CSE 5523: Problem Set 3

Due via Carmen dropbox at 11:59 PM 10/11.

1. For vectors $\mathbf{u} = \begin{bmatrix} 2 \\ 5 \\ 3 \end{bmatrix}$ and $\mathbf{v} = \begin{bmatrix} 4 \\ 2 \\ 1 \end{bmatrix}$:

   (a) [3 pts.] Calculate the inner product: $\mathbf{u}^\top \mathbf{v}$.

   (b) [3 pts.] Calculate the outer product: $\mathbf{u}\,\mathbf{v}^\top$.

   (c) [3 pts.] Calculate the Kronecker product: $\mathbf{u} \otimes \mathbf{v}$.

2. [20 pts.] PROGRAMMING:

   (In general for programming problems you should hand in the following as separate files:

   - a copy of each program file you write,
   - a representative sample of each input file you use,
   - a representative sample of each output you produce.

   Your programs should be as short as possible, and may be based on linear algebra and data analysis functions used in the lecture notes, but should not use higher-level packages or functions.)

   You think your two dimensions of data may have a 'sweet spot' of X values where Y values are high. Write a program called 'sweet.py' that takes the following as command-line arguments in the following order:

   (a) a filename of a csv file containing two dimensions (columns 'Y' and 'X') of training data points,

   (b) a filename of a csv file containing one dimension (column 'X') of test data points,

   Your program should output (to standard output) a csv file containing one dimension (column 'Y') of predicted values, using linear regression with a polynomial basis function with terms for x^2, x and 1. You may base your program on the linear regressor in the Lecture Notes 10. Run your regressor on the 'sweet-train.csv' and 'sweet-test.csv' data on the course web page and report the results.

3. [20 pts.] PROGRAMMING:

   You want a linear regressor using gradient descent with a very outlier-sensitive loss function:

   $$\mathsf{L}(y, \mathbf{x}^\top \mathbf{w}) = (y - \mathbf{x}^\top \mathbf{w})^4$$

   Write a program called 'sensitive.py' that takes the following as command-line arguments in the following order:

(a) a filename of a csv file containing two dimensions (columns ‘Y’ and ‘X’) of training data points,

and implements this regressor. Your program should output (to standard output) a csv file containing weights trained using linear regression with the outlier-sensitive loss function, as was done for the sample code in lecture notes 10. You may base your program on the gradient-descent linear regressor in that same lecture notes. You should also add (and manually adjust!) a parameter, as is done in the lecture notes for the log-cosh loss function, to make sure the regressor converges to a reasonable solution at a reasonable rate. Run your regressor on the simple outlier dataset from Section 10.7 of Lecture Notes 10 and report the results.