

Final Exam

600.464/664 Artificial Intelligence
Spring 2019

Name:


Instructions

- Please be sure to write both your name and section in the space above!
- Please be sure to read through the entire exam before you start, and be mindful of the time. If one question is taking too long, it may be worth moving on and coming back to the problem question(s) after the rest of the exam is complete.
- Remember that you are only allowed one sheet (both sides) of notes, everything else besides that and the test itself must be off of your workspace.
- Please show ALL relevant work for your answers and provide explanations when prompted. Failure to do either will result in loss of credit.
- You will have until 5pm to complete this test.


Basic Search

10 points

1. (5 points) Number the nodes (1, 2, ...) in the order they are explored in depth-first search. Process children left to right.



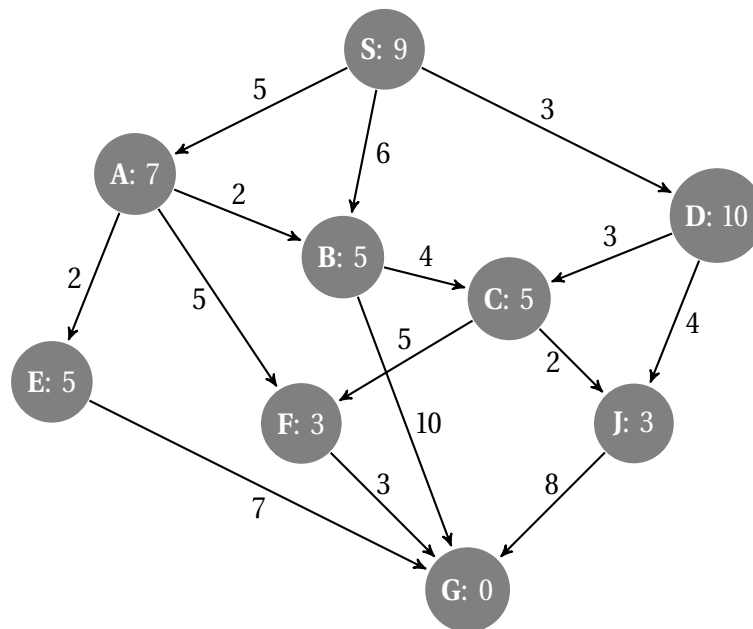
2. (5 points) Number the nodes (1, 2, ...) in the order they are explored in iterative deepening search. Enter multiple numbers, if a node is visited more than once. Process children left to right.



Informed Search

10 points

Consider the search space below, where **S** is the start node and **G** satisfies the goal test. Arcs are labeled with the cost of traversing them and the heuristic cost to the goal is reported inside nodes (so lower scores are better).



3. (10 points) For A* search, indicate which goal state is reached at what cost and list, in order, all the states popped off of the OPEN list. You use a search graph to show your work. Do not expand paths that revisit states at higher cost.

Note: When all else is equal, nodes should be removed from OPEN in *alphabetical* order.

Path to goal (cost):

States popped off of OPEN:

First Order Logic

20 points

4. (6 points) Convert the following sentences into first-order predicate calculus logic:

If a student gets a good grade in an exam in a class, they are happy.

If a student takes a exam for a class, studied for the exam in the class, and is smart, they get a good grade.

If a student takes a class, they take the exam.

Jane takes the class Artificial Intelligence.


Jane studied for the exam in Artificial Intelligence.

Jane is smart.

5. (6 points) Convert all rules to Conjunctive Normal Form (CNF). You do not need to restate rules that are already in CNF.

6. (8 points) Carry out a resolution proof of the statement *Jane is happy*. Hint: start with $\neg \text{Happy}(\text{Jane})$.

Consider the following plot of data points (circles and squares).




10. (5 points) Build a linear classifier $(x, y) \rightarrow shape$. You do not have to get the optimal parameter values for the model but they should be reasonable.

11. (5 points) Draw a decision tree for this data. Will it classify all examples correctly?

Reinforcement Learning

20 points

Consider the *non-deterministic* reinforcement environment drawn below. States are represented by circles, and actions by squares. The Probability of a transitions is indicated on the arc from actions to states. *Immediate* rewards are indicated above and below states. Once the agent reaches the **end** state the current episode ends.



12. (15 points) The current policy is to always take action **X**.

- (a) What paths could be taken?
- (b) What is each path's probability?
- (c) What is each path's reward?
- (d) What is the utility of each state?

13. (5 points) Assuming that under the current policy state **b** has higher utility than state **a**.

- (a) How will a greedy agent that has access to the transition probabilities change the policy?
- (b) Is the updated policy guaranteed to be the best policy?

14. (5 points) Co-reference resolution.

(a) Consider the following sentences.

The boy kicks the ball. He scores a goal.

To which noun (*boy*, *ball* or *goal*) does the pronoun *he* refer to?

How can a computer identify the right noun?

(b) Consider the following sentences.

Jane went to the store to eat a pizza. It was very tasty.

To which noun (*Jane*, *store* or *pizza*) does the pronoun *it* refer to?

How can a computer identify the right noun?

15. (10 points) N-Gram Language models. You are given the following text (tokenized, lowercased) as training data (with `<s>` as beginning-of-sentence marker and `</s>` as end-of-sentence marker).

`<s> the sun is shining . </s> <s> the girl is going to the park with the boy . </s> <s> the boy gives the girl a flower . </s> <s> she is happy . </s> <s> the boy is happy , too . </s>`

You use this data to train a bigram language model with maximum likelihood estimation. What score will it assign to the following sentence?

`<s> the girl is happy . </s>`