**CS 188 Final Exam, 8-11 am, Thursday May 18, 2000**

Make sure you have all 3 pages of this exam. The maximum score possible is 120 points. Good luck.

1. Design a Bayes net to model how a student is going to perform on the CS188 final exam. There are three binary random variables to consider

   - *Lazy* which is true if a student is lazy, false if s/he is not.
   - *Difficult* which is true if the final is difficult and false otherwise.
   - *Pass* which can be true or false.

   Suppose that 30% of the students are lazy, and 60% of the time, the final is difficult. A lazy student will pass an easy exam 50% of the time, and a difficult exam 20% of the time. A student who is not lazy will pass an easy exam 100% of the time, and a difficult exam 90% of the time.

   (a) Draw a suitable graph for representing the Bayes net and indicate clearly the information (eg. conditional probability tables or prior probabilities) required to completely specify the Bayes net.

   (b) Suppose you know that a student has failed the exam. What is the probability that s/he is lazy?

   (c) Now suppose you also know that the exam was difficult. In the light of this additional evidence, what is the probability that s/he is lazy?

2. In this problem, you are asked to model a simplified version of the weather and umbrella HMM, that we all know and love. The weather on any day can be either sunny or rainy, each with prior probability 0.5. From one day to the next, the weather remains the same with probability 0.7 and changes with probability 0.3. On a sunny day, the probability that you will observe the caretaker carrying an umbrella is 0.2, on a rainy day it is 0.6. The weather outside is hidden to you, all you can observe is the presence (or absence) of an umbrella.

   On the first day your observe the caretaker come in with an umbrella, and on the second day without an umbrella. Given both these facts,

   (a) What is the chance that the weather outside is raining on the first day?

   (b) What is the chance that the weather outside is raining on the second day?

3. The MAJORITY function of 3 inputs $x_1$, $x_2$, $x_3$ is defined to be 1 if at least two of the inputs are 1, and 0 otherwise. Can this function be represented by a single layer perceptron? Either prove that this is impossible or construct such a perceptron.

4. Speech recognition is commonly done using Hidden Markov Models. Outline the basic idea–you should be very clear about what the hidden and observed variables correspond to in such systems. Two problems that such systems have to deal with are due to different accents or due to co-articulation. How are these handled?

5. *Natural Language Processing*

    - State four different speech acts, each with an example.
    - Give an example of a sentence that can be parsed in at least three different ways.
    - Give an example of anaphoric reference.

6. In this exercise, you have to work out a solution to a simplified version of the gambler's problem you programmed for one of the homeworks.

    A gambler has the opportunity to make bets on the outcome of a sequence of coin flips. If the coin comes up heads, she wins as many dollars as she has staked on that flip; if it is tails she loses her stake. The game ends when the gambler wins by reaching her goal of attaining \$ 4, or loses by running out of money. On each flip, the gambler must decide what portion of her capital to stake, in integer number of dollars. The state is the gambler's capital $s \in 1, 2, 3$. The reward is zero on all transitions except ones in which the gambler reaches her goal, when it is $+1$.

    Let the probability of the coin coming up heads be $p$. I want you to solve for the utilities in two cases:

    (a) For $p = 0.6$, the optimal policy is to always bet \$1. Use value determination to calculate the utilities for each of the three non-terminal states.

    (b) For $p = 0.4$, the optimal policy is non-obvious. Implement value iteration to calculate the utilities. Start with an initial guess of $U_1 = 0.25$, $U_2 = 0.5$, $U_3 = 0.75$

7. Answer true or false, and give a short (one sentence) explanation for your answer. (No credit without the explanation.)

(a) Convolving an image with a Gaussian is a good way to detect edges in it.

(b) The Focus of Expansion of an optical flow field gives information about speed.

(c) Temporal difference learning needs a model of the environment.

(d) Given the optimal $Q(s, a)$ function for a Markov decision problem, the optimal policy is automatically determined.

(e) Corresponding to a single partial order plan, there can be more than one sequence of operators that will transform the initial state to the goal state.