Spring 2018, Midterm 2
CSCE 4613, Artificial Intelligence

1. (20 points) What is your name?

2. (10 points) Suppose you want to race in the Tour de France. You have already purchased a non-refundable round-trip plane ticket to France for $750 and an entry fee of $150. If you decide to proceed with your racing plans, you will need to purchase a high-performance speed bike for $35,000. You estimate that there is a 0.2 chance that you will win the race, and a 0.8 chance that you will lose. If you win, you expect to be offered $70,000 in endorsement contracts from foreign advertisers. If you sign with advertisers in the States instead, you can get $50,000. However, if you sign with foreign advertisers, you will have to pay extra taxes of $5000. You can decide not to report the foreign income on your Tax return, but you estimate that there is a 0.15 chance that you will be audited by the IRS and the fraud discovered. If that happens, you will have to appear in court and pay a fine for tax evasion. The total cost to you if that happens is $x, (but you will have already profited from your celebrity endorsements.) Please draw a tree to represent the decisions and contingencies of this problem.

3. (5 points) Continuing from the previous problem, for what values of x is it to your financial advantage to attempt to sign with foreign advertisers?

4. (5 points) If x is $0, should you still attempt to race, or just visit the Eiffel Tower (assuming you only care about maximizing financial return)?

5. (5 points) Which statement is more correct?

- When playing the Iterated Prisoner's Dilemma, minimax is a poor strategy.
- When playing the Iterated Prisoner's Dilemma, minimax is the best strategy.

6. (10 points) Consider the following two-player zero-sum turn-taking game:

- This game requires 4 marshmallows to be placed on a table.
- Player 1 goes first. Player 2 goes second.
- To take a turn, a player must draw 1 or 2 marshmallows and eat them. The player must hum the Star-Spangled Banner until the marshmallows are consumed.
- When there are no marshmallows left on the table, the game is over.
- Each player's score is computed as the number of marshmallows they ate minus the number of marshmallows their opponent ate. The Last marshmallow is worth 5 extra points for the person who eats it.

Please draw a game tree to represent all the possible outcomes in this game. Indicate the scores for both players at each leaf node.

7. (10 points) At each node in the tree that you drew for the previous problem, embolden (darken) the branch that represents the choice that each player will choose, assuming that each player will play only to maximize his or her own personal score, and will ignore the opponent's score. (Note that this is not the mini-max strategy, but it just happens to make the same decisions as mini-max in this particular game.) So, one branch extending from each node should be emboldened. If there is only one branch extending from a node, it should be emboldened. If there is ambiguity due to multiple branches yielding the same utility, choose the first one.

8. (5 points) In the marshmallow game from problem 6, what will be the final score?

9. (5 points) Can A* search offer computational improvements over breadth-first search if all the actions have uniform cost?

- Yes
- No

10. (5 points) Given these two "chromosomes" in a genetic algorithm, perform single-point cross-over to produce a new "chromosome":

CATAGATACAAG
TATACAATGTAG

11. (5 points) Why does random forest use decision trees that make random divisions?
(Choose one.)

- Random divisions cause the trees to make better predictions.
- Random divisions cause the trees to be more inconsistent with each other.
- Random divisions do the best job of reducing entropy in the data.
- Because human doctors give random advice.
- Bootstrap Aggregation specifies to use random divisions.

12. (10 points) The tree below represents a zero-sum turn-taking game. The branches that would be pruned by alpha-beta pruning if the choices are evaluated in left-to-right order are already marked. Please draw an "X" on the branches that would be pruned by alpha-beta pruning if the choices are evaluated in right-to-left order instead.

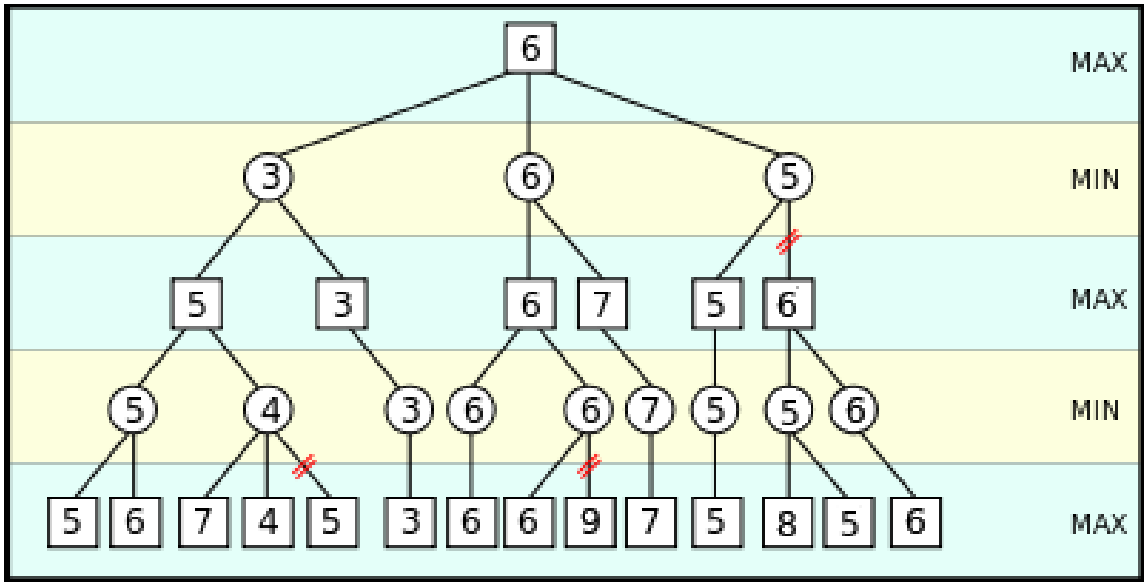

13. (5 points) How does bagging work?

- Train one random model, but make several copies, and test them all on different data.
- Train half of the models, then combine them to produce the remaining models. Use the model that yields the best predictive accuracy to represent the ensemble.
- Divide the data into two parts, called "train" and "test". Train each model on a different non-overlapping subset of the training data, and use the test data to identify the best model.
- Train each model with the same data, but then test them all on different data.
- Train each model with data that has been resampled (with replacement), then give each model one vote in the final prediction.
- It is theoretical ideal that does not actually work in practice.