a percentage.

```
>>> 100 * text3.count('smote') / len(text3)
```

How does this compare with a more common word, such as the word "a"?

```
>>> 100 * text3.count('a') / len(text3)
```

Try it Out:

1. How many times does the word "lol" occur in text5? What is the percentage of its occurrences in the text? [Warning: text5 is uncensored chat]

Think of another word to find occurrences and get the number of occurrences and its percentage in one of the texts. Save the word, the number of occurrences and its percentage in the text to post at the end of class.

Processing Text

In the first part of this lab, we counted words from text that had already been tokenized, i.e. separated into words. Now we'll look at some text examples that we will need to tokenize.

In addition to the examples that we imported for the NLTK book above, the NLTK has a number of other corpora, described in Chapter 2. In order to see these, type in

```
>>> import nltk
```

You can then view some books obtained from the Gutenberg on-line book project:

```
>>> nltk.corpus.gutenberg.fileids()
```

For purposes of this lab, we will work with the first book, Jane Austen's "Emma". First, we save the first fileid (number 0 in the list) into a variable named file1 so that we can reuse it:

```
>>> file0 = nltk.corpus.gutenberg.fileids() [0]
>>> file0
```

We can get the original text, using the raw function:

```
>>> emmatext = nltk.corpus.gutenberg.raw(file0)
>>> len(emmatext)
```

Since this is quite long, we can view part of it, e.g. the first 120 **characters**

```
>>> emmatext[:120]
```

NLTK has several tokenizers available to break the raw text into tokens; we will use one trained on news articles that separates by white space and by special characters (punctuation), but also leaves together some of these such as Mr.:

```
>>> emmatokens = nltk.word_tokenize(emmatext)
>>> len(emmatokens)
```

View the first 50 **tokens**

```
>>> emmatokens[:50]
```

We probably want to use the lowercase versions of the words. We define the variable `emmawords` using a Python list comprehension: “[w.lower() for w in emmatokens]”, where this expression is the list consisting of the all the items that are the lowercase of w “w.lower()” for all w that are in emmatokens.

```
>>> emmawords = [w.lower() for w in emmatokens]
>>> emmawords[:50]
>>> len(emmawords)
```

We can further view the words by getting the unique words and sorting them:

```
>>> emmavocab = sorted(set(emmawords))
>>> emmavocab[:50]
```

We can see that we will probably want to get rid of these special characters – Regular Expressions to the Rescue! (as in `xkcd_`), but we’ll work on that next week.

[Side note on Python: Here we defined `emmawords` by using a Python list comprehension. It would be equivalent to define `emmawords` by starting with an empty list and using a “for” loop that kept appending the lower case of each word in `emmatokens` .

Note that a Python “for” loop can be used to iterate over every element of a list.

Note the special syntax of Python for a multi-line statement: The first line must be followed by an extra “:”, and each succeeding line must be indented by some number of spaces (as long as they are all indented by the **same** number of spaces.) Just hit enter to finish the statement.

Start with an empty list.

```
>>> emmawords = []
```

For each word *w* in the list *emmatokens*, add the lower case of that word to the *emmawords* list.

```
>>> for w in emmatokens:  
    emmawords.append(w.lower())
```

```
>>> emmawords[:50]
```

End note.]

Try it out

2. Think of a word to find occurrences in the text for the book *Emma* and get the number of occurrences and its percentage in the text. Note that the list *emmatokens* is the list from which you want to count the number of occurrences of the word. Save the word, the number of occurrences and its percentage in the text to post at the end of class.

Exercise to submit this week:

Go to Blackboard and find the Discussion for the first week exercises. Create a post in which you:

- Make the title of the post contain the two words that you chose in the two “try it out” sections.
- In the body of the post, state your text, word, frequency and percentage that you got in the two “try it out” sections.

Reading Assignment: Python Variables, Strings and Lists

NLTK book: http://nltk.org/book_1ed/

Read the sections called “Indexing Lists”, “Variables”, and “Strings”.

Or

Python WikiBooks: http://en.wikibooks.org/wiki/Python_Programming

For more detail on these topics, read the sections on Variables and Strings, and on Lists, but note that these section have more detail than you need to know for this course.