


Name:

NetID:


1.) For the graph given below answer the questions. The cost to travel between nodes are given as/on the solid lines and the heuristic distances are on the dashed lines.



a.) Is this heuristic admissible? Why or why not?

b.) Is this heuristic consistent? Why or why not?

2.) Given the following graph, starting in Node S and trying to get to Node G, trace the steps the A* Algorithm would take. To make the graph less busy, the heuristic distances are given as a table.



*.) $N \rightarrow N' = d(N, N') + h(N', G)$

1.) $S \rightarrow A = 1 + 6 = 7$

2.) $S \rightarrow B = 4 + 2 = 6$

3.)

4.)

5.)

6.)


7.)

8.)


9.)

10.)

3.) Show which tuple of numbers gets chosen for each level of a 4 player minimax tree.




4.) Use Alpha/Beta Pruning on the tree below, showing the Alpha and Beta at each node and show which branches don't need to be explored.




5.) Answer the following two questions about Constraint Satisfaction Problems (CSPs).

a.) Explain in your own words how the backtracking algorithm for CSPs works and is implemented.



b.) How does forward checking or constraint propagation improve on basic backtracking?



6.) Given the following dataset layout, answer the two questions.



		S	X		Y
		S	X_1	X_2	Y
Train	S_1	3.9	4.2	0	
	S_2	2.7	8.5	1	
	S_3	4.1	4.0	0	
	...				
Test	S_{N-1}	3.1	7.9	1	
	S_N	4.0	3.8	0	

a.) How do we derive the mu and sigma for a gaussian PDF during the training or fitting step of a Naive Bayes classifier?

b.) What is the role of the gaussian PDF during the prediction step of using a Naive Bayes classifier and what assumption do we make by using a Gaussian distribution?

NOTE: I don't expect any math here, just use words and pictures if you like.


7.) What is the difference between a regression task and a classification task?

8.) Why is the Viterbi algorithm considered a “greedy” algorithm


9.) What would the equation look like for a multiple linear regression with three features?

10.) What are the two regularization techniques we learned and how do they affect the weights of a linear regression?


11.) What is the purpose of the AC-3 Algorithm?

A large, empty rectangular box with a black border, intended for the student's answer to question 11.

12.) What is the “learning” a Markov Babler does when given an input text?

A large, empty rectangular box with a black border, intended for the student's answer to question 12.

13.) What are MLE and MAP?

A large, empty rectangular box with a black border, intended for the student's answer to question 13.

14.) What is the difference between uninformed and informed searches?

15.) How does IDS simulate BFS behaviour with a DFS?

16.) What is a loss function?