CSE 415 Winter 2024 Assignment 4

Last name:____________________ First name:____________________________

UWNetID: ____________________

Due Wednesday night February 7 via Gradescope at 11:59 PM. You may turn in either of the following types of PDFs: (1) Scans of these pages that include your answers (handwriting is OK, if it's clear), or (2) Documents you create with the answers, saved as PDFs. When you upload to GradeScope, you'll be prompted to identify where in your document your answer to each question lies.

Do the following four exercises. These are intended to take 30-45 minutes each if you know how to do them. Each is worth 25 points. If any corrections have to be made to this assignment, these will be posted in ED.

This is an individual-work assignment. Do not collaborate on this assignment.

Prepare your answers in a neat, easy-to-read PDF. Our grading rubric will be set up such that when a question is not easily readable or not correctly tagged or with pages repeated or out of order, then points will be be deducted. However, if all answers are clearly presented, in proper order, and tagged correctly when submitted to Gradescope, we will award a 5-point bonus. (Thus the maximum points for A4 will be 105).

If you choose to typeset your answers in Latex using the template file for this document, please put your answers in blue while leaving the original text black.

Version 24-02-03d. (Includes an update to Q3b, and a fix to separate subquestions iii and iv in Question 4 part b, as well as corrections to "Weekend" and "Weekday" descriptive labels in Table 1 of Question 4.) If further corrections to this document are required, this will be announced in ED.

1 Basic Search

Use the following tree to answer the questions below comparing Breadth First Search, Depth First Search, and Iterative-Deepening Depth First Search. Assume that children are visited from left to right.

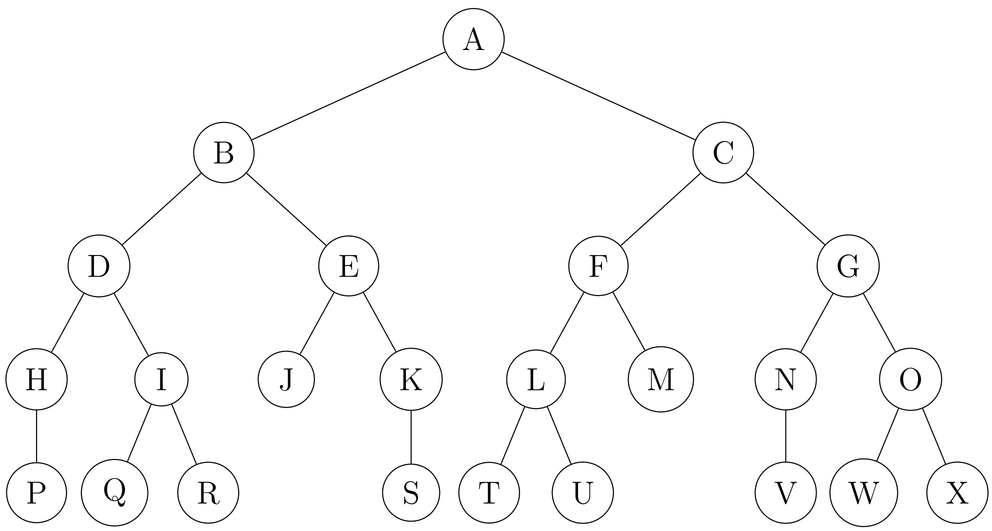

(a) (3 points) Write out the order that the nodes are expanded in using Breadth First Search, starting from A and searching to N.

(b) (3 points) Write out the order that the nodes are expanded in using Depth First Search, starting from A and searching to N.

(c) (3 points) Write out the order that the nodes are expanded in using Iterative-Deepening Depth First Search, starting from A and searching to N. If a node is repeated, make sure to include it each time it is expanded.

(d) (4 points) Which of the three search algorithms (BFS, DFS, IDDFS) has the smallest maximum size of the open list while searching from A to K? What is the maximum size of the open list for that algorithm?

(e) (2 points) True or False: BFS, DFS, and IDDFS will each return the same path on this tree starting at A searching to K.

(f) (2 points) True or False: Given the same start and goal state, BFS, DFS, and IDDFS will always return the same path for any search graph (not necessarily a tree).

(g) (4 points) Suppose you are running your search algorithm on an embedded system device with highly limited memory available. Which of the three search algorithms (BFS, DFS, IDDFS) would you choose to run to minimize the memory used by the algorithm? Justify your answer with a one sentence explanation.

(h) (4 points) Suppose you need to run a search algorithm but you are confident that the vast majority of your search goals will be near the root of the tree (for example, G or earlier in the alphabet on this tree). Which of the three search algorithms (BFS, DFS, IDDFS) would you use to find the shortest path in the least amount of time given these conditions? Justify your answer with a one sentence explanation.

2 Heuristic Search

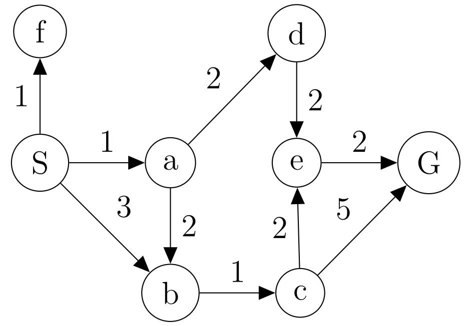

For the following questions, consider three heuristics h_1, h_2, h_3. The table below indicates the estimated cost to goal, G, for each of the heuristics for each node in the search graph.

state (s)	S	a	b	c	d	e	f	G
heuristic $h_1(s)$	6	5	2	2	2	2	6	0
heuristic $h_2(s)$	6	6	4	5	4	1	5	0
heuristic $h_3(s)$	6	6	4	3	4	2	8	0

(a) (2 point) What does it mean for a heuristic to be "admissible"?

(b) (3 points) Which heuristics among $\{h_1, h_2, h_3\}$ shown above are admissible? For heuristics which are not admissible, please identify the nodes where admissibility is violated.

(c) (2 point) What does it mean for a heuristic to be "consistent"?

(d) (3 points) Which heuristics are consistent? For heuristics which are not consistent, please identify the edges where consistency is violated.

(e) (2 points) What does it mean for one heuristic to dominate another heuristic?

(f) (2 points) Of the heuristics provided above, which would you consider the best heuristic? Explain your reasoning in terms of admissibility, consistency, and dominance.

(g) (4 points) Show the node expansion order going from S to G using $A*$ with the heuristic you identified in the previous question. Also provide the final (solution) path obtained.

(h) (1 points) Consider h_1 – what single change would you make to improve this heuristic?

(i) (4 points) Having made the change above, has your determination of the best heuristic changed? Explain your reasoning in terms of admissibility, consistency, and dominance.

(j) (2 points) Explain why, when using the $A*$ search algorithm, it is important to continue the expansion process until the goal state is removed from the OPEN list and becomes the current state, rather than terminating as soon as the GOAL state is found.

3 Adversarial Search and Alpha-Beta Pruning

3.1 Alpha-Beta Pruning

Use this search tree along with the table of state evaluation values to answer the questions below. By default, children are processed from left to right.

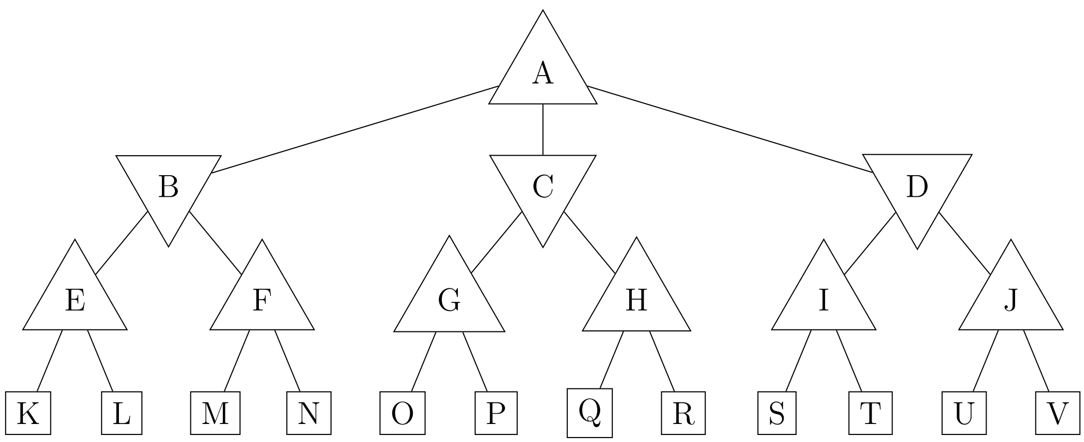

state	K	L	M	N	O	P	Q	R	S	T	U	V
evaluation	5	-3	6	8	-3	-5	6	-4	9	-5	2	1

(a) (8 points) Use minimax to perform adversarial search with alpha-beta pruning on the tree above. Fill in the values in the table below as you go. For α and β values, write the values that are passed from the state's parent to that state. For example, the value of state K will not be shown in the $\alpha - \beta$ values of state E, but would be reflected in the $\alpha - \beta$ values of state L. If a state does not have to be evaluated, do not write any values for it in the table.

(b) (7 points) [Note: This question was updated on Feb. 3. The old text is now crossed out ~~like this~~, and the new text is *emphasized*.] Let's assume that you ~~reverse~~ *change* the order of the leaf nodes in the above tree A (call that ~~reversed~~ tree B). The evaluation function f represents the utilities for the first tree A, and let the evaluation function g represent the utilities for the second tree B.

state	value	α	β
A		N/A	N/A
B			
C			
D			
E			
F			
G			
H			
I			
J			
K	5		
L	-3		
M	6		
N	8		
O	-3		
P	-5		
Q	6		
R	-4		
S	9		
T	-5		
U	2		
V	1		

~~After performing alpha beta on this reversed tree, what observations can you make about the two evaluation functions? Be sure to answer what effect switching the order of the evaluation function utilities has. What are some ways to improve the evaluation functions in general?~~

Is it possible to completely prevent pruning in the tree by just reordering the leaf nodes in the tree? If so, explain why. If not, what is the least number of leaf nodes whose values need to be changed to prevent pruning. In either case, describe your new tree. (Include a drawing if you think it will be more clear that way.)

3.2 Pruning with Chance Nodes

Although alpha-beta pruning cannot be applied directly to searching trees that contain chance nodes, reasoning like that inherent in alpha-beta search can sometimes be applied when static values are constrained to lie within given ranges of values. For example, if all leaf-node values must be values $f(s)$ such that $0 \leq f(s) \leq 10$, then node G can be pruned in the tree below, because Max can get 9 by moving left at A to B, and if Max goes to C and finds $f(F) = 3$, Max can reason that getting $f(G)$ is useless because if $f(G) = 10$ which is largest allowed, then the value at C is $(3+10)/2 = 6.5$, which is inferior to 9.

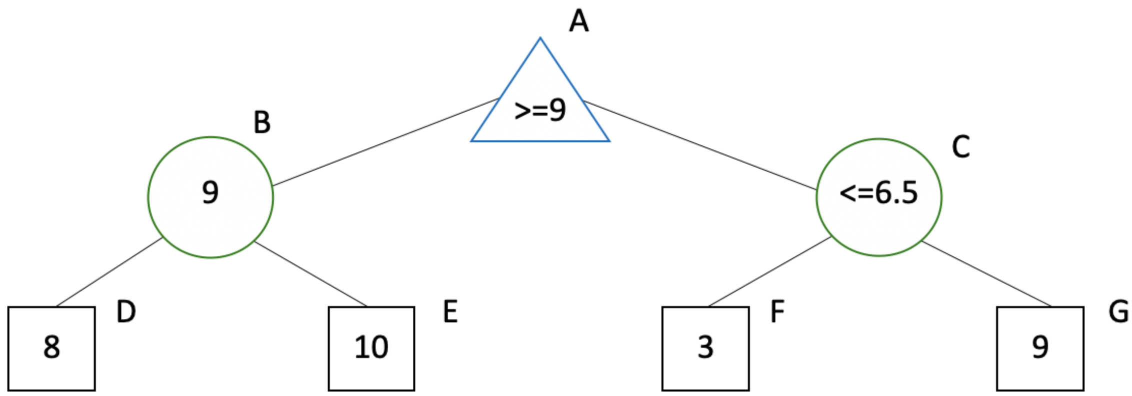

(c) (6 points) In the following game tree, determine where pruning can be performed using the same range assumption as above. Show where there are cutoffs. Also after pruning, change the value of one of the nodes such that the pruning will no longer occur.

(d) (4 points) Explain your reasoning for each cutoff.

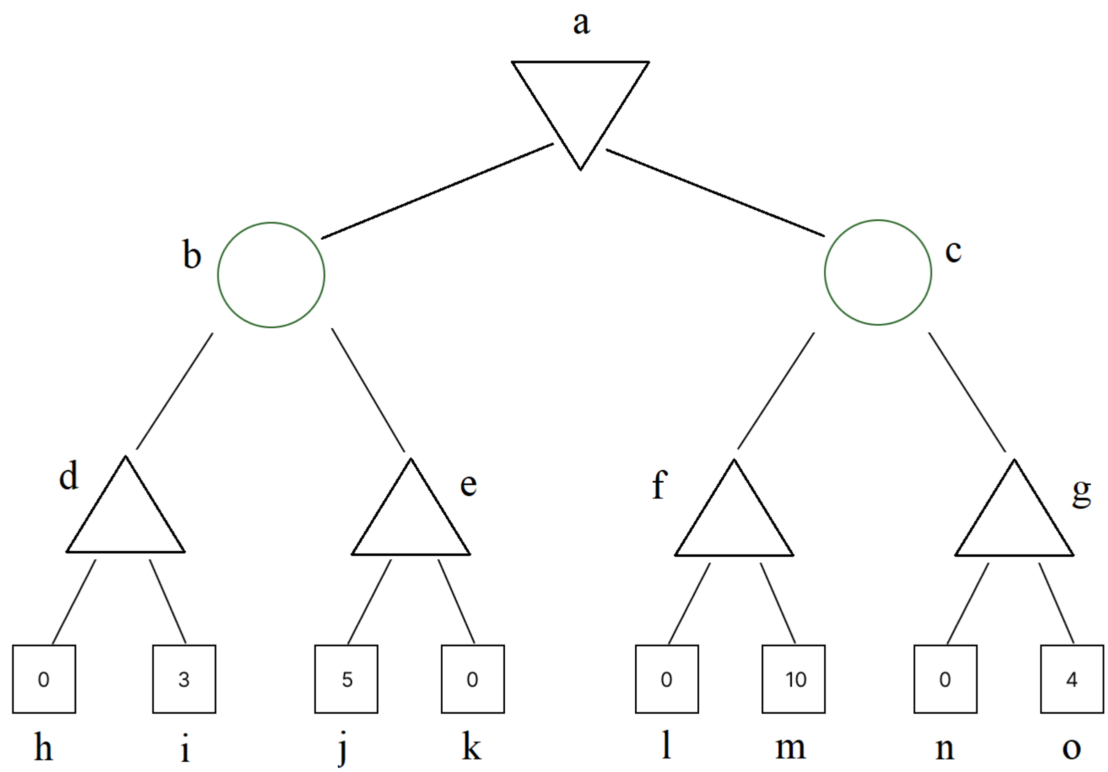
a
b
c
d
e
f
g
0
3
5
0
0
10
0
4
h
i
j
k
l
m
n
o

4 Probabilistic Inference and Factoring Joint Distributions

(a) (10 points) Consider a situation where we are modeling the sale of Ice cream (S), which is dependent on three variables: Weather condition (W) and Day of the week (D). You are given the joint probability distribution of these variables in the table below:

Weather (W)	Day (D)	Ice Cream Sales (S)	P(W,D,S)
Sunny (Su)	Weekday (Wd)	High	0.108
Sunny (Su)	Weekday (Wd)	Low	0.252
Sunny (Su)	Weekend (We)	High	0.216
Sunny (Su)	Weekend (We)	Low	0.024
Cloudy (Cl)	Weekday (Wd)	High	0.024
Cloudy (Cl)	Weekday (Wd)	low	0.096
Cloudy (Cl)	Weekend (We)	High	0.032
Cloudy (Cl)	Weekend (We)	Low	0.048
Rainy (Ra)	Weekday (Wd)	High	0.012
Rainy (Ra)	Weekday (Wd)	Low	0.108
Rainy (Ra)	Weekend (We)	High	0.016
Rainy (Ra)	Weekend (We)	Low	0.064

Table 1: Joint Probability Distribution of Weather, Day, and Ice Cream Sales.

Answer these questions based on the joint probability distribution table above:

(i) What is the marginal distribution P(W) ?

(ii) What is the marginal distribution P(W, D) ?

(iii) What is the marginal distribution P(W, S) ?

(iv) What is the marginal distribution P(D) ?

(v) What is the conditional distribution P(S | W, D)?

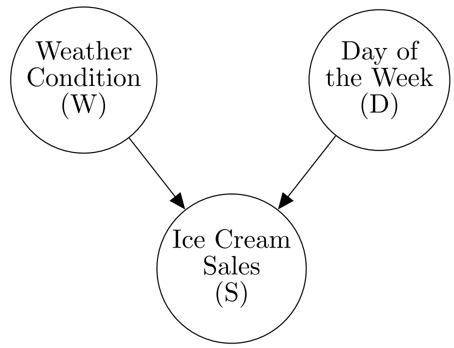

(b) (6 points) Consider the Bayes Net of the same Ice Cream Sales problem we were considering. Answer the following question based on it:

  (i) Write the equation for the joint distribution of the three variables as a product of the appropriate marginal and conditional distributions, according to the Bayes net graph.
  (ii) What is the probability of high ice cream sales on a weekend, given that the weather is sunny?
  (iii) Given that ice cream sales are high, what is the probability that the day is a weekend?
  (iv) What is the probability of having a sunny weekday with high ice cream sales compared to having a rainy weekend with low ice cream sales?

(c) (9 points) Compare the sizes of the representations for the full joint distribution in (a) and the Bayes net in (b) as follows:

  (i) Give an expression for the number of free parameters in Table 1 in terms of the sizes of the domains of the variables W, D, and S. You can use $|W|$ to represent the size of the domain of W, for example.
  (ii) Give the resulting integer value for your expression.
  (iii) Give a corresponding expression for the number of free parameters in the Bayes net in (b).
  (iv) Give the integer value of the expression for the Bayes net.
  (v) Is there any savings in storage, according to the number of free parameters? Explain.
  (vi) What kind of changes to this joint distribution and Bayes net could lead to a significant savings?