CS141 **Introduction to AI** **Professor Meinolf Sellmann**

Homework 8

Bayes and Markov

Due: 5:00pm on 4/14/08

Problem 8.1

Consider a Markov model with the state space $S = \{1A, 1B, 2A, 2B, 3A, 3B, T\}$ and action set $A = \{\alpha, \beta\}$. Here is the reward function:

	α	β
1A	0	100
1B	0	200
2A	0	50
2B	0	400
3A	0	0
3B	0	600
T	0	0

As for the probability transition function, $P(T|T) = 1$, and $P(T|\beta) = 1$. Transition probabilities of non-terminal states for the α action are:

	2A	2B	3A	3B	T
1A	.4	.6			
1B	.6	.4			
2A			.4	.6	
2B			.6	.4	
3A					1
3B					1

Perform policy iteration on this Markov model, showing the state values, action values, and recommended policy after each iteration.

Initialize the policy $\pi(s)$ to β for each state s.
Use a discount factor of 1 (in other words, no discount).
In the policy iteration step, if $Q(s, a)$ is the same for both actions, default the action to β.
Do as many iterations necessary for the policy to converge.

Problem 8.2

The informational entropy of a probability distribution is defined as the - expected-value $\log(p(X))$. (Use a base 2 log)

a) Calculate the entropy of a fair coin toss (one that is heads half the time and tails the other half.

b) Calculate the entropy of a biased coin that is heads 75% of the time.

c) Explain in your own words what you think entropy measures and explain why the answer to (1) is greater than the answer to (2).

Problem 8.3

a. For the Bayesian network given in Russel and Norvig, p. 494, apply variable elimination to the query $P(Burglary|JohnCalls = true, MaryCalls = true)$. Follow the walkthrough of calculations on p. 507 and give the result of each intermediate calculation. In other words, provide your computed values for $f_M(A)$, $f_{\bar{A}JM}(B, E)$,$f_{\bar{E}\bar{A}JM}(B)$, and $P(B|j, m)$.

b. Prove that the complexity of a running variable elimination on a polytree network is linear in the size of the tree for any variable ordering consistent with the network structure.