Machine Learning 15-681/781
Tom Mitchell, Fall 1998

Homework Assignment 4
Due at beginning of class on Nov 17, 1998

1. Exercise 5.2 from the textbook.
2. Exercise 5.4
3. Exercise 7.8
4. Exercise 6.5
5. This question asks you to consider the complexity of several probabilistic approaches to learning the Cup target concept summarized in Table 12.3 (page 342) of the textbook. Note all the attributes in this problem are boolean valued.
   - First consider using a Naive Bayes classifier to learn this target concept. Exactly how many distinct probability terms must be estimated from the training data to learn a Naive Bayes classifier for this problem?
   - Now consider learning a Bayesian belief network to recognize Cups, instead of a Naive Bayes Classifier. In particular, use the network suggested by the domain theory in Table 12.3, and interpret each rule in the domain theory as a specification of the parents of the corresponding network variable. (e.g., the first rule in the domain theory specifies that $Parents(Cup) = \{Stable, Liftable, OpenVessel\}$).

     Draw the Bayes network as a graph with directed edges. You may omit variables that are not mentioned by the domain theory.

     Exactly how many distinct probability terms must be estimated from the training data to learn the conditional probability tables for your network?
   - Both of the above approaches make conditional independence assumptions to reduce the complexity of estimating $P(Cup|BottomIsFlat, ConcavityPointsUp, Expensive, \ldots)$ from the training data. If no such assumptions are made, how many distinct probability terms must be estimated from the training data?