CSE 5526 - Autumn 2014

Introduction to Neural Networks

Programming Assignment 2

Due Thursday, Oct. 9

Grader: Yuzhou Liu
Email: liu.2376@osu.edu
Office: Dreese Labs 578
Office hours: Monday & Wednesday 3 - 4pm

Implement an RBF network for one input variable and one output variable using Gaussian basis functions. Generate a set of 75 x_p values from a uniform distribution on the interval [0, 1]. Then for each x_p, generate a d_p from the function $d_p = h(x_p) + u_p$ where $h(x) = 0.5 + 0.4\sin(2\pi x)$ and u_p is random variable with a uniform distribution over the interval [-0.1, 0.1]. Implement the k-means algorithm and use it to determine the centers of the Gaussian nodes. Set the variance of each cluster to the variance of the points assigned to it. If a cluster contains only one point, use the mean variance of all the other clusters as its variance. Use the LMS rule to learn the hidden-to-output weights (note a bias term/node needs to be added). Train it for 100 epochs. Train 10 models using all combinations of the following parameter settings:

1. Number of bases: 2, 4, 7, 11, 16.
2. Values of learning rate, η: 0.01 and 0.02.

Plot for each model a graph that shows the data points, (x_p, d_p), as separate points; the original function, $(x_p, h(x_p))$, as a line; and the predictions of the RBF network, (x_p, y_p), as a line. Comment on the network performance as a function of the number of bases and the choice of η.

Repeat the above experiments using the same variance for all clusters as described in lecture. Comment on the comparative advantages of the two methods of determining cluster variance.

Turn in the following to the project's dropbox on Carmen:

1. A 1-2 page summary report, plus the requested plots
2. Your source code
3. A script or executable that can be run with no arguments to generate the plots in your report and any other evidence that your implementation is working properly
4. A README file identifying and explaining the script/executable from 3.