

NLP Lab Session Week 5
September 29, 2011

POS taggers in NLTK

Installing NLTK data

If needed, do an `nltk.download()` and choose `nltk_book?` and save on the H: drive if you have space. Or, if you prefer, I can give you the dataset on a memory stick.

Getting Started

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 python file on the iLMS system, under Resources.

Labweek5examples.py

Open an IDLE window. Use the File-> Open to open the labweek4examples.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.

Set up a tagged corpus

In the last session, we looked at both the Brown corpus and the Penn Treebank corpus; we will work with the Penn Treebank corpus in this session to learn POS tag frequencies.

```
import nltk
from nltk.corpus import treebank
```

Optionally, we can look at some of the tagged words from the corpus.

```
treebank.tagged_words()[:50]
len(treebank.tagged_words())
```

POS Tagging

The process of classifying words into their parts of speech and labeling them accordingly is known as part-of-speech tagging, POS-tagging, or simply tagging. Parts of speech are also known as word classes or lexical categories. The collection of tags used for a particular task is known as a tagset. We will use two different parsers to tagging text.

First we initialize some variables from NLTK, the Penn Treebank. The first has the tagged sentences. We will use those in training later taggers and to compare for the

evaluation. The second just has the words (untagged text) that we will use to test our taggers.

```
treebank_tagged = treebank.tagged_sents()
treebank_text = treebank.words()
len(treebank_text)
treebank_text[:50]
```

To introduce the N-gram taggers in NLTK, we start with a default tagger that just tags everything with the most frequent tag: NN. We create the tagger and run it on text.

```
default_tagger = nltk.DefaultTagger('NN')
default_tagger.tag(treebank_text[:50])
```

The NLTK includes a function for taggers that computes tagging accuracy by comparing the result of a tagger with the original “gold standard” tagged text. Here we use the NLTK function “evaluate” to apply the default tagger (to the untagged text) and compare it with the gold standard tagged text.

```
default_tagger.evaluate(treebank_tagged)
```

Other simple taggers described in the NLTK book are the Regular Expression Tagger and the Lookup Tagger.

Next we train a Unigram tagger. It tags each word with the most frequent tag in the corpus. First we compare it to the same corpus that it learned from.

```
unigram_tagger = nltk.UnigramTagger(treebank_tagged)
unigram_tagger.tag(treebank_text[:50])
```

Next, we do separate the tagged data into a training set and a test set. This allows us to test the tagger’s accuracy on similar, but not the same, data that it was trained on. We take the first 90% of the data for the training set, and the remaining 10% for the test set.

```
size = int(len(treebank_tagged) * 0.9)
treebank_train = treebank_tagged[:size]
treebank_test = treebank_tagged[size:]
unigram_tagger = nltk.UnigramTagger(treebank_train)
unigram_tagger.evaluate(treebank_test)
```

Finally, NLTK has a Bigram tagger that can be trained using 2 tag-word sequences. But there will be unknown frequencies in the test data for the bigram tagger, and unknown words for the unigram tagger, so we can use the backoff tagger capability of NLTK to create a combined tagger. This tagger uses bigram frequencies to tag as much as possible. If a word doesn’t occur in a bigram, it uses the unigram tagger to tag that

word. If the word is unknown to the unigram tagger, then we use the default tagger to tag it as 'NN'.

```
t0 = nltk.DefaultTagger('NN')
t1 = nltk.UnigramTagger(treebank_train, backoff=t0)
t2 = nltk.BigramTagger(treebank_train, backoff=t1)
t2.evaluate(treebank_test)
```

Let's use this tagger to tag some example text. Define some example text, tokenize it, and apply the tagger.

```
text = "Three Calgarians have found a rather unusual way of leaving snow and ice
behind. They set off this week on foot and by camel on a grueling trek across the burning
Arabian desert."
tokens = nltk.wordpunct_tokenize(text)
taggedtext = t2.tag(tokens)
taggedtext
```

In order to look at some of the tags in a text, we can use the `findtags` function from the last session. (Note that a conditional frequency distribution first has a set of keys called the "conditions" that indexes a regular frequency distribution (with keys and values).)

```
def findtags(tag_prefix, tagged_text):
    cfd = nltk.ConditionalFreqDist((tag, word) for (word, tag) in tagged_text
                                   if tag.startswith(tag_prefix))
    return dict((tag, cfd[tag].keys()[:20]) for tag in cfd.conditions())
```

```
tagdictNN = findtags('NN', taggedtext)
```

```
#Find the most frequent words in this example that have one of the noun tags
for tag in sorted(tagdictNN):
    print tag, tagdictNN[tag]
```

Exercise:

For this exercise, you are to choose any document, including any of the untagged documents in the Gutenberg corpus, run the bigram POS tagger, and find the top frequency words by POS tag class. As usual, you may work in groups.

(There is one caveat on using the bigram POS tagger on these texts, which is that it was trained on text which had been separated by sentences. So we should really run a sentence detector on our raw text first, and we'll talk about how to do that in the future.)

1. Choose a document, as we did for the word frequencies, and tokenize it. Do not lowercase the words or stem them. For example, if I chose emma, the first document in

the Gutenberg corpus, I would get the raw text and tokenize it with wordpunct, ending up with something named emmatokens.

2. Now run the bigram pos tagger - it was named t2 in our development – and save the results in a variable (don't try to print the whole thing). For example, I would save mine in a variable called emmatagged.
3. Use the findtags function to define 4 dictionaries for nouns (prefix 'N'), verbs (prefix 'V'), adjectives (prefix 'J') and adverbs (prefix 'R'). Display these dictionaries.
4. Choose the results of one of your dictionaries and post it to the discussion for today's lab. Include the names of everyone in your group in the discussion title. The discussion entry should include the name of your document and what type of words you are giving.