

Introduction to Machine Learning - COS 324

Programming Assignment 1

Due date: 11:59:59pm on 3rd Oct. Electronic submissions only.

IMPORTANT:

1. Consulting with other students from this course is allowed. If you do so, clearly state whom you consulted with for each problem separately.
2. Searching the internet or literature for solutions is **prohibited**.

The purpose of this exercise is to use the Weighted Majority algorithm for text classification. You'll need to install Python 3 (and packages like numpy, scipy, scikit-learn, jupyter) before you begin. Here is how one might install numpy for instance:

```
$ pip install numpy
```

We recommend, but not require, that you install [Anaconda](#) that installs both Python and the required packages. Click [here](#) or [here](#) to learn the difference between Anaconda and a bare minimum Python distribution.

After you are done installing the prerequisites, launch the Jupyter Notebook from the assignment directory. Download the notebook from here: [class-text.ipynb](#)

```
$ jupyter notebook class-text.ipynb
```

The notebook contains instructions on how you must proceed. Here is a useful summary: there are two parts; Part (A) has 20 action prompts, and Part (B) has 14 action prompts. Each action prompt is marked with a sequence number and the token **ACT**. Each prompt is answerable in 5 lines of code or less – most requiring a line of code like:

```
np.argmax(label)
```

(This is, in fact, an answer to one of the prompts.)

Any changes made to the notebook (the substitution of prompts, the output of execution) can be saved by selecting the appropriate option in the notebook toolbar.

Once you have completed all prompts and executed all cells, save the notebook. At this point, it'll contain your answers and the output produced. You're required to submit the file *class-text.pynb* to the course dropbox.

Last but not least, *enjoy* working on your first machine learning algorithm.