

Computer Science 152, Fall 2010

Assignment 3

Due Tue. 21 September 2010

Create a 2-layer MLP program using the backpropagation training method, with an eye toward extending this to an arbitrary number of layers. That is, do not wire 2-layerness into your program.

Use the data from assignments 1 and 2. In particular, your network should be able to learn xor exactly (although it may take awhile).

- Use on-line training.
- Use appropriate sigmoids for all functions in training. For testing on discrete output examples, the final layer sigmoids should be replaced with hard limiters with threshold at the corresponding cross-over point.
- Initialize your weights to small random reals (from -0.05 to 0.05, say). If you do not randomize, they will all be updated the same, and you will not achieve the needed weight diversity to reach the training goal. If you make the weights too large, then the neurons may saturate, making it difficult to reach the goal.

While you don't need to have in the implementation this week, keep things open for using momentum and variable learning rate.

Draw comparisons between the performance of the MLP, and Perceptron and Adaline from previous assignments.