

15-381 Artificial Intelligence: Representation and Problem Solving

Homework 3

Out: 10/02/08

Due: 10/16/08

Instructions

This assignment is due on Thursday, October 16, 2008. The written portion must be turned in at the beginning of class at noon on October 16. Type or write legibly; illegible submissions will not receive credit. Write your name and Andrew ID clearly at the top of the assignment. The programming portion of this assignment must be electronically submitted by noon on October 16th. The submission instructions are described in the Programming section.

You are to complete this assignment individually. However, you are encouraged to discuss the general algorithms and ideas in the class in order to help each other answer homework questions. You are also welcome to give each other examples that are not in the assignment in order to demonstrate how to solve problems. But we require you to:

- not explicitly tell each other answers;
- not copy answers;
- not allow your answers to be copied.

In those cases where you work with one or more other people on the general discussion of the assignment and surrounding topics, we ask that you specifically record on the assignment the names of the people you were in discussion with (or “none” if you did not talk to anyone else). This will help resolve the situation where a mistake in general discussion led to a replicated error among multiple solutions. This policy has been established in order to be fair to everyone in the class. We have a grading policy of watching for cheating and we will follow up if it is detected.

Instructions for turning in programming component: The last question involves programming which should be done in MATLAB. Do *not* attach your code to the writeup. Make a tarball of your code and output, name it <userid>.tgz where your userid is your Andrew id, and copy it to /afs/andrew/course/15/381-f08/submit/hw3/<userid>/ .

Refer to the web page for policies regarding collaboration, due dates, and extensions.

1 [15 pts] Maximum Likelihood

Our favourite robot, R2D2, has been sent on a mission to the outer world to locate the latest weapon of the dark lord. On its way, it gets partly damaged by an enemy ship damaging its encryption mechanism.

R2D2 wants to send out the range of the enemy weapon. Since he is worried that the message might be intercepted, he instead sends a series of randomly generated values $y_1, y_2, y_3, \dots, y_m$ which are related to ϕ , the actual range of the weapon in the following manner: $P(Y = y|\phi) = \phi e^{-\phi y}$ (y 's are always positive).

Unfortunately, the Jedi knights, powerful though they might be, are somewhat lacking in their knowledge of math and ask the good students of 15-381 to help them out in this task. Can you help them defeat the dark lord?

1. Write down the likelihood(probability) of the data $y_1 \dots y_m$ assuming they were independently generated from the distribution.
2. Determine the value of ϕ that maximizes the likelihood of the data, the is the MLE estimate. Recall from the discussion in the class that the MLE estimate of a parameter is the value of the parameter that maximizes the likelihood of the data. One way to obtain the maximum value of a function, is to solve for its derivative being zero, estimate the value of the function at that point. This is guaranteed only to give you a stationary point(it could be a minimum point for example). You'll need to then show that it is a maxima and not a minima.

Warning: The Jedi can read your mind, but the TAs are somewhat lacking in that aspect(in other words, show us your working :))

2 [15 pts] HMM miscellany

1. Derive the update formula that was shown in class for $\beta_t(i)$.

$$\beta_t(i) = P(O_{t+1}, \dots O_T | q_t = s_i) = \sum_j a_{j,i} b_j(O_{t+1}) \beta_{t+1}(j)$$

Use $\beta_T(i) = 1$ for all i to initiate the function (T is the total number of values observed.)

Hint: Sum over all possible states in time t+1 and then use Bayes rule.

2. Show how to derive the equation for $S_t(i)$ (the probability of being in state i at time t) which was used in class as part of the EM algorithm.

$$S_t(i) \triangleq P(q_t = s_i | O_1, \dots O_T) = \frac{\alpha_t(i) \beta_t(i)}{\sum_j \alpha_t(j) \beta_t(j)}$$

Hint: Use chain rule as shown on slide. Then use chain rule again for the numerator. Note that the denominator is a sum of the numerator for all possible states.

3 [15 pts] Value Iteration

Consider an MDP with 6 states: Bathroom, Kitchen, Bedroom, Dining Room, Under-Attack, and Dead. This represents a domain of a robotic rat foraging for food in a house with four rooms (Bathroom, Kitchen, Bedroom, and Dining Room). In states Bathroom, Kitchen, Bedroom, and Dining Room, there are three available actions: stay in place (S), move horizontally (H), and move vertically (V). While foraging, the rat may be attacked by a robotic cat that also inhabits the house, which causes the rat to enter the Under-Attack state. From the Under-Attack state there is only one action, Die. In state Dead there is only one action, Stay Dead. The rewards and transition probabilities are as follows:

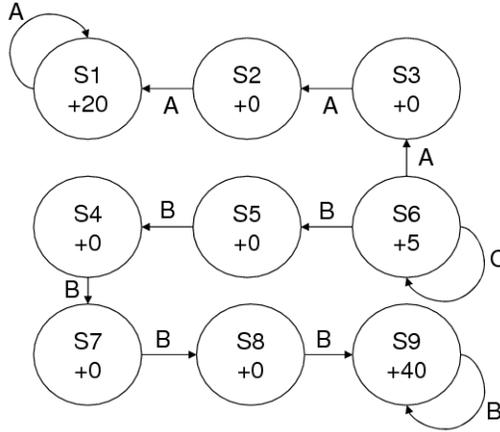
T(s,a,s')						
s, a	s'					
	Bathroom	Kitchen	Bedroom	Dining Room	Under-Attack	Dead
Bathroom, H	0	0.6	0.4	0	0	0
Bathroom, V	0	0.4	0.6	0	0	0
Bathroom, S	0.75	0	0	0	0.25	0
Kitchen, H	0.6	0	0	0.4	0	0
Kitchen, V	0.4	0	0	0.6	0	0
Kitchen, S	0	0.75	0	0	0.25	0
Bedroom, H	0.4	0	0	0.6	0	0
Bedroom, V	0.6	0	0	0.4	0	0
Bedroom, S	0	0	0.75	0	0.25	0
Dining Room, V	0	0.6	0.4	0	0	0
Dining Room, H	0	0.4	0.6	0	0	0
Dining Room, S	0	0	0	0.75	0.25	0
Under-Attack, Die	0	0	0	0	0	1.0
Dead, Stay Dead	0	0	0	0	0	1.0

s	R(s)
Bathroom	+4
Kitchen	+10
Bedroom	0
Dining Room	+2
Under-Attack	-50
Dead	0

1. What is the total number of possible policies?
2. Trace value iteration by hand on this problem. Initialize the value of each state with 0. As robotic rats are somewhat shortsighted, use $\gamma = 0.5$. Give the values for all states after each iteration. You may stop after 6 iterations.
3. What is the optimal policy given your values for the states?

4 [15 pts] Discount Rate

Consider the following MDP.



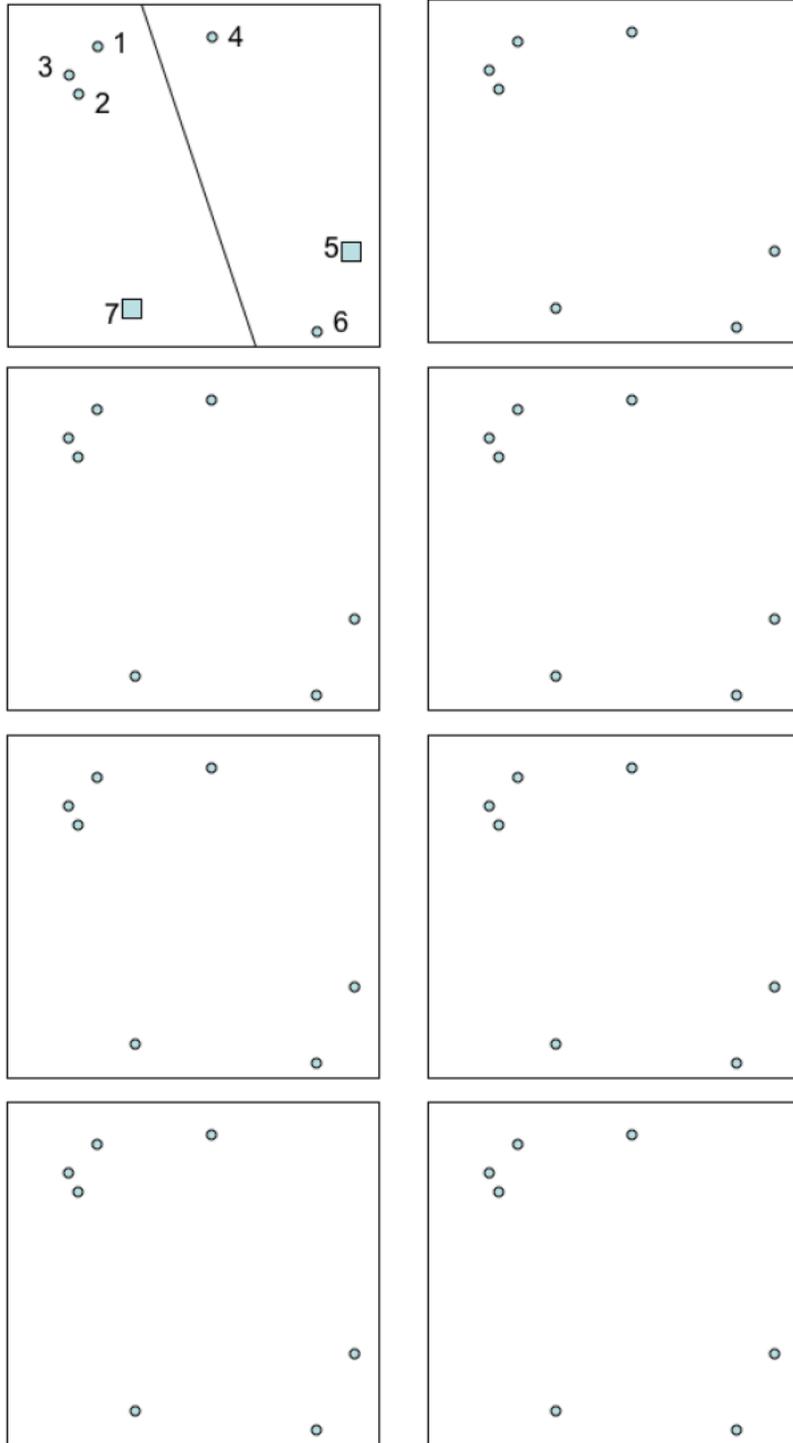
Actions A, B, and C can be taken from state S6. All transitions are deterministic, and the immediate rewards for the states are given in the nodes. Let π^* denote the optimal policy.

1. If $\pi^*(S6) = A$, would you expect γ to be small, moderate, or large? Why?
2. If $\pi^*(S6) = B$, would you expect γ to be small moderate, or large? Why?
3. If $\pi^*(S6) = C$, would you expect γ to be small, moderate, or large? Why?
4. Find the specific ranges of γ (if they exist) so that $\pi^*(S6)$ will be A, B, and C. (For example, one incorrect solution is $\pi^*(S6) = A$ when $\gamma \in [0, 0.9]$, $\pi^*(S6) = B$ when $\gamma \in [0.9, 0.99]$, and $\pi^*(S6) = C$ when $\gamma \in [0.99, 1.0]$; another incorrect solution is $\pi^*(S6) = A$ for all γ .) You may find it helpful to recall the formula for a geometric series:

$$\sum_{i=0}^{\infty} ar^i = \frac{a}{1-r} \text{ if } r \in [0, 1)$$

5 [10 pts] Clustering

1. Perform K-means clustering on the dataset given below. Circles are data points and there are two initial cluster centers, at data points 5 and 7. Draw the cluster centers (as squares) and the decision boundaries that define each cluster (we've already done this for the initial state, as shown in the first figure). If no points belong to a particular cluster, assume its center does not change. Use as many pictures as you need for convergence.



2. Give one advantage of hierarchical clustering over K-means clustering, and one advantage of K-means clustering over hierarchical clustering.

6 [30 pts] Forward-Backward in HMMS

Cecilia and Slackernerney, two graduate students in Markovford, have gotten into an argument over the weather. Cecilia thinks summer has ended and that we are into Autumn(Fall), while Slackernerney (who has been spending all his time in a room without windows) thinks that this is summer. Since they weren't able to resolve it themselves, they have decided to ask the fine students at CMU to help them decide.

A bit about Markovford: there are three seasons in Markovford – Summer(S), Autumn(A) and Winter(W). Given the season the previous day, the season on a day is conditionally independent of the season on all previous days. The weather is either Hot(H), Rainy(R) or Freezing(F). Given the season on any given day, the weather that day is independent of all other variables.

More formally, if we let q_i denote the season on the i^{th} day (taking values S,A,W) and O_i denote the observed weather pattern(one of H,R,F); the model $q_1 \dots q_N, O_1 \dots O_N$ is an HMM as covered in class.

1. Does the weather in Markovford reflect normal weather patterns? Why, or why not?
2. Implement forward-backward in MATLAB, for this HMM. In order to fully specify the HMM, you will need to fill the values for the CPT. These are in the tranprob and emitprob matrices of 'model' structure saved in input.mat (use load(input.mat) to get the structure; model.tranprob and model.emitprob are the respective probability distributions). readme.txt contains the numerical mapping of states. You will also need the initial state probabilities; please assume them to be uniform, ie $P(q_1 = A) = P(q_1 = S) = P(q_1 = W) = 1/3$
3. Cecilia has made 20 observation of the weather over the last 20 days(ie $O_1 \dots O_N$). This data is in the weather array of the model structure(model.weather).
 - (a) Apply inference, to compute the probability of Autumn for each of these days (ie $P(q_i = A|O_1 \dots O_N) \forall i$) and write down the result you get. You must include this output both in your written submission and in a file called ps3_6a.txt and include it in your tarball.
 - (b) Cecilia made the last observation today. Who do you think was correct – Cecilia or Slackernerney?
 - (c) What is the probability that it is currently winter($P(q_N = W|O_1 \dots O_N)$) in Markovford?

Submit your code and a readme describing briefly what each file does, along with the output listed above. Your code musn't modify input.mat so I can test it with other inputs(by replacing this input.mat with a different file)

For those who don't plan to use matlab, the same data is stored in input.txt in human readable format. readme.txt has an explanation of input.txt. Again, your code should take in input.txt as a command line argument.

A note of caution on the programming: The transition probabilities in the data are defined to be from state i to state j and not j to i. So you'll need to transpose the matrix if you want to use the equation in question 2.1 to compute β_t .