


CS 188 Final Exam, 5-8 pm, Wed Dec 16, 1998

1. 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 0.8 and changes with probability 0.2. On a sunny day, the probability that you will observe the caretaker carrying an umbrella is 0.3, on a rainy day it is 0.7. The weather outside is hidden to you, all you can observe is the presence (or absence) of an umbrella.
 - (a) If on the first day of your enforced seclusion, you observed the caretaker come in without an umbrella, what is the chance that the weather outside is raining?
 - (b) Now suppose that you observe that on the second day, the caretaker came in with an umbrella. What would you now say about the probability that the weather outside was rainy on the *first* day?
 - (c) Calculate $\alpha(q_2) = P(y_1, y_2, q_2)$.
2. Consider the weather-umbrella situation defined above. Suppose that you are given the 3-day sequence of observations *no-umbrella, umbrella, umbrella*. Calculate the most probable sequence of weather states for the 3 days using the Viterbi method (dynamic programming). You can get half the credit for the problem if you use a method based on explicitly finding all the entries of the joint probability distribution.
3. Consider the problem of planning/control for a mail delivery robot. This robot works inside Soda Hall—its job is to pick up mail from the front desk on the third floor and deliver it to any of the rooms on the same floor. Design suitable operators which could be used by a partial order planner to come up with a plan to deliver mail. How would you use feedback control in this task?
4. You have been hired by the Sierra Club to design a classifier for labeling each roadside plant that you see as either a tree or a bush (ignore other possible plant types) using height as a feature. You can assume that the probability distribution of height for each class is a Gaussian, with mean height being 6 feet for bushes and 12 feet for trees. The standard deviation σ is 3 feet in each case. 70% of the roadside plants are trees. If you have to use a classifier of the type—Declare a plant to be a bush if the height is less than h_0 , a tree otherwise—what should h_0 be?
5. Answer true or false, and give a short (one sentence) explanation for your answer. (No credit without the explanation.)

- (a) To decide when to stop training a multilayer perceptron using backpropagation, it is more important to look at the misclassification error on the training set instead of the test set.
 - (b) An inference rule can be sound but not complete.
 - (c) The Focus of Expansion of the optical flow field is not changed when the robot starts moving in a different direction, if it keeps its speed constant.
6. You find yourself in a tunnel shown below. There are two terminal states, a desirable one with value $+1$ and an undesirable one with value -1 . For each of the nonterminal states the reward function R is -0.1 . At any nonterminal state, you can choose to move left or right. However you may not necessarily achieve what you want. 60% of the time, you will actually move in the desired direction; 40% of the time you will be moved to the neighboring room in the other direction.



The optimal policy for this problem is obvious—move in the direction that takes you closer to the $+1$ state. Calculate the values associated with each of the nonterminal states.