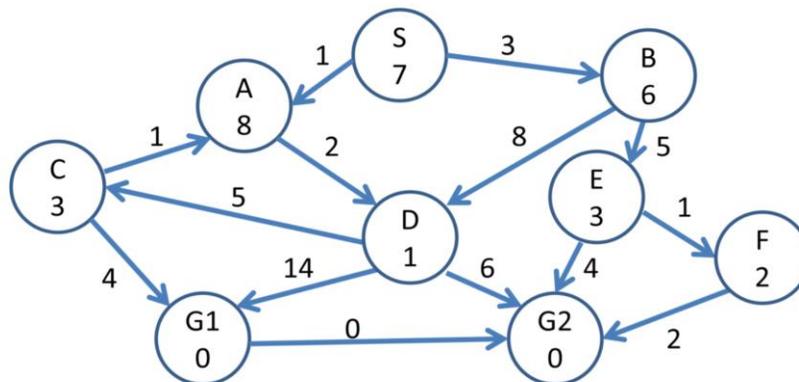


CS 4700: Foundations of Artificial Intelligence
Sample exam questions

1. True/False Questions:

- a. If the third state visited by both depth-first and breadth-first search is the same (starting from the same initial state with the same goal) then they must have visited the same second state.
- b. If you don't check for cycles then breadth-first search can fail to find a solution.
- c. If iterative deepening finds a solution for a given problem then depth-first search will also find that solution.
- d. If breadth-first search is able to find a solution to a search problem, then A* is guaranteed to find a solution.
- e. If f_1 , f_2 , and f_3 are all admissible functions, then $\frac{f_1}{6} + \frac{f_2}{3} + \frac{f_3}{2}$ is admissible.
- f. Alpha-beta pruning will never result in a different move being selected compared to plain minimax search, but the value it generates for the current state may differ from that of plain minimax search.
- g. A perceptron is guaranteed to learn any set of training data given a suitable learning rate.
- h. Stochastic gradient descent is a form of gradient descent that occasionally makes small random changes to the weights of your neural network.

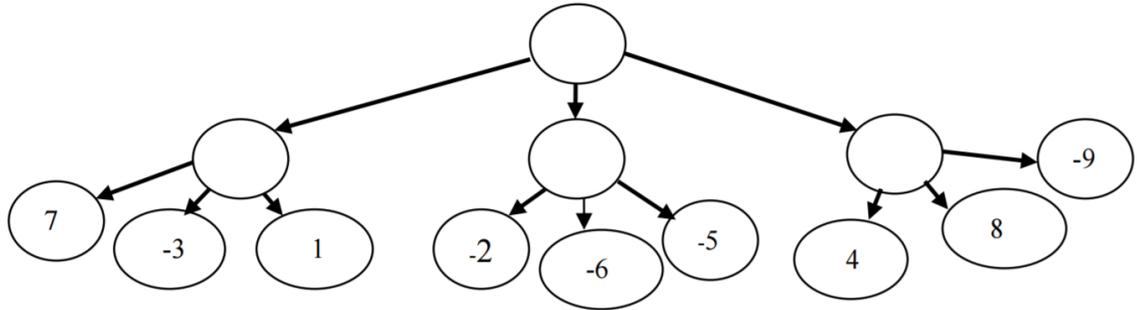
2. Consider the following search space, where S is the initial state and G1 and G2 are goal states, where the cost of the operator that takes you from one state to the next is given along the edge between the states, and where the value of the heuristic evaluation function h applied to the state is the number written inside the state. In cases of ties pick the state that comes earlier in the alphabet.



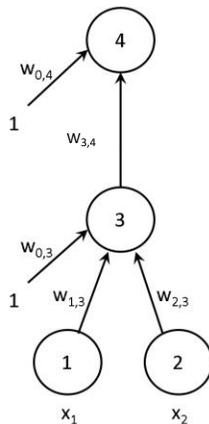
- a. Give the sequence of states that depth-first search visits.
- b. Give the sequence of states that iterative deepening visits.
- c. Give the sequence of states that hillclimbing with $f(s) = g(s) + h(s)$ visits (lower values are better).

- d. Give the sequence of states that A* visits if it uses $f(s) = g(s) + h(s)$.
- e. Is $h(s)$ admissible?

3. Apply the minimax algorithm to the game tree below, where it is the **opponent's** turn to move next and the leaf nodes are terminal nodes whose values are given in each node in the figure. Process this game tree working left-to-right.



- a. Write the values that minimax gives the intermediate nodes inside their circles.
 - b. Circle the outgoing arc of the root node that represents the move that minimax search would select for this game.
 - c. Put X's through the nodes that would be pruned by alpha-beta pruning.
4. Imagine you trained a perceptron and after learning was completed you discovered that the data were all labeled backwards: Every example that was labeled 1 should have been a 0, and vice versa. Sadly, you no longer have the data and have no ability to change the code that uses the perceptron to flip its answers. All you have access to are the weights of the perceptron. How would you change the weights in order to flip all the answers?
5. Consider the following two-node neural network:



For the purpose of this question nodes 3 and 4 use the sigmoidal activation function, and all of the weights are set to 1.0. In your answers do not compute the decimal representation of the value, keep it in symbolic form in terms of e .

- a. Given the example whose two inputs are $x_1=x_2=0.0$, what is the activation of node 4?
- b. Recall:
 - If u is a node in the output layer $\Delta_u = g'(In_u)(y-a_u)$
 - If u is a node in a hidden layer it is $\Delta_u = \sum_{v=0}^n w_{u,v}g'(In_u)\Delta_v$, where v iterates over all nodes that u feed into.
 - The weight update equation is $w_{u,v} = w_{u,v} + \alpha a_u \Delta_v$.

If the desired output for the example in part a is 0 and you run it through backprop:

1. What is Δ_4 ?
2. What is Δ_3 ?
3. If $\alpha = 1$ what is $w_{0,4}$ after the weight update?