Fall 2018, Final Exam
CSCE 4613, Artificial Intelligence

1. (5 points) What is your name?

2. (5 points) Consider the game tree to the right. What will be the final score if both players use a mini-max strategy?

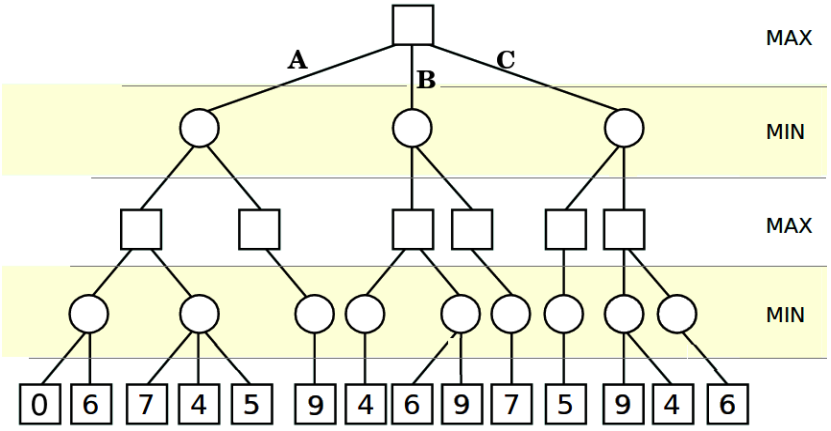

3. (10 points) Consider the game tree to the right. Which choice will the first player make if it uses a mini-max strategy(A B or C)?

4. (10 points) Identify the type of each algorithm as:

A. Reinforcement learning
B. Unsupervised learning
C. Supervised learning
D. A planning algorithm

___ An autoencoder (used for reducing the dimensionality of data).
___ A mouse is given cheese or electric shocks, depending on its choices.
___ D* Lite (an improvement upon A*)

___ Random Forest

___ Q-learning

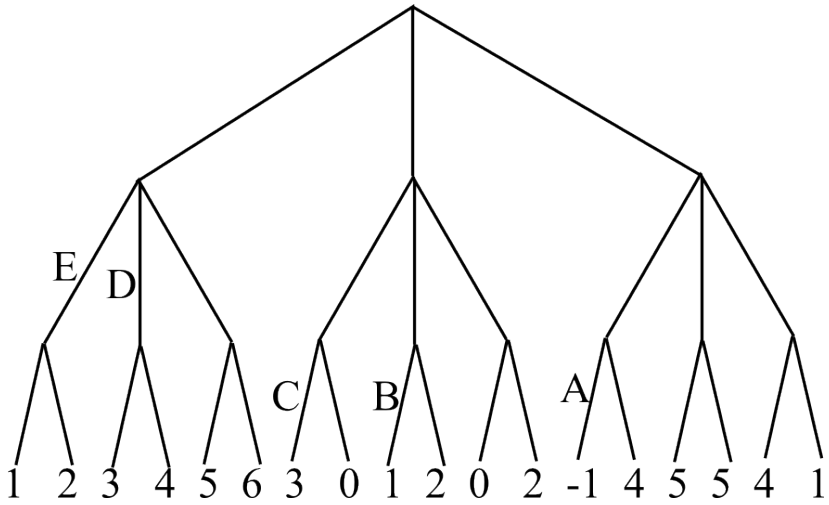

5. (5 points) The game tree to the right (and up a little) represents a zero-sum game with outcomes shown in the leaf nodes. Player 1 (who goes first) wants to minimize the outcome. Player 2 (who goes second) wants to maximize the outcome. Suppose this tree is evaluated in a depth-first manner from **right to left** with alpha-beta pruning. Which branch, A, B, C, D, or E would be evaluated. (That is, which one would not be pruned by alpha-beta pruning. Choose only one.)

$$Q(i,a) \leftarrow (1-\alpha^k)Q(i,a) + \alpha^k \left[r(i,a,j) + \gamma \max_{b \in \mathcal{A}(j)} Q(j,b) \right]$$

6. (10 points) On the Q-learning equation above, draw an arrow pointing at the variable that represents the current state. (If that variable occurs more than once in the equation, just pick one.) Underline the function that computes the immediate reward that the agent will receive. Circle the discount factor, which is responsible for determining how the agent prioritizes long-term versus near-term rewards. Draw a box around the learning rate, which controls how much the Q-table is updated from each experience. Draw a squiggley line under the part of the equation that evaluates future actions. Draw an 'X' over the symbol that represents the Q-table.

7. (10 points) Which of the following AI techniques,

A. A* search
B. A genetic algorithm
C. Mini-max with alpha-beta pruning
D. An ensemble of decision trees
E. Utility theory

would be best-suited for addressing each of these challenges:

___ Finding a building architecture that can withstand strong winds.
___ Diagnosing medical conditions

___ Solving a Rubick's Cube

___ Finding weights for an artificial neural network
___ Determining the business decisions that will minimize costs.
___ Finding a reasonable solution to the travelling salesman problem
___ Playing checkers

___ Finding the shortest path between two points in a video game
___ CAPTCHA breaking (optical character recognition)
___ Cheating at gambling

___ Maze running

8. (10 points) Please put an A next to the properties of A* Search and a G next to the properties of Genetic Algorithms, B for Both, or N for Neither.

____Guarantees to find the optimal solution

____Typically fast with a small number of dimensions

____Scales well with a large number of dimensions

____Requires a lot of memory (proportional to the size of the state space)

____Easy to parallelize

____Requires the user to supply a function that evaluates candidate solutions.

9. (20 points) Suppose you wake up in the middle of the night with a strong craving for graham crackers. Unfortunately, there are no graham crackers in your pantry! You will not be able to sleep until you satisfy your craving, or determine that both Walmart and Harps are sold out of graham crackers.

You know that Walmart is selling a box of graham crackers for $3, and Harps is selling the same brand of graham crackers for $2. Walmart is exactly 10 miles due east of your home. Harps is *n* miles due west of your home. You estimate that driving one mile costs you about $0.1 (including the cost of gas, the value of your time, the cost of wear on the car, etc.)

You estimate that there is a 0.3 chance that Walmart will be sold out of graham crackers, and a 0.4 chance that Harps will be sold out of graham crackers.

(Assume that sleep is more important to you than any amount of money, the roads are perfectly straight, there are no hills or stop lights, the two stores are in exactly opposite directions, Paris is the capital of France, and pineapple is a perfectly valid pizza topping.) For what values of *n* would it be more cost-effective to drive to Harps first?

Please draw a nice labeled tree to represent this problem before you start doing the math. The tree will help me follow your math, and will give me some basis for awarding partial credit. The root decision in your tree should be whether to drive to Walmart first or Harps first. If you get graham crackers, you can just go home. If not, you must visit the other store before returning home to get some sleep.

10. (10 points) Please describe how bagging (bootstrap aggregation) is done. (Don’t write a book about it, just describe the major operations.)

11. (5 points) Describe a problem that Q-learning could handle effectively that breadth-first-search might not handle effectively. (Don’t write a book about it, just describe the part of the problem that makes the difference.)