

# CS1114 Assignment 6

out: Saturday, April 18, 2009.  
due: **Friday, April 25, 2009, 5PM.**

---

In this assignment, you will be implementing an authorship detector which, when given a large sample size of text to train on, can then guess the author of an unknown text.

As usual, you can find the stub files in the directory `~cs1114/student_files/A6/`. There is also training and testing data in the `TrainingData` and `TestData` subdirectories.

## 1 Detecting authorship

The algorithm you will implement works based on the following idea: An author's writing style can be defined quantitatively by looking at the words he uses. Specifically, we want to keep track of his word flow - that is, which words he tends to use after other words.

To make things significantly simpler, we're going to assume that the author always follows a given word with the same distribution of words. Of course, this isn't true, since the words you choose when writing obviously depend on context. Nevertheless, this simplifying assumption should hold over an extended amount of text, where context becomes less relevant.

In order to implement this model of an author's writing style, we will use a Markov chain. A Markov chain is a set of states with the Markov property - that is, the probabilities of each state are independent from the probabilities of every other state. This behavior correctly models our assumption of word independence.

A Markov chain can be represented as a directed graph. Each node is a state (words, in our case), and a directed edge going from state  $S_i$  to  $S_j$  represents the probability we will go to  $S_j$  when we're at  $S_i$ . We will implement this directed graph as a transition matrix. Given a set of words  $W_1, W_2, \dots, W_n$ , we can construct an  $n$  by  $n$  transition matrix  $A$ , where an edge from  $W_i$  to  $W_j$  of weight  $p$  means  $A_{ij} = p$ .

The edges, in this case, represent the probability that word  $j$  follows word  $i$  from the given author. This means, of course, that the sum of the weights of all edges leaving from each word must add up to 1.

We can construct this graph from a large sample text corpus. Our next step would be finding the author of an unknown, short chunk of text. To do this, we simply compute the probability of this unknown text occurring, using the words in that order, in each of our Markov chains. The author would likely be the one with the highest probability.

You will implement the Markov chain model of writing style. We have given you some sample texts for you to train your model on, as well as some challenges for you to figure out.

## 2 Constructing the transition matrix

Your first step is to construct the transition matrix representing our Markov chain. First, you must read the text from a sample file. To do so, you can use the Matlab command `fopen` to open a file and `textscan` to read the text into a cell array (making sure to `fclose` the file once you're done with it).

```
f = fopen('filename.txt', 'r');
cell = textscan(f, '%s');
fclose(f);
D = C{1};
```

The cell array `D` above will contain all the words read by `textscan` in the format of strings.

Once you have the cell array of words, you must count the number of unique words appearing in the text corpus. You can do so with the use of a hash table, as we covered in section. Recall that you can create a Java hash table using the following Matlab code:

```
hashtable = java.util.Hashtable;
```

Recall that an entry into a hash table has a *key* and a *value*—the key is used as the index into the hash table, and the value is what gets stored there (for this class, we are not worried about how a hash table is implemented, although this is a very interesting topic in itself). With hash tables in Matlab (which, confusingly, are really Java hash tables, but this is also a topic for another day), you can add an entry using the `put` function, and retrieve an entry with the `get` function. Getting a key which hasn't been previously added returns an empty matrix:

```
hashtable = java.util.Hashtable;
hashtable.put('llama', 10);
hashtable.get('llama'); % Returns 10
hashtable.get('platypus'); % Returns []
```

For this assignment, we would like you to create two hash tables: one will map each word (encoded as a string) to a location (row/column index) in your transition matrix (we will call this the *dictionary* for the transition matrix), and the other will be the inverse mapping from the word's location to the word itself.

Once this is done, you can construct a transition matrix from the words in your cell array. You will want to create a sparse array using the Matlab `sparse` function. Along with the transition matrix, you will be creating a corresponding vector that contains

word frequencies (normalized by the total number of words in the document (including repeated words)). For instance, if 5% of the words in the document are `and`, the entry of this vector corresponding to the word `and` would be 0.05.

You will write this in a function called `build_transition_matrix`, which takes in a filename, and returns the transition matrix, the dictionary, the inverse dictionary, and the histogram of words.

⇒ Write the function `build_transition_matrix.m`.

### 3 Calculating likelihood

Once you have your transition matrix, you can calculate the likelihood of an unknown sample of text. We have included for you several pieces of literature by various authors, as well as excerpts from each of the authors. Your goal is to identify the authors of each excerpt.

To do so, you will need to calculate the likelihood of the excerpt occurring in each author's transition matrix. Recall that each edge in the directed graph that the transition matrix represents is the probability that the author follows a word with another. You can load the text of the excerpt into Matlab using `fopen` and `textscan` in the same way as before.

Since you will be multiplying numerous possibly small probabilities together, your calculated likelihood will likely be extremely small. Thus, you should compare  $\log(\text{likelihood})$  instead. Keep in mind the possibility that the author may have used a word he has never used before. Your calculated likelihood should not eliminate an author completely because of this.

⇒ Write the function `compute_text_likelihood.m`.

### 5 Finding the author with the maximum likelihood

Now that you can compute likelihoods, the next step is to write a routine that takes a set of transition matrices and dictionaries (represented as cell arrays), and a sequence of text (again, represented as a cell array), and returns the index of the transition matrix that results in the highest likelihood. You will write this in a function `classify_text`, which takes a cell array `tmatrices`, a cell array `dictionaries`, a cell array `histograms`, and the name of the file containing the test text, and returns a single integer `best_index`. This will allow you to test your code using our routines—we will be using this to test your code on new authors.

⇒ Write the function `classify_text`.

### 6 Testing your authorship prediction

Now that you can calculate the likelihood that a text is written by an author, take a look at the texts we have provided you. In the folder `TrainingData`, you will find literature by several authors, taken from Project Gutenberg. In the folder `TestData`, you will find five unknown excerpts from each of the authors. Create a test function which creates author models using the training data, and predicts the authors of the test data. Note that in addition to particular works, in each author's directory there is a file `all.txt` that contains all of the words in a single file.

⇒ Write a series of tests that adequately demonstrates the operation of your code. Also document any tests you have done, and the results you got.

## 7 Challenges

If you have finished that, you may be interested in a harder challenge. Shortly after the founding of the United States, various founding fathers collectively published a series of 85 essays called the Federalist Papers under the pseudonym “Publius.” It is generally accepted that the papers were written by Alexander Hamilton, John Jay, and James Madison. While historians have generally agreed on who wrote which essay, we would like to corroborate it with evidence from your analysis. In the folder `TrainingData`, you will find sample writing from Hamilton, Jay, and Madison. In the folder `TestData`, you will find the first 20 Federalist Papers. Can you identify the authors of each? It may be possible that two people collaborated on a paper. Can you tell? You will get one extra credit point for writing code that tests on this data set.

There are many possible ways to improve on the method described above. You will get an additional extra credit point for implementing an improved method for predicting authorship. This part of the assignment is completely open-ended—you are free to brainstorm your own method, but you must explain what you did to get the extra credit.

Finally, you will be turning in your code for this assignment, and we will be testing your code on a completely different data set (which you won't have access too). The code that performs the best in our authorship tests will be awarded two extra credit points (note that you will likely need to write an improved algorithm to get better results than the baseline method described in Sections 1-6).