CSE 599 Winter 2001

Homework Assignment # 4 (last one!)

Due Date: February 22 (two weeks from today Feb 8)

The purpose of this homework assignment is to familiarize yourself with three popular artificial neural networks, and to get a sense of the capabilities and weaknesses of each in the context of a pattern recognition task. The task is the following: Given a noisy binary image of a number between 0 and 9 (inclusive), classify the input by activating one of ten output units representing the numbers 0 through 9.
You will examine three types of networks: a linear threshold network (Adaline) with an input and output layer, a three layer backpropagation network, and a Hopfield network.

C code for these three networks can be found at:
http://www.cs.washington.edu/education/courses/599/01wi/admin/Assignments/nn.html
You will need to make minor modifications of the code for the backprop and Hopfield networks to fit our specific task.

a. Train the Adaline and backpropagation networks on a data set containing 10 noiseless images of the numbers 0 through 9, as given in the C code for the Adaline network (with image size fixed at 5 x 7 pixels). Now, generate a test data set containing 20 noisy images, 2 images for each number. Produce the first noisy image by inverting (changing from 0 to 1 or vice versa) the 3rd and 4th pixels in both the second and third row of the training image for each number, and the second noisy image by inverting the 3rd and 4th pixels in both the fifth and sixth row of the training image. For each of the networks, report the percentage of correct classifications for the test data set and give a print-out of network outputs for the 20 test images (in the form: input image → output).
b. Train the Hopfield net on the 10 noiseless images as in (a) above. Test the trained network on the 20 noisy images. What percentage of noisy images were mapped by the net to the correct stored noiseless image? Give a print-out of the results (in the form: input image → final output image).
c. Use the trained Hopfield net as a pre-processor. For each of the 20 test images, use the output of the Hopfield net as the input to the two trained networks in (a) and report the new percentage of correct classifications for the two networks on the noisy data set. Identify the test images that are now correctly/incorrectly classified as compared to your results in (a).

Extra-Credit: To get realistic estimates of classification rates, we typically need to train networks with a larger set of examples and vary the number of hidden units. Generate a data set of 100 noisy images of the numbers 0 through 9, and set 50 of these apart as the test set. Use the other 50 to train backpropagation networks with 10, 20, and 30 hidden units. What percentage of the 50 test images are correctly classified by each of these networks? Does having more hidden units always yield better classification?