

Final Exam

600.464/664 Artificial Intelligence
Spring 2020

Name:

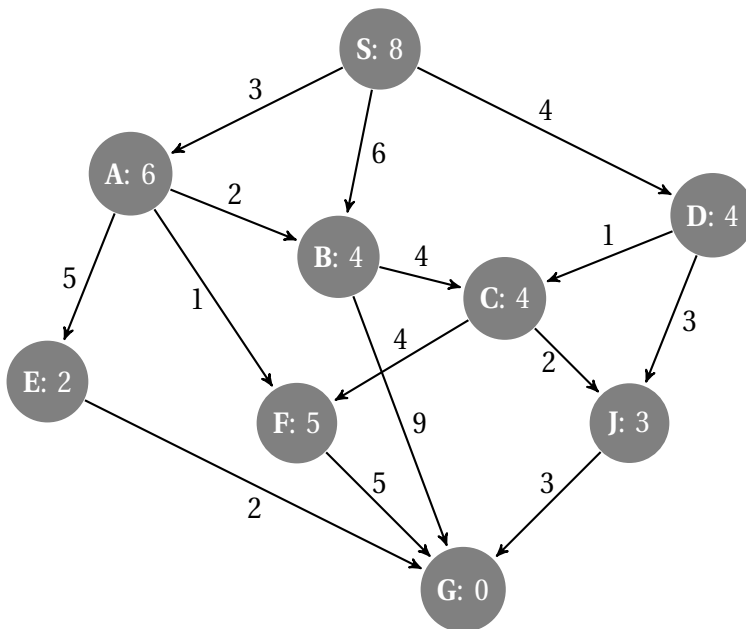
Instructions

- Please be sure to write your name in the space above!
- Please show ALL relevant work for your answers and provide explanations when prompted. Failure to do either will result in loss of credit.
- You may upload a edited pdf file (preferably), a scan of printout with hand-written answers, or any other document with answers.
- The exam is due by midnight on Gradescope.

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).



- (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):

Search graph (indicate order of states popped of OPEN list with numbers):

First Order Logic

15 points

2. (5 points) Convert the following sentences into first-order predicate calculus logic:

If a team wins a game, it is happy.

If a team plays a game, is better than the opponent, and is in good form, it wins the game.

The Blue Jays play a game against the Tigers.

The Blue Jays are better than the Tigers.

The Blue Jays are in good form.

3. (5 points) Convert all rules to Conjunctive Normal Form (CNF). You do not need to restate rules that are already in CNF.

4. (5 points) Carry out a resolution proof of the statement *The Blue Jays are happy.*

Constraint Satisfaction

15 points

Consider the following Mini Sudoku puzzle.
(numbers are assigned values, letters are names for cells)

A	B	1
C	2	D
E	F	G

Recall the rules of Sudoku:

- Each cell is assigned a number (in Mini Sudoku, the numbers are 1, 2, and 3).
- No row can have the same number twice.
- No column can have the same number twice.


These rules can be specified by constraints. You do not need to formally write these constraints down.

5. (8 points) Carry out backtracking search. Explore the search space in the alphabetic order of the cell names, and assign values in numerical order (i.e., first assign 1, then 2, then 3). Draw the search tree. Stop and backtrack at leaves where constraints allow no valid variable assignment.

6. (7 points) Now, use the minimum remaining values (MRV) heuristic. Fill in the following table. Start with a row that specifies the remaining possible values for each variable. Then commit one variable to a value following the MRV heuristic (in case of ties follow alphabetical order).


Step		A	B	C	D	E	F	G
	remaining values							
1	assigned value							
	remaining values							
2	assigned value							
	remaining values							
3	assigned value							
	remaining values							
4	assigned value							
	remaining values							
5	assigned value							
	remaining values							
6	assigned value							
	remaining values							
7	assigned value							

Consider the following plot of data points.



11. (5 points) Write down the formula for a linear classifier function $f(x, y) \rightarrow \text{value}$ where positive output *values* correspond to the sign + and negative *values* correspond to the sign -. The classifier should classify all training examples correctly.

Now consider the following plot of data points.




12. (10 points) Draw a neural network with weight values that classifies all values correctly. You may use the following activation function.

$$\text{threshold}(z) = \begin{cases} 1 & \text{if } z \geq 1 \\ -1 & \text{else} \end{cases}$$

Reinforcement Learning

15 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.



13. (15 points) Consider two possible policies: always take action **X** or always take action **Y**. For each policy, compute the answers to the following questions.
- What paths could be taken?
 - What is each path's probability?
 - What is each path's reward?
 - What is the utility of each state?

Natural Language Processing

15 points

We plan to build a database of which persons were hired at which company from the corpus of news paper articles.

The database table looks like follows:

Name	Position	Company
Joe Johnson	sales representative	IBM
Frida Feldman	lawyer	Microsoft
...

Consider the following sentences

- *Albert Altman was hired by Apple as a software developer.*
- *Booking.com hired Brianna Bayer as an accountant.*
- *Comcast hired as their new CEO Catherine Carter.*

14. (3 points) For each sentence, write a string pattern matching expression that extracts the relevant information (use any pattern matching formalism you like).

15. (3 points) Annotate each sentence with dependency relationships with appropriate labels.

16. (5 points) Define informally a pattern over dependency structures that allows you to extract the relevant information from all the example sentences for the database table.

17. (4 points) Hiring information may be expressed in many different ways in natural language. Below is a list of general problems in natural language processing. Give an example sentence that contains relevant hiring information for our database that demonstrates why each of these problems poses a challenge for our simple extraction patterns.

- Synonymy

- Hypernymy

- Co-reference

- Semantic inference