CPS 170: Introduction to AI	**Instructor:** Ron Parr
Homework 4	**Due:** Thursday April 2, 2010

1 Bayesian Networks (10 points)

Let H_x be a random variable denoting the handedness of an individual x, with possible values l or r. A common hypothesis is that left- or right-handedness is inherited by a simple mechanism; that is, perhaps there is a gene G_x, also with values l or r, and perhaps actual handedness turns out mostly the same (with some probability s) as the gene an individual possesses. Furthermore, perhaps the gene itself is equally likely to be inherited from either of an individual's parents, with a small nonzero probability m of a random mutation flipping the handedness.

(a) Which of the three networks in Figure 1 claim that $P(G_{father}, G_{mother}, G_{child}) = P(G_{father})P(G_{mother})P(G_{child})$?

(b) Which of the tree networks make independence claims that are consistent with the hypothesis about the inheritance of handedness?

(c) Which of the three networks is the best description of the hypothesis?

(d) Write down the CPT for the G_{child} node in network (a), in terms of s and m.

(e) Suppose that $P(G_{father} = l) = P(G_{mother} = l) = q$. In network (a), derive an expression for $P(G_{child} = l)$ in terms of m and q only, by conditioning on its parent nodes.

(f) Under conditions of genetic equilibrium, we expect the distribution of genes to be the same across generations. Use this to calculate the value of q, and, given what you know about handedness in humans, explain why the hypothesis described at the beginning of this question must be wrong.

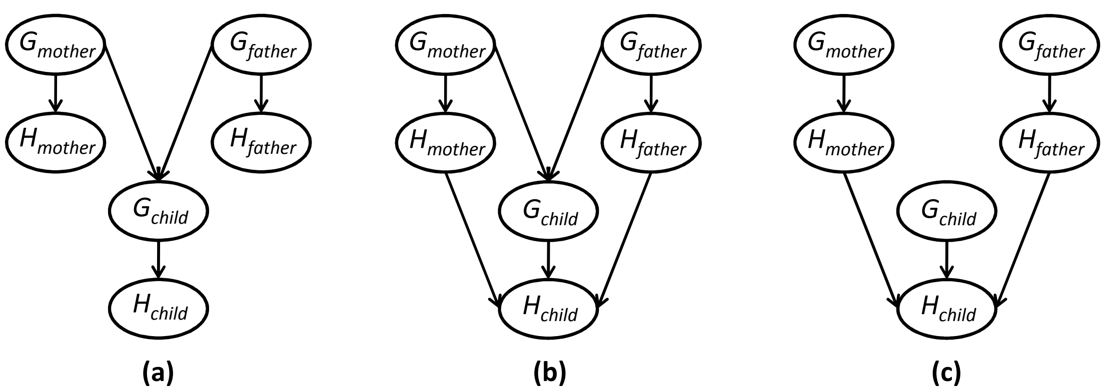

Figure 1: Three possible structures for a Bayesian network describing genetic inheritance of handedness.

2 Bayesian Networks (10 points)

Do problem 14.8, parts (a) through (d) (14.1 in the second edition).

3 Bayesian Networks (10 points)

Do problem 14.15 (14.7 in the second edition).

4 HMMs (10 points)

In this exercise, we examine what happens to the probabilities in the umbrella world in the limit of long time sequences.

(a) Suppose we observe an unending sequence of days on which the umbrella appears. Show that, as the days go by, the probability of rain on the current day increases monotonically toward a fixed point. Calculate this fixed point.

(b) Now consider *forecasting* further and further into the future, given just the first two umbrella observations. First, compute the probability $P(r_{2+k}|u_1, u_2)$ for $k = 1...20$ and plot the results. You should see that the probability converges towards a fixed point. Prove that the exact value of this fixed point is 0.5.

5 HMMs (10 points)

Do problem 15.4 (same number in the second edition).

6 HMMs and robotics (20 points)

For this problem you will implement a simple solution to the "kidnapped robot" problem. Your code should read in a map that uses "#" signs for walls and periods for empty spaces. For example, the following text file will correspond to the map in the class notes:

```
######
#....#
######
```

After the last line describing the map, your file should contain list of robot actions in $\{L, R, U, D\}^*$. For example, if robot action sequence is two right movements (as in the class notes), your file would be:

```
######
#....#
######
RR
```

Finally, assume that the last line in the file contains a number which is the robot's true state. This number will be hidden to the robot, but your code should use this to track the robot's state as the robot moves. Assume that the states are numbered in raster order (from left to right, top to bottom as they appear in the input file.

You program should assume an initial distribution which is uniform over all empty grid squares. You should also assume that if the robot bumps into a wall, it stays in the same place. Finally, you should assume (as in class) that the the robot has four proximity sensors that tell it if the robot is adjacent to a wall in each of the four ordinal directions.

The output from your code should be a distribution over states, expressed as a list of probabilities, one per line.

(a) Assuming deterministic motion and sensors, reproduce the result in the class notes in which the robot determines its position after two movements. (5 points)

(b) Assume that sensors produce the correct output with probability 0.9 and the opposite with probability 0.1. Further assume, that motions succeed with probability 0.9 and fail (leave the robot in the same position) with probability 0.1. If the robot movements are the same as in part (a), what is the distribution? (Note that this won't be the same every time you run the code because of the randomness in the actions and observations, so show us a few runs.) (5 points)

(c) Consider the follwing map:

```
#######
#.#.#.#
#.#.#.#
#.#.#.#
#.....#
#######
```

What is a sequence of actions that will concentrate most of the probability distribution in one state, thereby helping the robot figure out its location. Show your code running on this sequence. (For debugging purposes, we suggest you try this first without noise and then show it working with noise.) You may notice that a sequence which works perfectly without noise will not focus the robot's distribution with probability 1.0 when there is noise. Can you propose a way to concentrate the probability distribution even more? (10 points)