

Sep 26, 06 11:26	hw2.txt	Page 1/2
<p>CS578 Fall 2006 Empirical Methods in Machine Learning and Data Mining Homework Assignment #2 Due: Tuesday, October 17, 2006 at 2:55pm *before* class</p> <p>The goal of this assignment is to experiment with artificial neural nets trained with backpropagation, early stopping, and N-fold cross validation. For this assignment use the hwl.dta dataset used in the first assignment. The goal is to predict the boolean variable in col 1 from the other attributes. You will have to recode the data so it can be given as input to a neural net. One way to do this is to convert N-arity nominal attributes to N boolean inputs on the net. Another way is to use only one input for each N-arity nominal attribute by treating the values as ordered numbers. As in HW1, you will experiment with both codings.</p> <p>You may implement backprop yourself, or use a commercial/public domain implementation. Note that it probably will take more time to install, learn to use, and modify someone else's implementation than to program backprop yourself, so we encourage you to code backprop yourself. In fact, implementing bp yourself counts as extra credit. If you decide to use someone else's package, it is up to you to install it and make sure that it will support the experiments needed for this assignment. One public domain package you might want to consider that runs on a variety of platforms is SNNS: Stuttgart Neural Network Simulator. There also is a Matlab toolbox for neural nets that is supposed to be pretty good.</p> <p>In this assignment we will use Accuracy, RMSE, and ROC Area as the performance measures. C-code for calculating ROC Area is available on the course web page. The program is called perf.c, and calculates a number of performance measures such as Accuracy, RMSE, and ROC Area. (ROC will be discussed in class next week.)</p> <p>EXPERIMENTS:</p> <p>0: Code each attribute by splitting nominal attributes into boolean attributes. This is already done for you in the data for HW1. Make sure that the output target is scaled 0-1 so that you can use a sigmoid output unit.</p> <p>Draw one random sample of 10,000 points from the data set to use for your experiments. Save the remaining 15,000 points as a final, final test set.</p> <p>1: For neural nets, you need train sets (backprop sets), early stopping sets (technically part of the train set), and test sets. Use N-fold CV from the 10k cases for the train/test sets. The early stopping set should be held out of the train set. One way to do this is to split the data into 5 2k folds. Do backprop on folds 1-3 (60% of the train data), use fold 4 for early stopping (20% of the train data), and test on fold 5. Repeat this process 5 times for 5-fold CV. There are other ways to do this and you don't have to use 5-fold CV. Carefully explain how you choose to do N-fold CV. A diagram or table showing the splits might be helpful.</p> <p>2: Train fully-connected feedforward neural nets using vanilla backpropagation with momentum. Every backprop implementation defines learning rate and momentum somewhat differently, and the definitions often vary when using batch mode (updating once per epoch -- full pass through the training set) or when updating per pattern, so you'll have to experiment with the parameter settings to find values that work well with your code. You can use batch mode, per pattern, or per group of pattern updating. (HINT: If the nets are fully trained after less than 200 passes through the train set you're probably training too fast. If the nets are taking more than 10^6 passes through the</p>		

Sep 26, 06 11:26	hw2.txt	Page
<p>train set you're probably training slower than necessary.) Feel free to use any number of hidden units that seems to work. I often use 16 or 32 hidden units when starting on a new problem.</p> <p>3: Compute accuracy, RMSE, and ROC Area on the train, stopping, and test sets. Show graphs of performance vs number of epochs for the train and early stopping sets. The performances on the test sets should be reported at the early stopping point. Are the early stopping points for accuracy, RMSE, and ROC Area the same?</p> <p>4: Try an alternate coding of the input attributes such as converting nominal attributes to N-valued inputs (e.g. code as integers). Repeat the experiment above. How does performance with this alternate coding compare with the boolean coding?</p> <p>5: Using the boolean coding, experiment with different numbers of hidden units. You might try 1,2,4,8,16,32,64,... or even 1,4,16,64,... What size net yields best generalization performance?</p> <p>EXTRA CREDIT -- do one or more of the following:</p> <ul style="list-style-type: none"> - experiment with different size training sets (generate a learning curve). how does the performance change with the size of the training set? does the optimal net size change? - there are different ways to convert the nominal attributes to boolean attributes. try something clever. one idea might be to use the decision trees trained in HW1 to help you decide how to code the nominal attributes for the neural nets. - recode the inputs so that they are symmetric about 0 (i.e., go negative and positive). does this change performance? - N-fold CV leaves you with N or more trained neural nets. compare the average prediction of these nets with the performance of each of the nets alone to see which works better. use the final, final test set of 15k cases as the test set for this comparison. - do a study of the effect of altering the learning rate and momentum on the generalization performance of the nets - try nets with two or more hidden layers. do they perform better than nets with one hidden layer? are they harder to train? - is there a difference in generalization performance and training time between batch mode updates and per pattern or per group of pattern updating? - compare weight decay with early stopping. does one perform better than the other? is one easier to use than the other? - use cross-entropy to train the net. does training on cross-entropy yield better accuracy than training on squared error? - do feature selection to find a subset of the features that seems to perform better than using all the features - do a sensitivity analysis to figure out what inputs the trained nets use most. sensitivity analysis can be done by looking at derivatives of the output of the net with respect to the inputs, or by experimenting with injecting noise into the inputs one at a time - implement vanilla backprop with momentum for fully-connected feedforward neural nets containing one hidden layer and trained with squared error <p>Hand in a brief (5 pages or less for the main assignment) summary of the results with enough documentation so that we can see what you did and how you did it. Do not write a paper. This is homework, not a class project. You'll probably want to use the neural net code for the class project, so effort spent now to write good code or become familiar with the package you use should pay off later. Attach a copy of any code you write to the report you hand in.</p> <p>Have fun.</p>		