Final Examination

CS540: Introduction to Artificial Intelligence

December 19, 2006 Instructor: Jerry Zhu

CLOSED BOOK
(Two letter-size notes allowed. You don't need to hand them in)

LAST (FAMILY) NAME: _____

FIRST NAME: _____

Problem	Score	Max Score: 5 points each
1	_____	
2	_____	
3	_____	
4	_____	
5	_____	
6	_____	
7	_____	
8	_____	
9	_____	
10	_____	
11	_____	
12	_____	
13	_____	
14	_____	
15	_____	
16	_____	
17	_____	
18	_____	
19	_____	
20	_____	
Total	_____	100

1. Iterative deepening search

Run iterative deepening search on the following graph, starting from node 1. Node 4 is the goal node. Write down the node numbers as they are expanded. Do NOT use any loop avoidance method. Everything being equal, expand node with the lowest number.

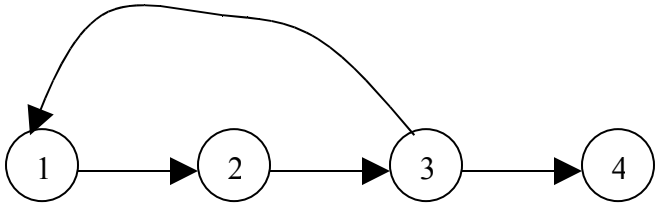

2. Optimization

In the following electric network, each node is connected by resistors (the edges) of unit resistance to some other nodes. Each node (except node x) is fixed at the voltage shown in the circle. The heat generated by the electric network is

$$E = \sum_{i \sim j} (v_i - v_j)^2 ,$$

where the sum is over all resistors, i and j are the two nodes connected by the resistor, and v_i, v_j are the voltages at node i, j respectively. What should the voltage x be to minimize the heat?

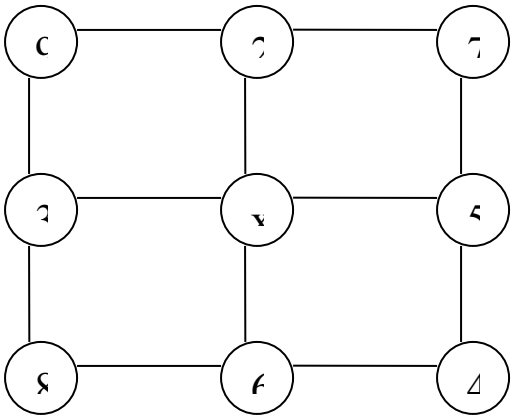

3. Heuristic search

In standard A* search the objective function at each node n is f(n)=g(n)+h(n), where g(n) is the cost from start to this node, and h(n) is an admissible heuristic from n to goal. Now let us use a different objective function:

$$f(n) = (1-w)*g(n) + w*h(n)$$

where w>=0.

a) What kind of search does this perform when w=0?
b) What about when w=0.5?
c) What about when w=1?
d) For what values of w is this algorithm guaranteed to be optimal?

4. Nash equilibrium

Circle all pure Nash equilibria in the following zero-sum game. A is the max player and B is the min player.

		B: I	B: II	B: III
A	I	-4	-7	-3
	II	9	1	7
	III	-6	-1	5

5. Matrix normal form with Chance node
Write down the matrix normal form of the following zero-sum game with chance node.

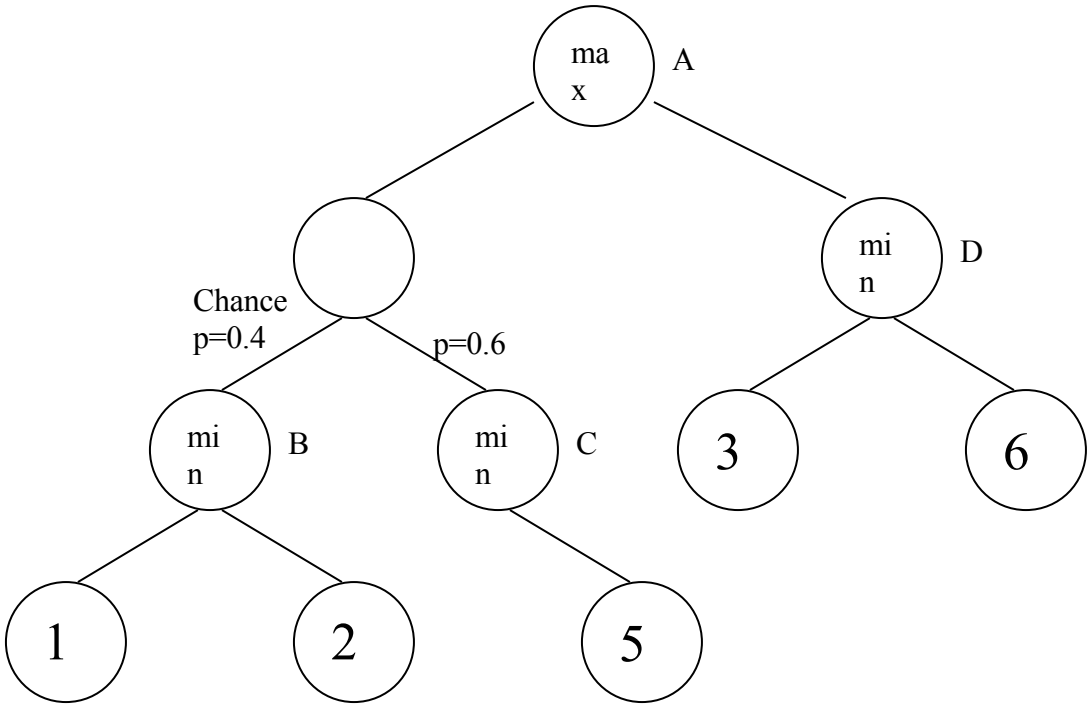

6. Strictly dominated strategy

Perform iterative elimination of strictly dominated strategies on the following matrix normal form for non-zero sum game. Each cell contains (A's score, B's score). Note in this notation, both players want to maximize their own score.

	B		
A	2, 4	3, 7	4, 5
	1, 2	5, 4	2, 3
	4, 0	2, 8	5, 3
	3, 6	4, 0	1, 9

7. Propositional logic

The propositional logic connective NAND is defined as:

P	Q	P NAND Q
T	T	F
T	F	T
F	T	T
F	F	T

Express (P OR Q) using and only using NAND, P, Q. (Hint: you need 3 NAND's. What is P NAND P?)

8. Unification
Write down the Most General Unifier for the following pairs of FOL clauses.

A. See(I, Duck, Telescope), See(x, y, z).

B. Like(x, Hate(Mary, Mary)), Like(Hate(y, y), x).

C. Like(Hate(y, y), x), Like(x, Hate(Mary, John)).

D. Wiser(Father(x), x), Wiser(Father(y), Tom).

E. Knows(Self(x), x), Knows(y, y).

9. Given KB:

$P \Rightarrow Q$

$Q \Rightarrow R$

Use resolution refutation to prove $\neg P \vee R$. Clearly write down all the steps.

10. K-nearest-neighbor

Given the following dataset:

{(x=1, y=+), (x=2, y=-), (x=4, y=-), (x=8, y=+), (x=16, y=-)}

a) What is the leave-one-out accuracy with 1-nearest-neighbor classifier? Use Euclidean distance.
b) What is the leave-one-out accuracy with 3-nearest-neighbor classifier? Use Euclidean distance.

11. Entropy

Compute the entropy (in bits) of the following loaded 6-sided die: p(side1) = ½, p(side2)=1/4, p(side3)=p(side4)=1/8, p(side5)=p(side6)=0. Show your steps.

12. Speech recognition
In one sentence, describe the idea of Captcha.

13. Mutual information
Given the dataset

x	y
1	+
1	-
1	-
1	+
0	-
0	-
0	+
0	+

What is the entropy H(y)? What is the conditional entropy H(y|x)? What is the mutual information I(x;y)? Use bits as the unit.

14. K-means Clustering

Given dataset {1, 2, 3, 6, 8, 10, 20, 60, 100} and three cluster centers c1=2, c2=7, c3=12

a) List the initial cluster assignment (i.e. which point is in which cluster).
b) Compute the new cluster centers after a)
c) List the cluster assignment after b)
d) Compute the new cluster centers after c)
e) List the cluster assignment after d)

15. Support Vector Machines (SVM)

Let X=(x1,x2) be a point in 2-dimensional space. Given two points

X1=(0,0), y1=+1

X2=(100,100), y2=-1

X'W + 1=0 defines a line, where W=(w1,w2). What values of w1 and w2 separates X1, X2 with the largest margin?

16. Perceptron
In question 7 we defined NAND. Let T=1 and F=0. Using a perceptron with step function activation (a(x)=0 if x<=0, 1 otherwise), implement NAND. You need to 1) draw the perceptron, 2) mark the inputs, 3) define the weights.

17. Neural network
In a neural network with x inputs, one hidden layer with y units, and one output layer with z units, assuming each input is connected to all hidden units, and each hidden units is connected to all output units, how many weights are there in the neural network altogether?

18. Probability

In a bag there are two envelopes, one has a red ball (worth $100) and a black ball (worth nothing), the other has two black balls. You randomly chose an envelope.

a) From the envelope your friend randomly took out one ball, and told you "it's black". Given the chance, should you switch to the other envelope? What is the probability that you get the red ball if you switch?
b) Instead of a), your friend looked into the envelope, and told you "there is at least one black ball". Given the chance, should you switch to the other envelope? What is the probability that you get the red ball if you switch?

19. Statistics

In your 10-day vacation in Alaska, you kept the following log on the weather and whether you saw a bear that day:

(rain, bear) 1 day
(¬rain, bear) 2 days
(rain, ¬bear) 6 days
(¬rain, ¬bear) 1 day

a) Write down the joint probability table, estimate probability with frequency.
b) Compute the marginal probability P(bear).
c) Compute the conditional probability P(¬bear | rain)

20. Bayes Net
Given the Bayes Net, compute P(C|B). Show your steps.