CS 688
Programming Assignment 1
due Nov. 2$^{\text{nd}}$ 2009

Brian Hrolenok

November 2, 2009

Problem Description

Implement the learning algorithm for the Perceptron discussed in class. (You can use whatever language you prefer, but I strongly encourage you to try MATLAB.) Use your algorithm to do crossvalidation on the following set of data (click here). The data corresponds to points in two spirals, each set of points belonging to a different class (2 class problem). Discuss your results.

Extend your implementation to a two-level neural network, using backpropagation, and try the same data again. Compare the results with the Perceptron implementation.

Implementation

I decided on a pure C implementation for this assignment. No external source code or libraries were used. I referenced the lecture slides from chapter 4 of Machine Learning by Tom Mitchell, and Section 20.5 from Russell and Norvig's Artificial Intelligence: A Modern Approach. The bias is not explicitly represented, so I supplied an input set to 1 for all training and testing.

Results

For testing, I used the leave-one-out approach, and averaged across all of the data sets. On both the Perceptron and the feed-forward two-layer neural network I used a learning rate of 0.1, 400 training repititions, and randomized the initial weights. On the feed-forward NN, I used 10 hidden nodes.

The Perceptron will correctly classify about 50% of the testing data, which is reasonable, as the dataset is highly non-linear. The feed-forward two-layer NN on the other hand, can correctly classify about 90% of the testing data.

Code

The source code is broken into three files, *ann.h*,*ann.c*, and *main.c* which can be found at `http://mason.gmu.edu/~bhroleno/`. Compilation requires GCC and the standard C libraries (stdlib.h,stdio.h,math.h, time.h). Further instructions can be found in the README provided in the tarball.