

CS 221 Information Retrieval
Assignment 1 - Text Processing Functions
Due date: 1/20

This assignment is to be done in groups of 2. You cannot use code written by your classmates. Use code found over the Internet at your own peril -- it may not do exactly what the assignment requests. If you do end up using code you find on the Internet, you must disclose the origin of the code. **Concealing the origin of a piece of code is plagiarism.** Use the Message Board for general questions whose answers can benefit you and everyone.

Project skeleton:

<http://www.ics.uci.edu/~lopes/teaching/cs221W13/assignments/Assignment1.zip>

General Specifications

1. You can use any programming language, but Java is strongly encouraged. This spec has all sorts of helpers for Java. Also, the next homework will use a Java crawler, so you may want to use this homework to brush up your knowledge of the language.
2. If you use Java, your solution must fill out the program skeleton provided.
 - a. Fill in each method according to its Javadoc specification.
 - b. Feel free to create additional methods / classes where necessary
3. If you don't use Java, you should produce a similar skeleton to start with and fill it out. You should also be very precise with instructions for how to run your program – what programs are needed, versions, etc. If the grader can't run your program, your grade will reflect that.
4. We will test your program with our own text files.
5. At points, the assignment may be underspecified. In those cases, make your own assumptions and document them.

Part A: Utilities (20 points)

Write a method that reads in a text file and returns a list of the tokens in that file.
Write a method to print out frequency results.

Package: `ir.assignments.two.a`
File: `Utilities.java`
Method: `tokenizeFile(File)`
Method: `printFrequencies(List<Frequency>)`

Part B: Word Frequencies (20 points)

Count the total number of words and their frequencies in a token list.

Package: `ir.assignments.two.b`
File: `WordFrequencyCounter.java`
Method: `computeWordFrequencies (ArrayList<String>)`

Part C: 2-grams (30 points)

A 2-gram is two words that occur consecutively in a file. For example, *two words*, *words that* and *that occur* are all 2-grams from the previous sentence.

Count the total number of 2-grams and their frequencies in a token list.

Package: `ir.assignments.two.c`
File: `TwoGramFrequencyCounter.java`
Method: `computeTwoGramFrequencies (ArrayList<String>)`

Part D: Palindromes (30 points)

A palindrome is a words or phrase that read the same in both directions. For example: *eye* is a palindrome and so is *Do geese see god*.

Count the total number of palindromes and their frequencies in a text file.

Package: `ir.assignments.two.d`
File: `PalindromeFrequencyCounter.java`
Method: `computePalindromeFrequencies (ArrayList<String>)`

Once you have implemented your palindrome counting algorithm, please perform a short analysis of its runtime complexity (does it run in linear time relative to the size of the input? Polynomial time? Exponential time?) This analysis should go in the **analysis.txt** file in this package.

Submitting Your Assignment

Your submission should be a single zip file submitted to EEE. This zip file should match the skeleton zip file provided with the assignment, with the addition of your implementations of the four sections. If there is anything you wish to communicate to the TA, such as implementation assumptions made, or how to run your program, this should be placed into the README.txt file provided in the skeleton zip file.

Evaluation Criteria

Your assignment will be graded on the following four criteria.

1. Correctness
 - a) How well does the behavior of the program match the specification?

- b) How does your program handle bad input?
- 2. Efficiency
 - a) How quickly does the program work on large inputs?
- 3. Aesthetics
 - a) Is the program clearly documented and well written?

For this assignment, there will not be face-to-face meetings with the grader.