

NLP Lab Session Week 3
September 10, 2014

Starting a Python and an NLTK Session

Open a Python 2.7 IDLE (Python GUI) window by going to All Programs->Python 2.7 -> IDLE (Python GUI).

In this lab session, we will work together through a series of small examples using the IDLE window and that will be described in this lab document. However, for purposes of using cut-and-paste to put examples into IDLE, the examples can also be found in a set of python files on the iLMS system, under Resources. I loaded this file both with the extension .py and .txt to see which one was best for you to download from Blackboard.

Download either: LabWeek3examples.py or LabWeek3examples.txt and save it in a folder on your H: drive where you will keep materials for this class. If you downloaded the one with extension .txt, rename it to have .py.

Open an IDLE window. Use the File-> Open to open the labweek3examples.py file. This should start another IDLE window with the program in it. Each example line can be cut-and-paste to the IDLE window to try it out.

Getting Started in Python and NLTK

First let's load nltk.

```
>>> import nltk
```

Next, we want to set up the emma text again for processing. Go to the file with LabWeek3examples.py and copy and paste the following lines to get started. This gets the text of the book Emma, separates it into tokens with the wordpunct tokenizer, and converts all the characters to lower case.

Notes: Paste one line at a time! Any line that starts with the character # is a python comment and doesn't need to be pasted. (This can be any text that explains to a person what is going on. Python doesn't try to process it.)

```
>>> from nltk import FreqDist
>>> print nltk.corpus.gutenberg.fileids()
>>> file0 = nltk.corpus.gutenberg.fileids()[0]
>>> emmatext = nltk.corpus.gutenberg.raw(file0)
>>> emmatokens = nltk.word_tokenize(emmatext)
>>> emmawords = [w.lower() for w in emmatokens]
```

For purposes of the lab, we create a list with only the first 101 words that follow the title and author.

```
>>> shortwords = emmawords[11:111]
>>> shortwords
```

As before, we create a frequency distribution of the words, using the NLTK FreqDist module/class.

```
>>> shortdist = FreqDist(shortwords)
>>> shortdist.keys()

>>> for word in shortdist.keys():
    print word, shortdist[word]
```

Again 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.)

Digression on sorting dictionaries and frequency distributions by value

In Python, we can sort dictionaries in order of their values in a couple of different ways. First let's create a dictionary and use the functions keys() and items() to look at the list of keys and the list of [key, value] pairs using the function items().

```
>>> d = {'a': 13, 'b': 4, 'c': 9}
>>> d.keys()
>>> d.items()
```

The python “sorted” function takes a list of values and returns a sorted list. If the list has more complex values, like our [key, value] pairs, then it also allows us to give a function that tells which parts of the list to sort by. So to sort by the values, there is a special function called itemgetter() that lets you get either the first part (index 0) or the second part (index 1) of the pair.

```
>>> from operator import itemgetter
>>> sorted(d.items(), key=itemgetter(1))
```

Another syntax that you often see to sort by the values of a dictionary, used the lambda notation to create a function that picks out the second part of the pair.

```
>>> sorted(d.items(), key=lambda x:x[1])
```

What about the NLTK Frequency Distributions? If you have made a frequency distribution with FreqDist, then NLTK allows you to treat it like a dictionary, where you can use the keys() and items() function, and then you can sort it like any python dictionary. And the sorted function allows you to put a keyword argument that specifies descending order by giving “reverse=True”).

```
>>> from nltk import FreqDist
>>> word_fd = FreqDist(shortwords)
>>> word_fd.keys()
>>> word_fd.items()[:10]
>>> sorted(word_fd.items(), key=itemgetter(1))
```

More Specialized Frequency Distributions (What is a word?)

We note that the word tokenization produces tokens that have special characters in them. Let's remove all the tokens that have special characters only special characters. Notes that this leaves some special characters, such as words with embedded hypens, for example "lower-case" and "need-based".

We'll use a regular expression that matches any token that contains all non-alphabetical character. We'll be covering regular expressions in class next week, but for now, we can just use this one.

```
>>> import re
# this regular expression pattern matches any word that contains all non-alphabetical
# lower-case characters [^a-z]
>>> pattern = re.compile('^[^a-z]+$')
```

Apply the pattern to the string '**' to see if it matches. Note that the result of the .match function can be used as a Boolean expression, either true or false, so you can use it in an "if" test.

```
>>> nonAlphaMatch = pattern.match('**')
```

```
# if it matched, print a message
>>> if nonAlphaMatch: 'matched non-alphabetical'
```

In Python, we can make a function that can take any argument, process it and return a result. For an example function, we make a function called "alpha_filter" that takes word as an argument and returns True if it contains all non-alphabetical characters and False otherwise.

Copy and paste this entire function definition into the Idle window.

```
>>> # function that takes a word and returns true if it does not consist solely
# of non-alphabetic words

def alpha_filter(w):
    # define the regular expression pattern to match a sequence of non-alphabetical tokens
    pattern = re.compile('^[^a-z]+$')
    if (pattern.match(w)):
        return True
```

```
else:  
    return False
```

Apply the new function to shortwords to include only those that don't match the filter:

```
>>> alphashortwords = [w for w in shortwords if not alpha_filter(w)]  
>>> alphashortwords
```

Our next step is to remove some of the common words that appear with great frequency. This is usually done by making a list of the words to remove, known as a *stop word* list. Every token is compared with the stop word list and not entered into the frequency distribution if it appears on the list.

For purposes of the lab, we'll use the stop word list given by NLTK. In the future, we'll use a list of stopwords from a file if you want to design your own custom stop word list. Note that what the stopword list should be will depend on the analysis that is using the word frequency list. An example is the inclusion of pronouns; if you are looking at the frequencies of topic words, you remove pronouns, but if you are looking at words contributing to literary style, you would include them.

```
>>> stopwords = nltk.corpus.stopwords.words('english')  
>>> len(stopwords)
```

Test if a word is in a list by using the Python keyword "in":

```
>>> word = 'the'  
>>> if word in stopwords:  
    print 'Stop!'
```

Now we define a function to make a frequency distribution from a list of tokens that has no tokens that contain non-alphabetical characters or words in the stopword list. We will pass the stop word list into the function as a second argument. (Now the function is defined to take two arguments and return one result, which is a frequency distribution.)

We can also filter our shortwords list to not include any stopwords that we defined above.

```
>>> stoppedshortwords = [w for w in shortwords if not w in stopwords]  
>>> stoppedshortwords
```

Note that this big stop word list removes a lot of the non content-bearing words in the list.

Bigram Frequency Distributions

Another way to look for interesting characterizations of a corpus is to look at pairs of words that are frequently collocated, that is, they occur in a sequence called a bigram.

First, we just simply look at the bigrams that can be defined.

```
>>> nltk.bigrams(shortwords)
```

But we will obtain a lot more functionality by using the functions from the `nltk.collocations` package. One of the places to obtain information about this package is in this section of the NLTK documentation called the HOWTOS. We note that there are also Trigram functions in addition to the Bigram functions shown here.

<http://www.nltk.org/howto/collocations.html>

We start by making an object called a `BigramCollocationFinder`. The finder then allows us to call other functions to filter the bigrams that it collected and to give scores to the bigrams. We start by scoring the bigrams by frequency:

```
>>> finder = BigramCollocationFinder.from_words(shortwords)
>>> scored = finder.score_ngrams(bigram_measures.raw_freq)
>>> scored
```

For any finder, we can also apply various filter functions. First let's apply our `alpha_filter` that we created earlier. It uses a filter that is applied to the individual words.

```
>>> finder.apply_word_filter(alpha_filter)
>>> scored = finder.score_ngrams(bigram_measures.raw_freq)
>>> scored
```

Next we can filter out the stopwords if we wish.

```
>>> finder.apply_word_filter(lambda w: w in stopwords)
>>> scored = finder.score_ngrams(bigram_measures.raw_freq)
>>> scored
```

Another filter would remove words that only occurred with a frequency over some minimum threshold.

But to show its effect, we first start over with a newly defined finder.

```
>>> finder = BigramCollocationFinder.from_words(shortwords)
>>> finder.apply_freq_filter(2)
>>> scored = finder.score_ngrams(bigram_measures.raw_freq)
>>> scored
```

Another word level filter would be to remove words whose length is small. Or we can apply filters that work on both words of the bigram. Suppose that we want to filter out bigrams in which the first word's length is less than 2. Then we could use an ngram filter that is applied to both words of the bigram.

```
>>> finder.apply_ngram_filter(lambda w1, w2: len(w1) < 2)
>>> scored = finder.score_ngrams(bigram_measures.raw_freq)
```

```
>>> scored
```

Mutual Information and other scorers

Look at the paper by Church and Hanks and observe their definition of the Association Ratio. Note that technically the original information theoretic definition of mutual information allows the two words to be in either order, but that the association ratio defined by Church and Hanks requires the words to be in order from left to right wherever they appear in the window

In NLTK, the mutual information score is given by a function for Pointwise Mutual Information, where this is the version with out the window.

```
>>> finder = BigramCollocationFinder.from_words(shortwords)
>>> scored = finder.score_ngrams(bigram_measures.pmi)
>>> scored
```

See the documentation for other scoring functions.

Bigrams from the entire Emma book

Finally, let's put together our work on bigrams to display bigrams from the entire emma book. We'll repeat many of our steps for emmawords, instead of shortwords.

These two scorers show the bigrams in order by frequency.

```
>>> finder = BigramCollocationFinder.from_words(emmawords)
>>> scored = finder.score_ngrams(bigram_measures.raw_freq)
>>> len(scored)
67625
>>> sorted(scored, reverse=True)[:30]
```

Apply the non-alphabetical and stopword filters.

```
>>> finder.apply_word_filter(alpha_filter)
>>> finder.apply_word_filter(lambda w: w in stopwords)
>>> scored = finder.score_ngrams(bigram_measures.raw_freq)
>>> len(scored)
16070
>>> sorted(scored, reverse=True)[:30]
```

This scorer shows the bigrams in order by Pointwise Mutual Information.

```
>>> scored = finder.score_ngrams(bigram_measures.pmi)
>>> len(scored)
16070
```

```
>>> sorted(scored, reverse=True)[:30]
```

Exercise for Week 3:

For this exercise, you may work in groups of 2-3, or you can choose to do your own work.

- Choose a file that you want to work on, either one of the files from the book corpus, or one from the Gutenberg corpus.
-
- Make a bigram finder and experiment with whether to apply the filters or not. Run the scoring with both the raw frequency and the pmi scorers and compare results.

To complete the exercise, choose one of your top 20 frequency lists to report to show to the class. Write an introductory sentence of paragraph telling what text you chose and what bigram functions you used. Put this and the frequency list in a discussion posting in the blackboard system under the Discussions tab.