Task 2: Naive Bayes Classifier

In the previous task we only used selected words for the classification, and we relied on the manual work that went into the lexicon to decide which words were likely to be in reviews with a particular sentiment. The sentiment lexicon is general: it is supposed to be applicable to any English text. We will now create a statistical classifier, learning which words are positive and negative in movie reviews. Unlike the lexicon approach, this classifier uses **all** of the words in the text.

Step 0: Data preparation.

Use the dataset from Task 1 for today's training and development. You should split the 1800 document dataset from Task 1 so there are 200 documents in your development set (balanced: 100 positive, 100 negative) and 1600 in the training set (also balanced). There are a further 200 reviews that you do not have access to at the moment (**held-out** data) that you will use for more formal testing and evaluation in a subsequent session. Code to make the split has been provided in `DataSplit.java` (note that this code refers to the developement data as the `validationSet`).

Having made the split, you can rerun the `simpleClassifer` from Task 1 on the 200 examples in the development set. Later you can see whether you get an improvement with the Naive Bayes classifier you implement today. N.B. everything you need to implement this step can be found within `Exercise2Tester.java`. Familiarise yourself with this class first.

Step 1: Parameter estimation.

Your first task in developing the Naive Bayes classifier is to estimate probabilities from observations of the data, namely: the proabilities of the classes in the dataset $P(\text{POS})$, $P(\text{NEG})$; and the probabilities of each word in the dataset appearing with a particular class label $P(\text{word}|\text{POS})$ and $P(\text{word}|\text{NEG})$. See pg 5 of Jurafsky and Martin's draft chapter https://web.stanford.edu/~jurafsky/slp3/4.pdf for more information.

Write a program that estimates these probabilities from the training data. Then compute the logs of these probabilities.

Step 2: Classification.

Now use the probabilities to apply the classification argmax formula of Naive Bayes to the development set:

$$c_{NB} = \underset{c \in POS, NEG}{argmax} \{logP(c) + \sum_{i \in positions} logP(w_i|c)\}$$

positions is the set of indexes into the all words in the document. $P(c)$ is the probability of a document belonging to a given class. For a balanced data set, like the one here, this probability is the same for each class but this will not always be the case. The argmax formula gives you a decision (a classification) for each document, either negative or positive.

How well did this program perform? As before, we can compare the system's performance to the truth but we should use only the development set to do this.

Step 3: Smoothing.

Have you noticed any issues when using the log probabilities calculated in Step 1?

When using a Naive Bayes classifier, you need to consider what to do with unseen words – words which occur in the development/test data but which do not occur in the training data at all. You will also need to consider what to do with words which were seen only in one class of document, either positive or negative ones. Write down possible solutions to this problem.

Modify your calculation of log probabilities to implement add-one (Laplace) smoothing – add one to all the counts.

Example: Say a word *jabberwocky* appeared in your training set only once, in a positive review. These are the counts you should record.

	positive	negative
unsmoothed	1	0
add-one	2	1

How does your classifier perform now?

Once you have successfully developed and tested a system which incorporates add-one smoothing, you may submit your code to the online tester. Remember that, even if you do not base your code on the Exercise2Tester.java, you will find it useful to look at that to see what we expect you to submit. Once your code has passed, you may contact a demonstrator to obtain Tick 2.