UIN: Name:

CSCE-625 Final exam

The exam is composed of 5 questions each weighing 20%
Every question is followed by the expected answer format in parenthesis. Please follow this format.
You may use a simple calculator. You may not use a computer or a phone. Write your final answers in ink (no pencils).
Submit all pages of this exam.
Write your UIN and name at the top of each page.

1. (20pts) Consider the following two-player, zero-sum, turn-based game where the MAX player is first to play, then it's the MIN player's turn, and finally MAX makes another move. Outcomes' utilities are presented at the leaves.
   a. What will be the game outcome if both players are fully rational? (numerical value)

      Answer: 10

   b. What subtrees are pruned by an alpha-beta run from left to right? (set of letters) Specify subtrees by the affiliated preceding letter. E.g., the subtree rooted in the bottom left MAX node is denoted by the letter "b". A subtree is pruned iff all its composing nodes are never visited while the successor of its rout is visited.

      Answer: l

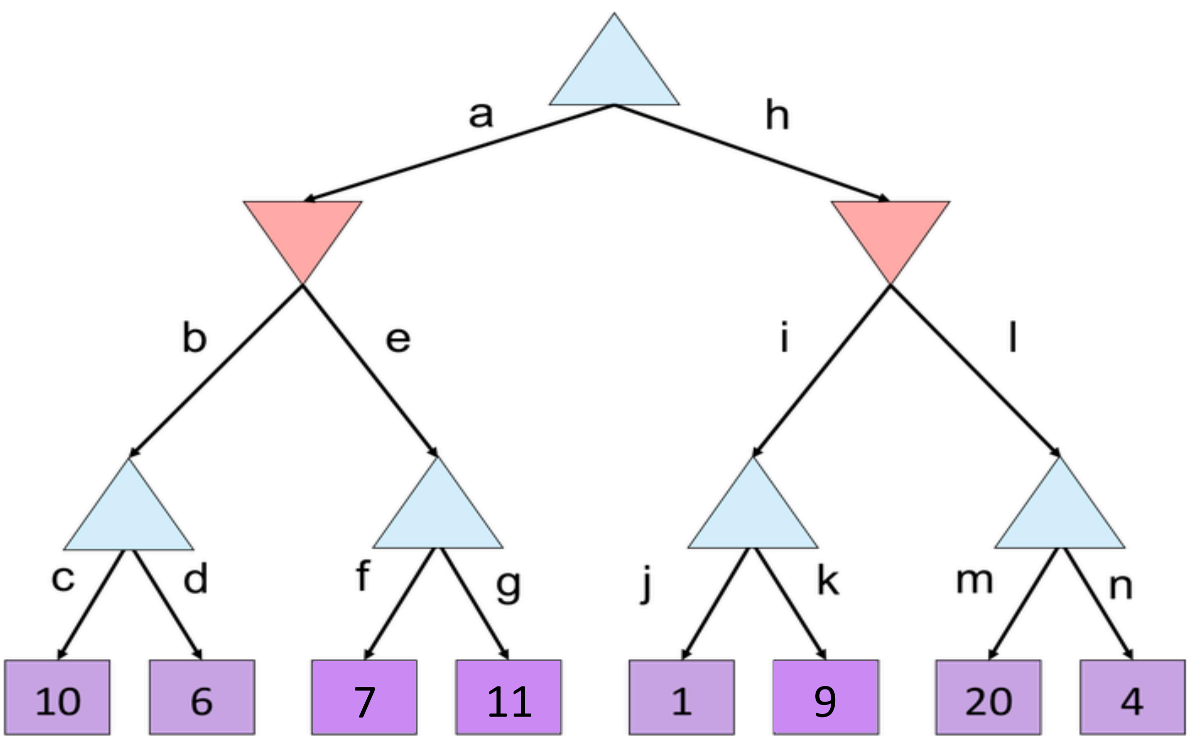

UIN: Name:

2. (20 pts) Consider running the A* algorithm with duplicate detection for finding the shortest path (least number of steps) leading from state S to state G in the state space provided below. State identifier is given in blue, heuristic value is given in red. Assume tie braking in favor of earlier letters in the ABC… order.
   a. What is the solution returned by A*? (ordered set of states)

      Answer: S,b,l,G

   b. What states are expended by A* and in what order? (ordered set of states)

      Answer: S,a,b,d,c,l,G

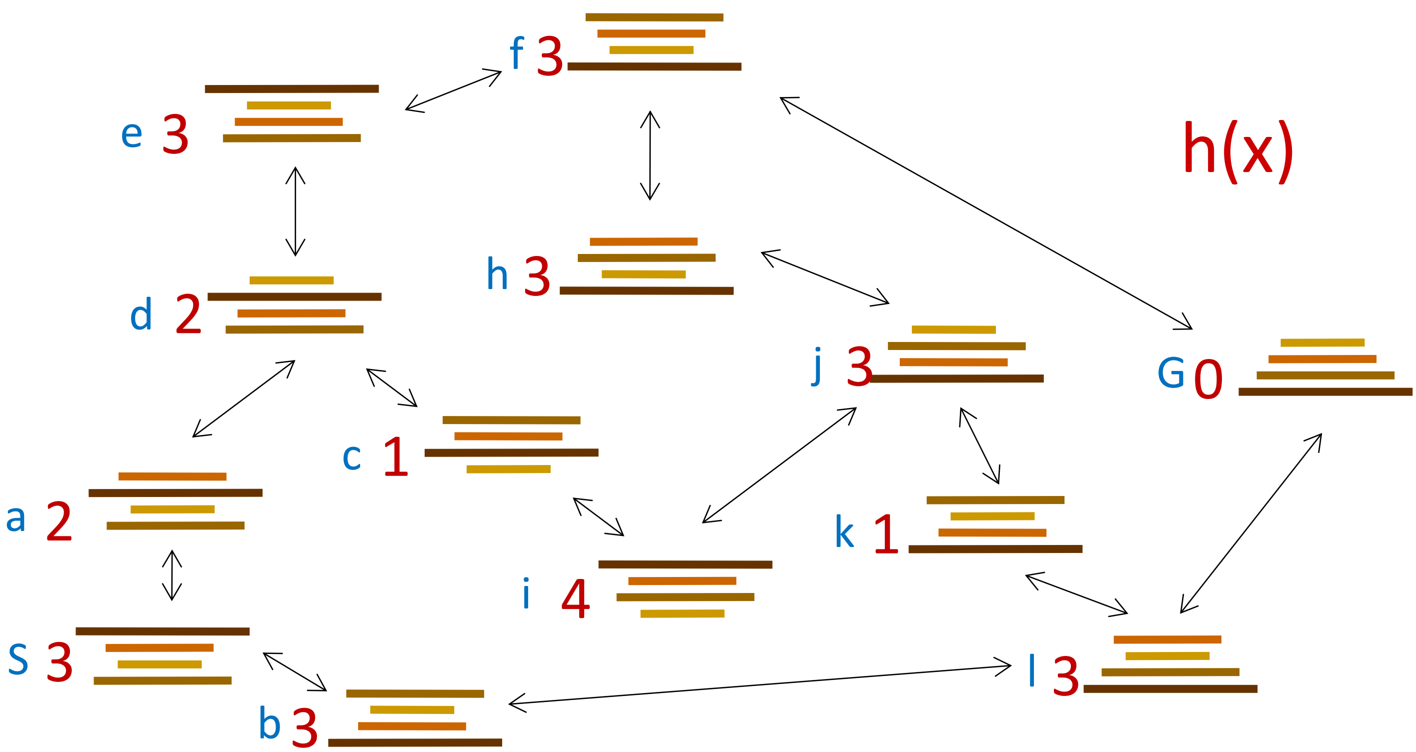

UIN: Name:

3. (20pts) Assume an MDP with 4 states (a,b,c,d). At each nonterminal state (a,b) the agent can choose one of two actions {West, East}. The model is unknown so we will use TD learning in order to learn Q values and find the optimal policy. Assume a discount factor (gamma) of 0.9, learning rate (alpha) of 0.1, and that Q values are initiated to 0. Two episodes are performed (episodes given as [state1,action1,reward2,state2,action2,reward3,state3]):

I. [a,East,-1,b,East,10,c]

II. [a,East,-1,b,West,-100,d]

a. What are Q(a,East), Q(b,East) after the first episode? (2 numerical values with 1 decimal digit)

   Answer: -0.1, 1

b. What are Q(a,East), Q(b,East) after the second episode? (2 numerical values with 2 decimal digit)

   Answer: -0.1, 1

$$sample = R(s,a,s') + \gamma \max_{a'} Q(s',a')$$

$$Q(s,a) \leftarrow (1-\alpha)Q(s,a) + (\alpha)\,[sample]$$

or

$$Q(s,a) \leftarrow Q(s,a) + \alpha\big(sample - Q(s,a)\big)$$

Or

$$Q(s,a) \leftarrow Q(s,a) - \alpha(Q(s,a) - sample)$$

UIN: Name:

4. (20pts) Assume a training set (x,y) of 2 sentences:
x1="win the vote" y1=Politics
x2="win the election" y2=Politics

Assume a multiclass perceptron that considers the count of each word as a feature in a vector [#win, #game, #vote, #the]. E.g., x1 is represented by the feature vector [1,0,1,1]. Assume the following initial weights vectors that are affiliated with each class:

w_{SPORTS}

win	: 1
game	: 1
vote	: 0
the	: 1

$w_{POLITICS}$

win	: 0
game	: 0
vote	: 1
the	: 0

w_{TECH}

win	: 1
game	: 0
vote	: 0
the	: 0

a. What are the values of the three weight vectors (w_s, w_p, w_t) after training the multiclass perceptron on x1? (3 vectors)

Answer: $[0\ 1\ -1\ 0]^{\mathsf{T}}\ [1\ 0\ 2\ 1]^{\mathsf{T}}\ [1\ 0\ 0\ 0]^{\mathsf{T}}$

b. What are the values of the three weight vectors (w_s, w_p, w_t)) after training the multiclass perceptron on x1 and then x2? (3 vectors)

Answer: $[0\ 1\ -1\ 0]^{\mathsf{T}}\ [1\ 0\ 2\ 1]^{\mathsf{T}}\ [1\ 0\ 0\ 0]^{\mathsf{T}}$

UIN: Name:

5. (20pts) Consider the HMM depicted below. The hidden variable is the presence of rain in a given day and the observed variable is the presence of an umbrella in a given day. Assume that the initial probability for rain ($Rain_0$) is p(+r)=0.4, p(-r)=0.6.
   a. What is the probability of observing rain in day 1 given that an umbrella is observed in day 1? (numerical value with 3 decimal digits)

      Answer: 0.793

   b. What is the probability of observing rain in day 2 given that an umbrella is observed in day 1 and in day 2? (numerical value with 3 decimal digits)

      Answer: 0.879

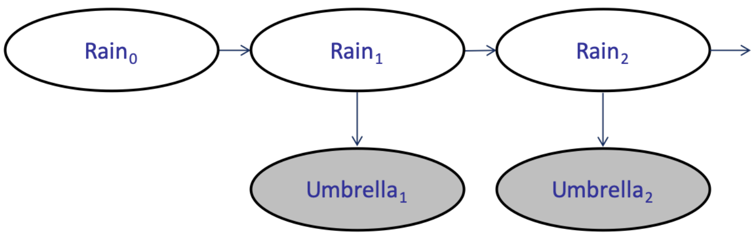

R_t	R_{t+1}	$P(R_{t+1}\|R_t)$
+r	+r	0.7
+r	-r	0.3
-r	+r	0.3
-r	-r	0.7

R_t	U_t	$P(U_t\|R_t)$
+r	+u	0.9
+r	-u	0.1
-r	+u	0.2
-r	-u	0.8

Passage of time: $B'(X_{t+1}) = \sum_{x_t} P(X'|x_t)B(x_t)$

Observation: $B(X) \propto P(e|X)B'(X)$