CSE 417T: Homework 1

Due: September 12 (Wednesday), 2018

Notes:

- Please check the submission instructions for Gradescope provided on the course website. You must follow those instructions exactly.

- Please complete and submit the following two stub Matlab files for Problem 2.
  `http://classes.cec.wustl.edu/~cse417t/hw1/perceptron_experiment.m`
  `http://classes.cec.wustl.edu/~cse417t/hw1/perceptron_learn.m`

- In addition to submitting the above Matlab files, you need to separately submit a write-up with written solutions for all problems (including reports for programming problems) to Gradescope. Your score will be based on the report. The code you submit is used for checking correctness and for running plagiarism checkers.

- Homework is due **by 11:59 PM on the due date.** Remember that you may not use more than 2 late days on any one homework, and you only have a budget of 5 in total.

- Please keep in mind the collaboration policy as specified in the course syllabus. If you discuss questions with others you **must** write their names on your submission, and if you use any outside resources you **must** reference them. **Do not look at each others' writeups, including code.**

- There are 4 problems on 2 pages in this homework.

- **Keep in mind that problems and exercises are distinct in LFD.**

- All graphs should have clearly labeled axes. The Matlab `hist` function should be useful.

Problems:

1. (30 points) LFD Problem 1.3

2. (30 points) Consider the following experiment on perceptron learning for random training sets of dimension 10:

   - Generate an 11-dimensional weight vector $\mathbf{w}^*$, where the first dimension is 0 and the other 10 dimensions are sampled independently at random from the uniform $(0, 1)$ distribution (the first dimension will serve as the threshold, and we'll just set it to 0 for convenience).

   - Generate a random training set with 100 examples, where each dimension of each training example is sampled independently at random from the uniform $(-1, 1)$ distribution, and the examples are all classified in accordance with $\mathbf{w}^*$.

- Run the perceptron learning algorithm, starting with the zero weight vector, on the training set you just generated, and keep track of the number of iterations it takes to learn a hypothesis that correctly separates the training data.

Write code in Matlab to perform the above experiment and then repeat it 1000 times (note that you're generating a new $\mathbf{w}^*$ and a new training set each time). We have provided two stub files that you should complete for this purpose. The files have comments that explain their inputs and outputs. You need to submit both your codes and the report of this problem.

Once you have your code working, plot a histogram of the number of iterations the algorithm takes to learn a linear separator (you should include this plot in your writeup). How does the number of iterations compare with the bound on the number of errors we derived in the previous problem? Note that this bound will be different for each instantiation of $\mathbf{w}^*$ and the training set, so in order to answer this question, you should analyze the distribution of differences between the bound and the number of iterations. Plot and submit a histogram of the **log** of this difference, and discuss your interpretation of these results.

3. (20 points) LFD Problem 1.7

4. (20 points) LFD Problem 1.8