F21 AI - Final Exam

This is an open exam, meaning that you are free to use books, notes, computers, and the internet. However, any form of collaboration is prohibited. Complete the exam, and email it as a PLAIN TEXT under the subject "AI Final" to nzhou@brooklyn.cuny.edu by 4:30pm on Friday, December 10.

Q1

Indicate whether each of the following statements is true or false.

1. Breadth-first search is a special form of best-first search where the cost g(n) is the number of steps from the initial state to state n.
2. Depth-first graph search is may be incomplete even if the state space is finite.
3. If A* uses an inadmissible heuristic, then a solution found by it may not be optimal.
4. Best-first search is a special form of A* search.
5. For a CSP (Constraint Satisfaction Problem), depth-first tree search can guarantee finding a solution if the CSP is satisfiable.
6. In Poker, it's more likely to have a full-house hand (three of a kind and two of another kind) than a straight hand.
7. For a random variable representing an event, the more its entropy is, the more uncertain the event is.
8. Although linear regression problems can be solved analytically, it is more practical to use gradient descent learning to obtain parameters when data matrices are large.
9. There are functions that cannot be approximated by using a neural network, even with multiple layers of neurons.
10. In natural language processing, ambiguity may occur in all phases, including speech recognition, syntax analysis, semantic analysis, and pragmatic analysis.

Q2

Consider the *Tower of Hanoi problem* as searching for a plan in the state space for transforming a given initial state to a given goal state. The actions in the plan must obey the following rules:

1. Only one disk can be moved at a time.
2. Only the top disk on any peg can be moved.
3. Larger disks cannot be stacked above smaller disks.

Apply the A* algorithm to the following problem instance, using the number of the misplaced disks as the heuristic. Show the states in *frontier* and the $f(n)$ value of each of the states after three expansions.

Each state is represented as a list of three lists, each of which gives the numbers of the disks from top to bottom on a peg. For example, the initial state in the above Figure is represented as `[[1,3], [2], []]`.

Q3

A Latin square of order n is defined to be an $n \times n$ matrix made out of n^2 integers 1, 2, ..., n^2 with the defining characteristic that each of the n symbols occurs exactly once in each row and each column. For example, the following is a Latin square of order 3.

```
3  2  1
1  3  2
2  1  3
```

1. Give a CSP model for finding a Latin square of order n.
2. (Extra 5 points) Implement the model in a language of your choice.

Q4

Consider the following Bayesian network:

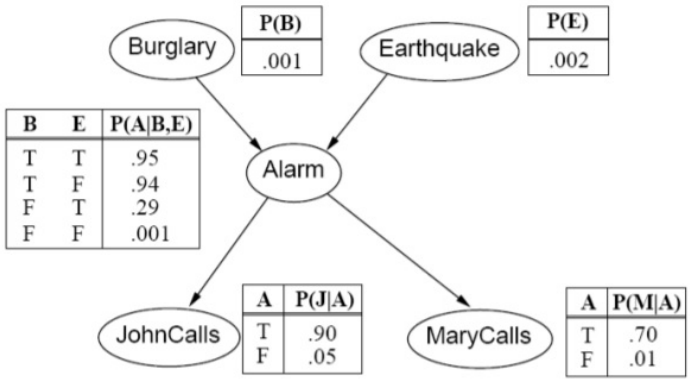

Compute each of the following:

1. P(Burglary = T, Earthquake = T)
2. P(Burglary = T | Alarm = T)

You can only give a formula without doing the actual computation.

Q5

Consider the following feedforward neural network. Assume that the bias weights of H_3, H_4, and O_5 are, respectively, W_{03}, W_{04}, and W_{05}, and the bias inputs are always 1.

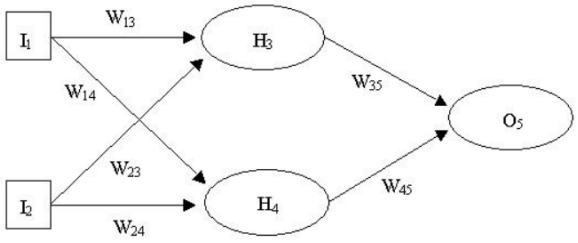

1. Write a function in a language of your choice that simulates the neural network. The function takes I_1 and I_2 as inputs, and returns O_5 as the output. Assume that the activation function used by the first layer is g_1, and the activation function used by the second layer is g_2. The function computes the output based on the inputs, the weights, and the activation functions.

2. Assume the following weights:

W_{03}	-1
W_{13}	2
W_{23}	2
W_{04}	3
W_{14}	-2
W_{24}	-2
W_{05}	-3
W_{35}	2
W_{45}	2

Also assume that both g_1 and g_2 are the step function, which outputs 1 if the input is greater than or equal to 0, and 0 otherwise. Give the function represented by the neural network as a truth table or a formula.