CS 4300/6300 **HW8: Markov Models, HMMs, and VPI**

Hand in via gradescope.

1 Markov Chains

Alice decides that she wants to keep re-taking her AI class every semester for the rest of eternity (she *really* likes AI). We're interested in modeling whether she passes the class or not as a Markov chain. Suppose that if in semester t she passes the class, then in semester $t+1$ she passes the class with probability 0.8 (maybe she gets bored and forgets to pay attention). On the other hand, if she doesn't pass in semester t then she'll pass it with probability 0.4 in semester $t+1$.

1. Suppose that in semester $t = 0$ Alice passes the class with probability 0.5. Compute the probability that she passes in semester $t = 1$ and semester $t = 2$.

2. Compute the stationary distribution of this chain.

2 Alice and the Crazy Coke Machine

Alice is up late studying for her AI class final exam and wants to stay hydrated (and caffienated!). Unfortunately, her school bought a really wonky soda machine. All you can do with this soda machine is put money in and hope it gives you the type of drink you want. It carries three types of soda: Coke (C), Diet Coke (D) and Sprite (S).

Alice has been monitoring the soda machine for a while and has figured out that it behaves as an HMM. It has two internal states (call them A and B). When asked for a soda from state A, it gives a coke with probability 1/2, and a diet coke and a sprite each with probability 1/4. On the other hand, when asked for a soda in state B, it gives a diet coke with probability 1/2, a sprite with probability 1/3 and a coke with probability 1/6. Each day it changes state. It transitions from state A to A with probability 4/5 and from state B to B with probability 3/5.

The machine formally works as follows. It is in some state s_t on day t. Every time someone puts in money, it dispenses a soda according to the probability rules set out above. At the beginning of each day, it (randomly) transitions to a new state s_{t+1} according to the transition probabilities above.

1. Yesterday, Alice's friend Bob says he bought an undisclosed, but large, amount of soda and as a result he said he was 80% confident that the machine was in state A. Today, Alice puts money in and gets a Sprite out. Assuming she believes her friend Bob, what is Alice's probability distribution over soda machine states now?

3 (CS 6300 only) Decision Networks and VPI

A	P(A)
e	0.5
l	0.5

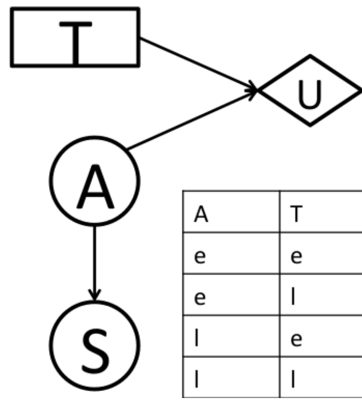

S	P(S)
e	
l	

A	S	P(S|A)
e	e	0.8
e	l	0.2
l	e	0.4
l	l	0.6

A	T	U(A,T)
e	e	600
e	l	0
l	e	300
l	l	600

S	A	P(A|S)
e	e	
e	l	
l	e	
l	l	

Your parents are visiting you for graduation. You are in charge of picking them up at the airport. Their arrival time (A) might be early (e) or late (l). You decide on a time (T) to go to the airport, also either early (e) or late (l). Your sister (S) is a noisy source of information about their arrival time. The probability values and utilities are shown in the tables above.

1. Fill in the above empty table entries by computing $P(S), P(A|S)$ and then compute the quantities below.

   (a) $EU(T = e)$

   (b) $EU(T = l)$

   (c) $MEU(\{\})$

   (d) Optimal action with no observations

2. Now we consider the case where you decide to ask your sister for input. You will walk through the calculation of $VPI(S)$ by computing the quantities below.

   (a) $EU(T = e|S = e)$

   (b) $EU(T = l|S = e)$

   (c) $MEU(\{S = e\})$

   (d) Optimal action with observation $\{S = e\}$

   (e) $EU(T = e|S = l)$

   (f) $EU(T = l|S = l)$

   (g) $MEU(\{S = l\})$

   (h) Optimal action with observation $S = l$

   (i) $VPI(S)$