

COMP4702/COMP7703 - Machine Learning

Prac 9 – Assessing algorithms, Bagging and Boosting

Aims:

- To perform some experiments to highlight issues involved in assessing the performance of predictors.
- To explore the performance of bagging and boosting ensemble techniques.
- To produce some assessable work for this subject.

Procedure:

Assessing Algorithms

In this prac we will use the Weka software. Weka has an implementation of multilayer perceptrons trained with backpropagation (in the “functions” folder of classifiers). Open this up and read the information associated with the MLP model (including how to create/modify an MLP using the GUI). In the first part of this prac we will experiment with training MLPs on two of the datasets used previously – the diabetes and sonar datasets. Choose a model and associated parameters (e.g. number of hidden units, learning rate, momentum, number of epochs) and use this setup **unchanged** through questions x-x below.

- Note: you will need to use the csv file converter in Weka to load these datasets. This is easy, except that Weka insists that the first line of the file should be a header (containing the names of each attribute in the datafile). As a quick fix to get around this issue, I have duplicated the first datapoint in each dataset! This results in nonsensical attribute names, but we don't have this information and it's not important for our purposes.
- **Q1:** Train your network 10 times from random weight initializations (using the default 66% split of training/test sets) and record the min, max, mean and standard deviation of your 10 results (in terms of percentage correct/incorrect) on each dataset (diabetes and sonar). Note: to do different random weight initializations, you will need to vary the random seed parameter of the MLP (right-mouse click on the classifier “bar” once you're selected MLP). For these experiments, use random seeds 0,1,...,9.
- Pick the experiment from Q1 that was the closest to your mean result (i.e. some random seed value). Using this setup:
 - **Q2:** Perform 10-fold cross-validation and record percentage correct/incorrect result.
 - **Q3:** Do 10 random partitions of the dataset (66% training/test split) and retrain 10 times and record the min, max, mean and standard deviation of your 10 results (in terms of percentage correct/incorrect). Note: to do repartition the dataset in Weka, you need to go to the Preprocess tab and use the “Random” filter (under filters-unsupervised-instance). Click apply to apply the filter, then return to the Classifier tab to use the new split. There is another random seed value for this filter – use values 42,43,...51 here.
 - **Q4:** Perform a bootstrapping process by training your classifier using resampled versions of the dataset, again recording min, max, mean and standard deviation. Note: use the “Resample” filter repeatedly in Weka to do this, random seeds 1,2,...,10.

You should also have a look at some of the output that WEKA produces and connect it back to lecture notes (confusion matrix, true/false positives/negatives, precision, recall).

Bagging and Boosting

To do this part of the prac, you will need to refer to the highly readable paper (available on the course material webpage):

D. Optiz and R. Maclin. Popular ensemble methods: an empirical study. *Journal of Artificial Intelligence Research* v.11 (1999) pp.169-198.

In this paper, the authors produce a large set of experimental results, including MLPs and bagged and boosted ensembles of MLPs. The datasets used include most of the datasets we have used during earlier pracs.

- **Q5:** For the sonar, diabetes and ionosphere datasets, attempt to reproduce the results reported in the paper – specifically in the first column of Table 2. You should be able to find most of the experimental details used by the authors in the paper – provide comments if you have to make any assumptions for your experiments. Record your results and comment briefly on any difference between them and the results in the paper.

Weka contains implementations of both bagging and AdaboostM1. To use them, you select one as the classifier from the “meta” folder. Once you’ve done this, right-clicking on the classifier “bar” opens a window detailing the parameters of the method, including what base classifier to use. Any of the models in Weka can be selected in the same way as for a single classifier. In our case we will build ensembles of MLPs. Read the information in Weka to make sure you understand what you’re doing!

- **Q6:** For the datasets used in Q5, attempt to reproduce the results in columns 3 and 5 of the paper. Provide comments if you have to make any assumptions for your experiments. Record your results and comment briefly on any difference between them and the results in the paper.
- **Q7:** Try bagging and boosting with a single-layer perceptron (set the “Hiddenlayers” parameter to 0 for an MLP in Weka). Record your results for the datasets used in Q5 and Q6 and comment briefly on how they compare to your results from Q5 and Q6.