ENEE 757 – Security Analytics Homework

Homework Due: 19 October 2015 at 11 am.

Submission Instructions: Write two Python programs, following the instructions below. Submit them from the GRACE machines, along with the corresponding output files, using the following commands:

```
$CLASS/submit 2015 fall AAAA BBBB 0101 2 clustering_problem.zip
$CLASS/submit 2015 fall AAAA BBBB 0101 2 spark_problem.zip
```

where `$CLASS` is `/afs/glue.umd.edu/class/fall2015/enee/757/0101/bin/`. You must replace `AAAA BBBB` with your own college and course number (`enee 757` or `cmsc 818v`).

1 Homework overview

The learning objective of this homework is for students to gain first-hand experience with some data analytics techniques that are commonly used to solve security problems. Specifically, *clustering algorithms* allow you to group data items into clusters of similar items, and expose the salient patterns in the data. For example, clustering algorithms may be applied to the entries in a system log to identify normal and anomalous behaviors, without having to specify manually what these behaviors should look like. However, the standard clustering algorithms do not scale well to high dimensional spaces and large data volumes. *Locality sensitive hashing* (LSH) is a technique for approximate clustering and nearest-neighbor search. For example, locality sensitive hashing may be applied to streaming Twitter posts to identify posts that are similar to a corpus of documents containing exploit code. The *Spark* data analytics platform allows you to perform some of these operations efficiently at scale.

2 Initial setup

Use the same Ubuntu VM as in the first homework. The VM includes all the tools you need for this homework. If you need to download the VM again, you can find it here:

`http://www.umiacs.umd.edu/~tdumitra/courses/ENEE757/Fall15/homeworks.html`

This is the machine I will use for testing your submissions. If your submission doesn't work on that machine, you will get no points. It makes no difference if your submission works on another Ubuntu version (or another OS).

You can find links to several tutorials at the address above; read these documents if youre stuck. I also encourage you to ask questions on our Piazza message board.

Starter files. Starter files are available on the class web page:

`http://www.umiacs.umd.edu/~tdumitra/courses/ENEE757/Fall15/homeworks.html`

3 Task 1: Identifying salient program behaviors with cluster analysis

The materials you need for this task are provided in the starter files, under the `unsupervised_learning` directory.

In this task, you are given a sample of 3000 machines, monitored over a one-week period. The data is in a comma-separated file, called `Error_Machines.csv`. Each row in this file corresponds to a machine, and has 11 columns:

- Column 1: Count of application crashes during the sample period.
- Columns 2–11: How many times the top 10 applications were launched during the sample period.

Your task is to cluster the machines based on the application usage. Note that it is difficult to determine in advance how many clusters the data will have, so you should experiment with different numbers of clusters.

TF-IDF. In many clustering applications, it is a good idea to start by selecting the most useful features before applying a clustering algorithm. One way to do this is to compute the *term frequency-inverse document frequency* (TF-IDF) values of the application instances. The TF-IDF value of an application App_i on a given machine M_j is the usage frequency of the application on that machine (term frequency) multiplied by the logarithmically scaled fraction of the machines where the application is present:

$$\text{tf-idf}(App_i, M_j) = \text{F}_{M_j}(App_i) \times \log \frac{|M|}{|\{M_k : App_i \text{ runs on } M_k\}|}$$

where M is the set of all machines and $\text{F}_{M_j}(App_i)$ is the usage frequency of app App_i on machine M_j.

The intuition behind this formula is that the term frequency will highlight the applications that are most frequently used on a machine, but does not distinguish between the frequent applications that appear on all the machines (and that are not useful for clustering) and the frequent applications that appear on a small set of machines. We therefore weight this value with the inverse

document frequency, which measures whether an application is frequent or rare across all the machines. The resulting TF-IDF values allow us to highlight the applications that are specific to each machine.

Clustering. There are several widely used clustering algorithms. For this homework, your task is to apply *K-Means clustering* and *agglomerative hierarchical clustering* to feature vectors with the TF-IDF values of application usage. You will also evaluate the clustering results by computing the *silhouette score*.

Starter files. The starter files include the `Error_Machines.csv` data file and a partially completed Python program called `clustering_exercise.py`. Complete this program by adding the appropriate code to replace all the `None` values. You should program in Python and use the functions provided by the Scikit-learn library (`http://scikit-learn.org/`).

Submitting. Create an archive, called `clustering_problem.zip`, which contains the completed `clustering_exercise.py`. The starter files include a shell script, called `pack_files.sh`, that will create this archive for you.

Submit the `clustering_problem.zip` file as described at the beginning of this handout.

4 Task 2: Evaluating document similarity in real time

The materials you need for this task are provided in the starter files, under the `locality_sensitive_hashing` directory.

In this task, you are given a pre-existing model built from a set of documents of interest (e.g. documents that include exploit code), and you must check, in real time, which of the documents from an incoming stream of data are most similar to the documents of interest. A document D is represented as a set of the words $W_D = \{w_1, w_2, ...\}$ that appear in the document.

Locality Sensitive Hashing. A standard measure of the similarity between two documents D_1 and D_2 is the Jaccard index $J(D_1, D_2) = \frac{|W_{D_1} \cap W_{D_2}|}{|W_{D_1} \cup W_{D_2}|} \in [0, 1]$. The Jaccard index of two documents is 1 if the documents are identical and 0 if they don't have any words in common. However, the set intersection and union needed to compute the Jaccard are expensive operations, which makes it difficult to use this similarity measure when dealing with large data sets.

Instead, it is possible to approximate the Jaccard index with *MinHashing*. Given the set W of all possible words that may appear in the documents, this technique requires a random permutation function $h : W \mapsto \{1, ..., |W|\}$. In other words, $h(w)$ assigns a unique rank to word w; in practice, you can use a collision-resistant hash function for h. The MinHash of a document D is first word from W_D in the ranking given by h: $MinHash_h(D) = \arg\min_{w \in W_d} h(w)$. For two documents D_1 and D_2, the probability of their MinHashes being the same is equal to their Jaccard index $J(D_1, D_2)$.

The MinHash is often not enough to conclude whether two documents are similar. The idea behind *locality sensitive hashing (LSH)* is to compute multiple hashes and map documents into buckets; similar documents are likely to be mapped to the same bucket. With LSH, you need $n = b \times r$ different hash functions and you compute n MinHashes for each document. You then arrange these

Table 1: Locality Sensitive Hashing Example

	h	D_1	D_2	D_3	Buckets
Band 1	$MinHash_{h_1}$	w_1	w_1	w_1	$[D_1, D_2]\,[D_3]$
	$MinHash_{h_2}$	w_1	w_1	w_2	
Band 2	$MinHash_{h_3}$	w_5	w_5	w_3	$[D_1, D_2]\,[D_3]$
	$MinHash_{h_4}$	w_8	w_8	w_4	
Band 3	$MinHash_{h_5}$	w_1	w_7	w_7	$[D_1]\,[D_2, D_3]$
	$MinHash_{h_6}$	w_6	w_6	w_6	

MinHashes into b bands, each band having r rows. Table 1 illustrates LSH for three documents, D_1, D_2, and D_3. We compute six MinHashes, partitioned in three bands with two rows per band. Each band has its own buckets (clusters). Two documents are in the same bucket if their MinHash values match for all the rows in the band (recall that in this case the MinHash of a document is a word). In our example, the MinHashes for documents D_1 and D_2 match in both Bands 1 and 2, so they are in the same bucket in those bands, while in Band 3 documents D_1 and D_2 are in the same bucket. Two documents are considered similar if they appear together in at least one bucket; the documents have b chances of appearing in the same bucket. The probability that two documents D and D' are considered similar, when using LSH, is $1 - (1 - J(D, D')^r)^b$.

Matching an LSH model. Your task is to write a program to find the candidate subset of documents that match the content of a pre-built model using LSH. All operations must be implemented in *Spark*, using parallel collections (*RDDs*). Your program should load documents from disk and distribute them amongst the Spark workers. Each worker is responsible for doing text normalization (removal of multiple white spaces & non-ASCII characters, lowercase conversion, etc.). The workers then compute the LSH for each of the documents and check the content match against a set of documents in the model. The interface returns None for documents with no match; these entries should be filtered out before saving the results to the output folder. The final list should contain the ids of the files that are considered candidate matches and can be further verified using an exact match algorithm. Your program should save this list to disk.

Starter files. The starter files include a partially completed Python program, located at `/homework/minhash_homework.py`. Complete this program by adding the appropriate code to replace all the `None` values in the `__main__` function. The section where you need to add your code is delimited by comments; you must complete 7 steps, in order (each step depends on the completion of the previous steps). The data files are in the `data` subdirectory. Invoke the program like this:

```
$ cd /home/seed/Documents/spark/spark/bin/
$ ./spark-submit homework/minhash_homework.py > data/solution.out
```

You can also find these files in the VM, under `/home/seed/Documents/spark/spark/bin`. The output of your program will be located in `data/output_folder`. The output folder must be deleted between consecutive runs:

```
$ rm -r data/output_folder/
```

Submitting.

Make sure you un-comment the *print* instructions after completing each step. The program stdout will be evaluated as part of your grade. Create an archive, called `spark_problem.zip`, which contains the `data\output_folder`, the `data\solution.out` and your completed `minhash_homework.py` script. The starter files include a shell script, called `pack_files.sh`, that will create this archive for you (run `chmod +x pack_files.sh` if you can't execute it). Submit the `clustering_problem.zip` file as described at the beginning of this handout.