CSC412 – Assignment #4

Due: March27, 10am at the **START** of class
Worth: 15%
Late assignments not accepted.

1 Fully Observed Trees

In this assignment you will train a fully observed tree model on word-document vectors taken from USENET articles.

1.1 What to do

- Using the data provided in `a4newsgroups.mat`, train a fully observed tree model with the optimal structure. The data has 100 (binary) features and 16242 training cases.
- The matrix `documents` contains one column per USENET article.
  Each of the 100 rows is a binary variable indicating whether a certain keyword appears in the posting or not.
  The cell array `wordlist` has the 100 words corresponding to these rows.

1.2 What to hand in

- Any *compact, clear and unambiguous* representation of the optimal (undirected) tree structure found by your algorithm. For example, you might print out all pairs of words `(word_i--word_j)` corresponding to edges chosen by the tree learning algorithm. Or you might want to draw a picture of the graph (see below).
  DO NOT hand in lists of numbers corresponding to pairs of features chosen – convert everything back to words using the `wordlist` cell array.
- A histogram (with at least 100 bins) of the log likelihoods of the training cases under your optimal model.
  Compute and print out the min,max,mean and median log likelihood.
  Which training case (give its number) has the worst (lowest) log likelihood?
  Print out the list of keywords appearing in this posting.
  What percentage of postings mention both "god" and "ftp"?
- After finding the optimal tree structure for the 100 variables, you may choose any node you wish as the root and form a directed graph. Different choices of root node will result in different output when you visualize a directed tree, but of couse the undirected model structure always remains the same. Chose a root word. Hand in a *compact, clear and unambiguous* representation of a directed tree structure created by directing all arcs in the optimal structure away from a your chosen root. For example, you might print out a breadth-first list of directed links `(word_i->word_j)` in the tree. Alternately, you might want to draw a picture (see below).
- Hand in your MATLAB code, properly commented (4 pages maximum).
  Also email this code to us at `csc412@cs` since we may want to run it.

1.3 Some hints

- Be careful when computing the mutual information weights for pairs of variables that have zero counts for some of their joint settings. In this case the mutual information should be zero, but a careless calculation will lead to division by zero or log of zero errors in the code.

- To save you some trouble, I've written the function `mwst.m` for you, which finds minimum or maximum weight spanning trees given a weight matrix.
- Here is a trick for computing the 2 by 2 joint count table of two binary column vectors `v1` and `v2`:
  ```
  counts12 = reshape(hist(v1+2*v2,0:3),2,2);
  ```
  but make sure you understand it before using it!
- The data is almost a megabyte in MATLAB , stored as a sparse matrix, so be careful about making copies of it in memory. You should be able to do everything you need without making another (sparse or nonsparse) copy of the whole dataset.
- In case you are interested, the array `newsgroup` tells you the toplevel newsgroup category of each posting, `1=comp.*, 2=rec.*, 3=sci.*, 4=talk.*`, although we won't use it in this assignment.

1.4 Reminders

Amongst other things, your code will need to:

- Compute the marginal count for every feature being on or off summed across all documents.
- Compute the joint counts for every pair of features being on together, off together, on-off or off-on, summed across all documents.
- Convert these counts into mutual information weights.
  Be careful to do this properly for joint counts which are zero!
- Find the best spanning tree over the features given these weights. Depending on how you compute the weights, you will either need a minimum or maximum weight spanning tree.
  (This will in turn affect how you call `mwst.m` if you choose to use it.)
- Select a root for the tree and direct all edges away from it to get a directed graph.

1.5 Automatic Graph Layout

- If you want to be really fancy, you can try using the function `dottree.m` and the automatic graph layout programs `dot` and `neato` to draw pictures of your final structures. It is very cool looking, but you will not get any more or any fewer marks for using this method or a simpler method to display your results.
  Caution: do not waste too much time trying to get this to work. Get the rest of the assignment finished first.
- Call `dottree(tree,'fname',wordlist)` and it will write out two files, `fname.u` and `fname.d`.
  Now call the `neato` and `dot` programs to convert these files into pictures.
- Information and downloads for `dot` and `neato` are available at `http://www.graphviz.org/`.