CS 683 Artificial Intelligence
Fall 2010
Homework 3
Due Monday Nov 15th by 5 pm

General Notes

- **Show your calculations:** if your final answer is correct, your calculations assure me that you got it yourself. If the final answer is wrong, your calculations help me find out where you went wrong and give you as much credit as possible for the parts you got right.
- You can turn in your homework handwritten (leave it in the TA's mailbox by 5pm) or typed and emailed to the TA, again by 5pm. If you're going to write it by hand, please make sure your answers are legible.

In both the MDP and POMDP questions, we assume a grid world with the following layout. For ease of reference, we'll give each state a name from a to k.

a	b	c	d
e	■	f	g
h	i	j	k

Taking an action in this grid world has probabilistic outcomes. The action 'Left' moves left with probabiliy 0.8, or up or down with probability 0.1 each. If a move will result in the agent hitting a wall (or going into the black forbidden cell), the agent does not move.

a) MDP

For the above grid world, assume the following:
1) The cost of an action is -0.2
2) The discount factor is 0.9

0.4	0.96	1.41	2
-0.15	■	0.58	-3
-0.69	-0.28	0.08	-0.5

The above shows the values of the states after running the value-iteration algorithm for a few iterations. Your task is to **calculate the value at the end of the next iteration for state h**.

b) POMDP

For this question, the initial belief state is given below.

0.1111	0.1111	0.1111	0
0.1111		0.1111	0
0.1111	0.1111	0.1111	0.1111

Show the updated belief state after the agent executes the action **'Left'** in the initial belief state.

<u>Notes:</u>

- Your final answer should be a grid with numbers in it.
- Show where the number in each cell comes from (i.e. show where the updated belief about every state came from).
- As you can see, there are no observations in our grid world. So you can say there is only 1 observation and it has probability 1 in all states.
- In calculating the updated belief, do not worry about normalization.

c) HMM

This question asks you to determine the most likely sequence of (hidden) states responsible for a given sequence of observations.

Suppose you have a friend who lives in an area where the weather has strange patterns and the carrying of an ubmrella is sometimes advisable. However, carrying an ubmrella is cumbersome and your friend doesn't do it all the time. You don't know the weather in your friend's country, but your friend texts you daily with wether she carried an umbrella that day. After 4 days (and 4 text messages), your task is to make an educated guess as to the most likey weather in that far away country over the last 4 days (e.g. sunny, sunny, rainy, sunny). The possible weather states are

<u>S</u>unny, **<u>R</u>**ainy, **<u>W</u>**indy, **<u>C</u>**loudy

You have some knowledge of the weather patterns in the area where your friend lives, so you know the probability of, for example, a sunny day given that the previous day was rainy. Your knowledge is given in the following transition probability table (an explanation of state I is given in the notes below):

Day	Next Day	Prob
I	S	0.3
I	W	0.7

S	R	0.2
S	W	0.8
R	C	0.7
R	R	0.3
W	C	0.1
W	S	0.9
C	C	0.7
C	S	0.3

You also know your friends habits when it comes to carrying an umbrella. For each kind of weather, you know her probability of carrying an ubmrella in that weather. This knowledge is given in the following table:

Weather	P(U|Weather)
S	0.3
R	0.7
W	0.5
C	0.8

The 4 text messages you receive give you the following sequence of observations, where U=carrying umbrella and ~U=not carrying one:

U ~U ~U U

Determine the most likely sequence of weather conditions in your friend's area.

Notes:

- For this question we'll be using an HMM with a known start state that is a dummy state we'll call **I**. I only has outgoing edges. We'll say that at state I, you can receive only 1 observation, let's call it I_o, with certainty, so $P(I_o|I)=1$. This special observation cannot be seen at any other state, so for any state k that is not I, $P(I_o|k)=0$.
- From the above, the observation sequence you'll do your calculations with is actually (I_o, U, ~U, ~U, U).
- Somewhere in your answer, I expect to see a table with rows for states and columns for time unit 0 (corresponding to the dummy observation I_o) and the 4 time units. Time unit 0 will have 1 for state I and 0 elsewhere.
- In this table, I need to see which previous state(s) you got the entry of a given state from.