**Homework #4 (30p)**
**Due:** Tue. Nov. 25 in class

**Note:** Please hand print your answers. If we cannot read them we will assume that they are wrong.

1. Consider the following problem: The US security rules forbid passengers to carry weapons on airplanes. A Swiss-army knife is found in a backpack that belongs to John Doe, an airline passenger, during a security check at the *Dulles International Airport*.
   a. Represent these facts as first-order logic clauses.
   b. Use backward-chaining to prove that John Doe broke US security rules.

2. Sam, Clyde, and Oscar are elephants. We know the following facts about them:
   1. Sam is pink.
   2. Clyde is gray and likes Oscar.
   3. Oscar is either pink or gray (but not both) and likes Sam.

   Use resolution refutation to prove that a gray elephant likes a pink elephant; that is, prove $(\exists x, y)[Gray(x) \wedge Pink(y) \wedge Likes(x, y)]$.

3. Consider the following seven clauses:
   a. $\neg A \vee \neg B \vee \neg C$
   b. $\neg A \vee B$
   c. $\neg A \vee C$
   d. $\neg B \vee C$
   e. $\neg B \vee A$
   f. $\neg C \vee A$
   g. $\neg C \vee B$

   where "$\neg$" stands for logical negation and "$\vee$" stands for logical OR.

   a. (10p) In the following truth table $F$ stands for logical false and $T$ stands for logical true. Fill the table and see if you can find an assignment of variables that satisfies all 7 classes.

   b. (10p) Enumeration using truth tables is not practical for large number of clauses. The GSAT algorithm for finding a model to satisfy a set of clauses uses hill-climbing search. It starts by choosing a random assignment of truth values to all variables. At each iteration one variable assignment which increases number of satisfied clauses, i.e. the number of clauses which are true, is chosen. The algorithm stops when no assignment increases the number of satisfied clauses. It sometimes ends in a local maximum. Use GSAT to find a global maximum, i.e. an assignment that satisfies all clauses. Use $A = T, B = F, C = F$ as initial value.
   c. (5p) Show that there is an assignment of truth values for A, B, and C that is a local but not a global maximum of the number of clauses satisfied by that assignment.

| A | B | C | $\neg A \vee \neg B \vee \neg C$ | $\neg A \vee B$ | $\neg A \vee C$ | $\neg B \vee C$ | $\neg B \vee A$ | $\neg C \vee A$ | $\neg C \vee B$ |
|---|---|---|---|---|---|---|---|---|---|
| F | F | F | | | | | | | |
| F | F | T | | | | | | | |
| F | T | F | | | | | | | |
| F | T | T | | | | | | | |
| T | F | F | | | | | | | |
| T | F | T | | | | | | | |
| T | T | F | | | | | | | |
| T | T | T | | | | | | | |

4. A patient takes a lab test for a certain type of cancer and the result comes back positive. However, the test returns a correct positive result in only 98% of the cases in which the disease is actually present, and a correct negative result in only 97% of the cases in which the disease is not present. Furthermore, 0.008 of the entire population have this cancer.
What is the actual probability that the patient has cancer?

5. Suppose you are going home, and you want to know what is the probability that the lights are on given the dog is barking and the dog does not have any bowel problem. If the family is out, often the lights are on. The dog is usually out in the yard when the family is out and when it has bowel troubles. And if the dog is in the yard, it probably barks. The Bayesian network for the example is given in the figure:

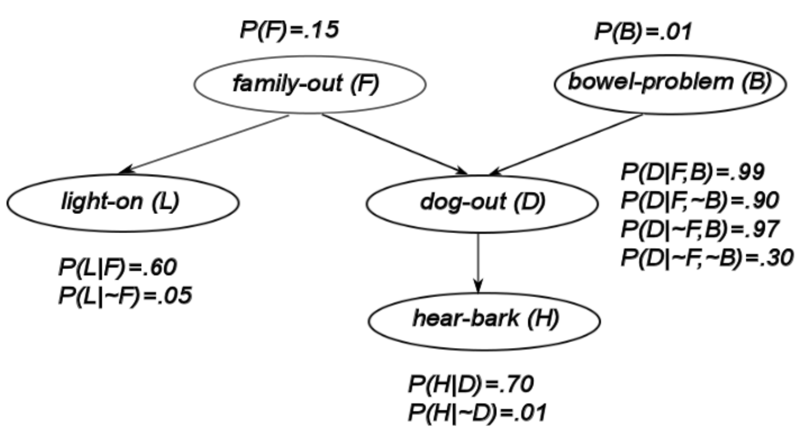


What can you infer when all the evidence points the same way?

Should you assume the family is out if the light is on, but you do not hear the dog?

What if you hear the dog, but the light is not on?

Note: Is there a way to answer both questions with a single query?

6. Orville, the robot juggler, drops balls often when its battery is low. In previous tests, it has been determined that the probability that it will drop a ball when its battery is low is 0.9. Whereas when its battery is not low, the probability that it drops a ball is only 0.01. The battery was recharged not so long ago, and our best guess (before looking at Orville's latest juggling record) is that the odds that the battery is low are 10 to 1 against. A robot observer,

with a somewhat unreliable vision system reports that Orville dropped a ball. The reliability of the observer is given by the following probabilities:

$$p(\textit{observer says that Orville drops} \mid \textit{Orville does drop}) = 0.9$$

$$p(\textit{observer says that Orville drops} \mid \textit{Orville does not drop}) = 0.2$$

Draw the Bayes network, and calculate the probability that the battery is low given the observer's report.