

CS6601 Final Exam  
Professor Thomas Ploetz  
Fall 2023

1. Before solving the exam, you should read the Ed exam posts for any updated information or clarifications on the exam.
2. You should work on this pdf in parallel with the canvas exam to ensure your work is always backed up.
3. Enter your answers as specified into each questions' provided solution space
4. When you submit the current working copy of your exam download a new copy of the pdf will include any updates or clarifications made to the exam. So submit regularly
5. Make sure your answers follow the instructions given in the problem area.
6. Unless expressly instructed, round all your final answers to 6 decimal places. Do not round intermediate results.
7. If you are asked to round to a different number of decimal places, do it or you will lose points.
8. We do this as many questions are auto grade based on regular expressions which have exponential combinations of solutions. We cannot accommodate all possible combinations. So please help us by closely reading the instructions.
9. Hint: E.g., For rounding to 6 decimals you can use the `round(answer, 6)` function in Python
10. This pdf is for your convenience and may be submitted through Fall 2023 Final Calculations. This pdf is not graded and is not a substitute for the canvas exam. However, any regrades or requests for additional points must be based on submitted calculations, and are at the discretion of the question author.
11. The exam is open-book, open-note, and open video lectures, with no time limit aside from the due date of the exam. No internet use is allowed, except for e-text versions of the textbook, this semester's course materials, Ed, and any links provided in the exam itself.
12. No resources outside this semester's class should be used. Do not discuss the exam with anyone except the instructional staff. That means no posting on Ed, Slack, Discord, Groupme, or any social media platform. Do not post publicly about the exam. If there is a question for the teaching staff, please make it private on Ed and tag it as directed in Ed by question subject and number.
13. Please make separate posts for each question area (e.g., Search, CSP, Patterns over time, Logic and Planning), in general grouping question areas together will have no response from the staff as they focus on responding to their question area

Points breakdown is provided in the canvas exam

### Q 1. Search Part A

In the heart of Silicon Valley, a team of AI researchers at Alphabet-a has just unveiled their latest creation: Gemini, a groundbreaking multimodal language model with the ability to interpret and execute complex tasks. Unlike its predecessors, Gemini is not just a master of words; it's proficient in understanding and generating actionable plans, and has been very popular since its global release.

Gemini has also been integrated into a new line of sophisticated household robots which are in their beta testing phase. In the hands of one eager beta tester, there is one of these robots about to face challenges that will put Gemini's capabilities to the test. Tasked with maintaining domestic tranquility, the robot is ready for its first command. The beta tester, with high expectations, gives the directive: "Make the bathroom sparkle like never before".

In order to fulfill the command, Gemini evaluates several potential actions, some of which are shown in Figure 1. These outputs were sampled via top-k sampling with  $k = 2$ , meaning all but the 2 most probable words are ignored in the search tree. Normalized probabilities are shown. To determine how Gemini will make decisions, Alphabet-a engineers have decided to apply Uniform Cost Search, with a cost function given by:

$$C = -\ln(p)$$

Where  $C$  represents the cost,  $p$  represents the probability of a particular word being chosen, and  $\ln$  is the natural logarithm. For example, if  $p = 0.50$ , then  $C = -\ln(0.50) = 0.693147181$ .

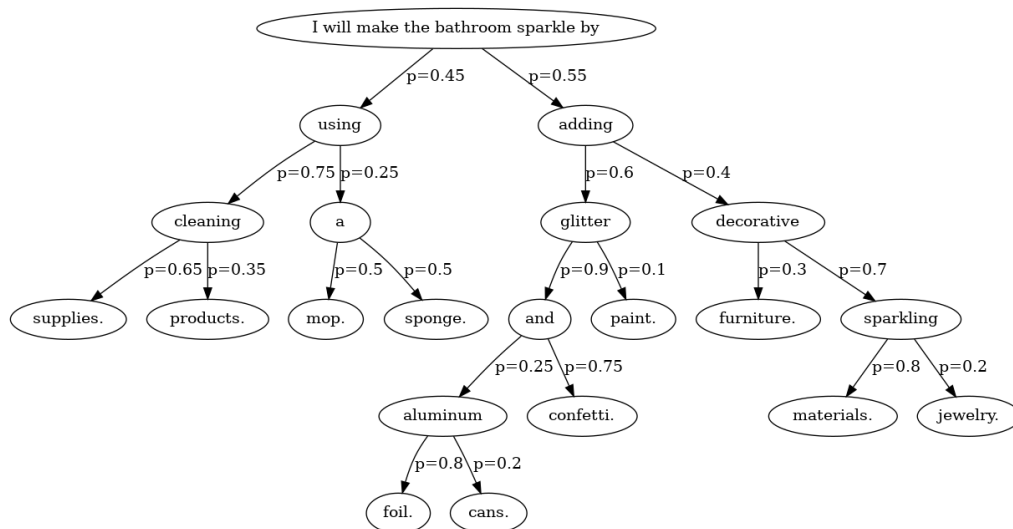


Figure 1: Gemini's search space for Uniform Cost Search.

1. What is the order in which the nodes are explored? (Exclude the root node)

You must indicate the search has ended by entering "search ended". Ties are broken by choosing the leftmost path (1 point)

2. What is Gemini's final path? (Exclude the root node)

You must indicate the path has ended by entering "path ended". Ties are broken by choosing the leftmost path (1 point)

3. What is the total cost of the final path? (1 point)

4. Which of the following describes a reason for using negative log probabilities as costs in the search algorithm? (1 point)

- To prioritize words with higher probabilities by assigning them lower costs.
- To ensure that less probable words are chosen more frequently.
- To simplify the calculation of probabilities into whole numbers.
- To maximize the reward for the most efficient path.

5. Is Uniform Cost Search complete and optimal with respect to the cost function for the tree in Figure 1? (1 point)

- It is both complete and optimal because it always finds the least cost solution
- It is complete but not optimal because it doesn't consider a heuristic
- It is optimal but not complete because it can get stuck in loops
- It is neither complete nor optimal due to the absence of a heuristic

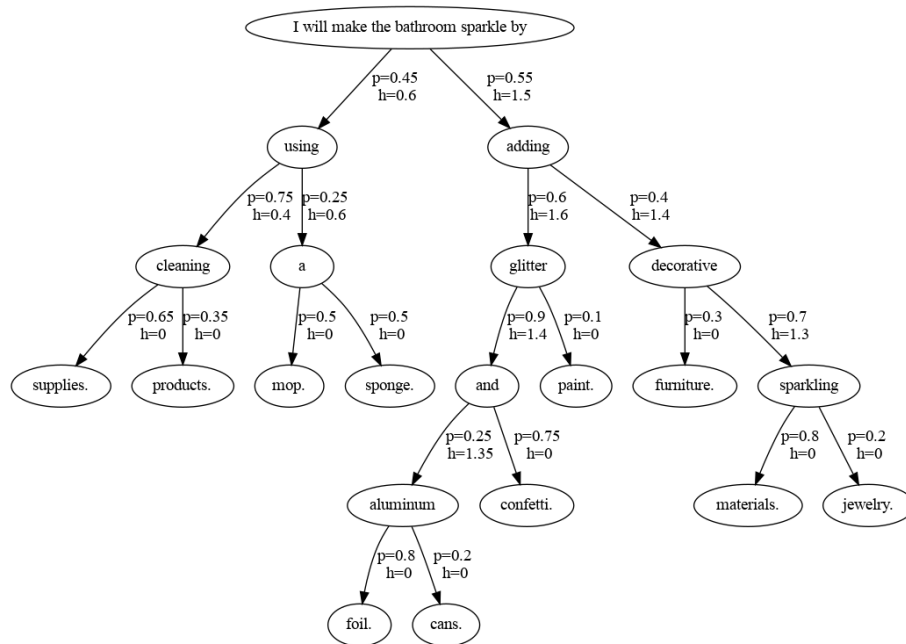


Figure 2: Gemini's search space after applying RLHF.

Q1. Part B

After some initial testing, the AI safety team at Alphabet-a identifies potential misinterpretations by Gemini. To address this, they introduce a heuristic  $h$  derived via Reinforcement Learning from Human Feedback (RLHF), representing the expected human satisfaction with the outcome. This heuristic is designed to guide the search towards solutions that align more closely with human values and expectations, and approaches zero as Gemini gets closer to the user's goal. The engineers decide to update their search strategy to include this heuristic, keeping the same costs from before. They opt for the A\* search algorithm, using the negative log probabilities for the cost as before and their new RLHF heuristic to estimate the remaining cost to reach the goal (Figure 2).

6. What is the order in which the nodes are explored? (Exclude the root node). You must indicate the search has ended by the option "search ended". Ties are broken by choosing the leftmost path (1 point)
  
7. What is Gemini's final path? (Exclude the root node). You must indicate the search has ended by the option "search ended". Ties are broken by choosing the leftmost path (1 point)

8 What is the total cost of the final path? (1 point)

9 The addition of the RLHF heuristic is intended to: (1 point)

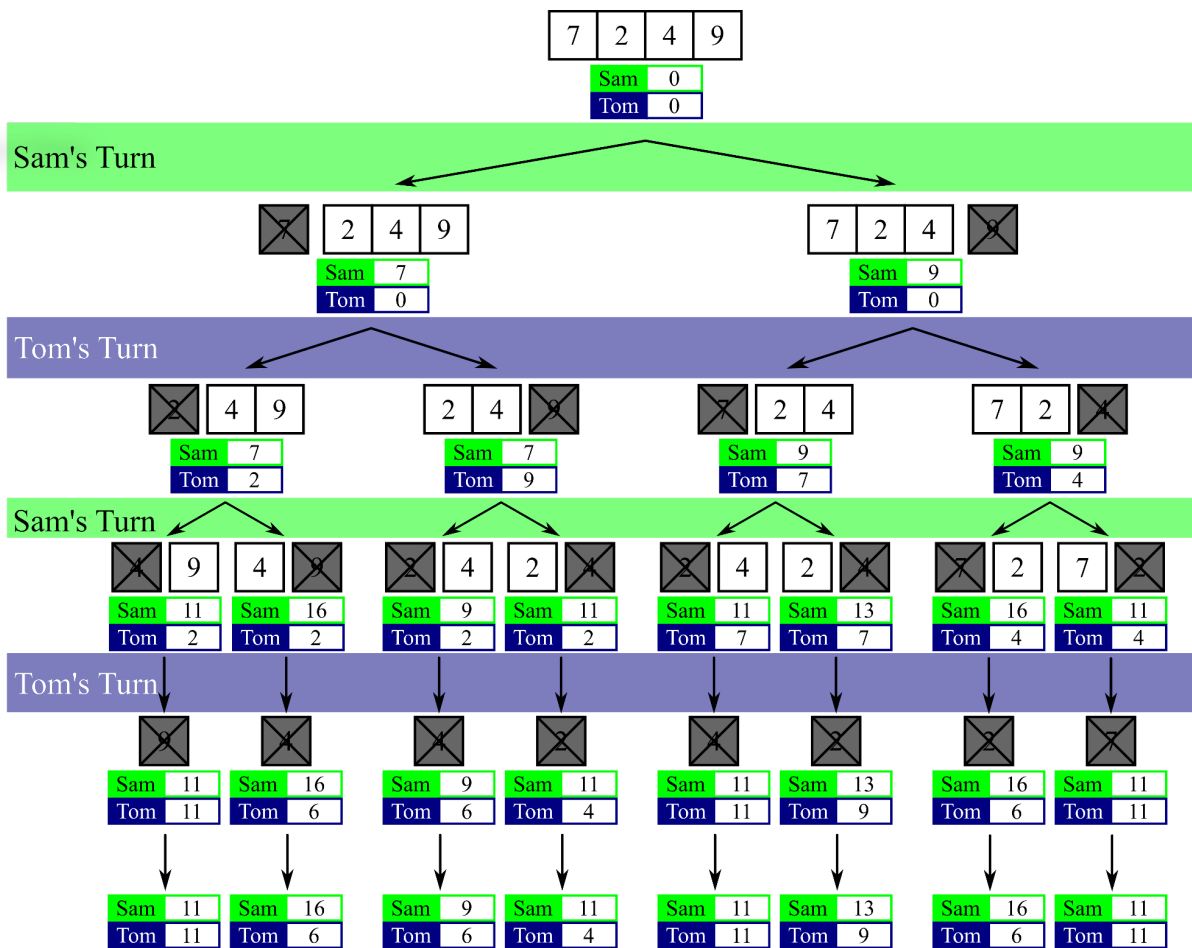
- Predict future states beyond the current search horizon.
- Increase the probability of words being chosen by lowering their costs.
- Align the search more closely with human satisfaction.
- Eliminate the need for negative log probability costs in the search.

10 Which of the following statements is true? (1 point)

- The RLHF heuristic is consistent.
- The RLHF heuristic is admissible.
- The RLHF heuristic is consistent and admissible.
- The RLHF heuristic is neither consistent nor admissible.

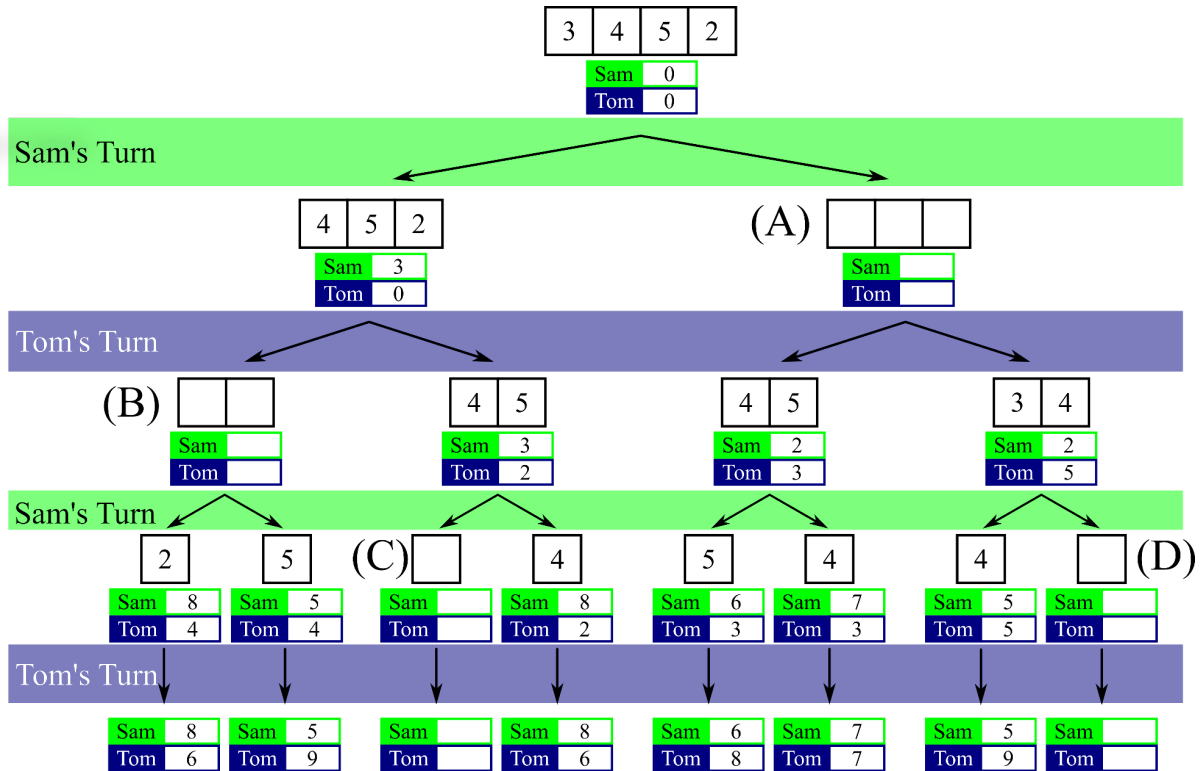
## Q2. Game Playing Part A

Two players are playing a game which starts with an array populated with an even number of elements. Each player takes turns removing an element from the left or right end of the array and adds it to their score. The game starts with both players having no points. The game ends when the array is empty. At the end of the game, the player with the maximum score wins the game. As an example, let's say Sam and Tom are playing a game that starts with the array [7, 2, 4, 9], and let's say that Sam takes the first move. The resulting tree would be:



Note: In this game tree with a branching factor of 2, a choice of left receives the leftmost value from the array, and a choice of right receives the rightmost value from the array. Note that the chosen array elements are restricted and not available for the next ply or layer.

Assume Sam has the first turn and the starting array is [3,4,5,2], to answer the following questions



11 Sam has decided to build a tree to find the optimal way to play and win the game. Can you help Sam complete his tree? Please enter the scores for Sam and Tom at each of the steps queried. Enter as integers (1.5 points)

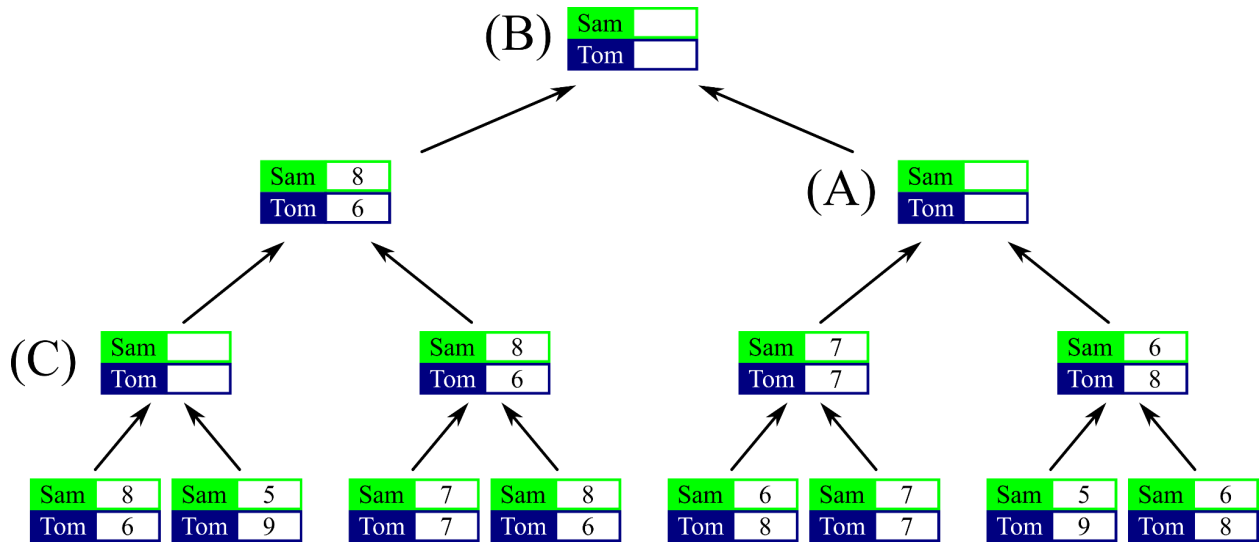
- A. Sam: \_\_\_\_\_, Tom: \_\_\_\_\_
- B. Sam: \_\_\_\_\_, Tom: \_\_\_\_\_
- C. Sam: \_\_\_\_\_, Tom: \_\_\_\_\_
- D. Sam: \_\_\_\_\_, Tom: \_\_\_\_\_

Q2. Game Playing Part B

12 Which of the following is/are true? (1 point)

- Given that both players play optimally and an initial array, the outcome of such a game is always predictable
- If we apply the minimax algorithm with Sam as the maximum player then it will find the shortest path in the above game tree
- The minimax algorithm will provide Sam with a better solution to play the game than alpha-beta pruning for this particular game as it will explore all the nodes in the game tree
- Both alpha-beta pruning and minimax algorithm will provide Sam with the same solution to play the games

Let us say that we are rooting for Sam, and thus we hope she wins. By strictly looking at the scores and the minimax tree, we obtain the following tree:



13 Complete the missing values in the minimax tree. (1.5 points)

A. Sam: \_\_\_\_\_, Tom: \_\_\_\_\_

B. Sam: \_\_\_\_\_, Tom: \_\_\_\_\_

C. Sam: \_\_\_\_\_, Tom: \_\_\_\_\_

14 If we apply alpha-beta pruning to the above tree with Sam as the maximum player, how many leaf nodes are pruned?

(1 point)

15 If both players play optimally, who will win this game? (1 point)

- Sam
- Tom



### Q3 Optimization

A research grant has been funded by the US department of transportation to optimize a high-speed train infrastructure between 6 key cities in the Southeast. You are tasked with creating a Genetic Algorithm to determine the best acyclic route. Ignore any infrastructure currently in place between these cities. The cities are:

Atlanta (ATL)  
Savannah (SAV)  
Nashville (NAS)  
Charlotte (CLT)  
Raleigh (RAL)  
Birmingham (BHM)

To do this, we will randomly generate a population of paths, to be parents. Each parent will be denoted as a chromosome, composed of a set of genes, where each gene is named with a city abbreviation. The cities are stitched together to create a path that the train can take. The parent population will be:

Parent 1: NAS → ATL → CLT → RAL → SAV → BHM  
Parent 2: CLT → NAS → ATL → SAV → BHM → RAL  
Parent 3: CLT → BHM → ATL → SAV → RAL → NAS  
Parent 4: RAL → NAS → SAV → BHM → ATL → CLT

We then evaluate their fitness with a custom fitness function. For a project such as this, the fitness of a path would be affected by length, geography, building costs, zoning, noise ordinances, and legal costs. In this problem, we will only use length, meaning the shorter the length the more fit the candidate is. To calculate the fitness we will use this function, which is the inverse of the length:

$$f(\text{fitness}) = \text{length}^{-1}$$

Now let's evaluate each candidate based on this function. Here are the distances:

ATL ↔ SAV: 250 miles  
ATL ↔ NAS: 250 miles  
ATL ↔ CLT: 245 miles  
ATL ↔ RAL: 400 miles  
ATL ↔ BHM: 150 miles  
SAV ↔ CLT: 250 miles  
NAS ↔ CLT: 400 miles  
NAS ↔ BHM: 185 miles  
CLT ↔ RAL: 170 miles  
RAL ↔ SAV: 330 miles  
BHM ↔ SAV: 350 miles  
RAL ↔ BHM: 500 miles  
CLT ↔ BHM: 350 miles  
RAL ↔ NAS: 450 miles

NOTE: If the distance is not listed here, then the path between the cities is not possible for legal or geographical reasons.

### Q3.1 Optimization

Provide the fitness for each parent: (.75 points ea.)

16: Fitness of Parent 1: \_\_\_\_\_

17: Fitness of Parent 2: \_\_\_\_\_

18: Fitness of Parent 3: \_\_\_\_\_

19: Fitness of Parent 4: \_\_\_\_\_

Now we will perform crossovers[1] from two pairs of existing parents to create a new generation of child paths. We will then check to see if each new child path is a viable child. Our criteria for viability is that the train can reach each city in the path exactly once. We will also check the fitness of each viable child. We have randomly selected a crossover point at index 3 (zero-based indexing) for all the children. Parents 1 & 2 will partner to create Child 1 (P1|P2) and Child 2 (P2|P1), and Parents 3 & 4 will partner to create Child 3 (P3|P4) and Child 4 (P4|P3). What is the fitness of each child? Please put -1 if the child is not viable, because a path does not exist or the sequence does not go to every city exactly once:

1. See "Evolutionary algorithms", AIMA 4th Ed., Russell & Norvig, Chptr. 4.1.4, Figure 4.6 (4 Queens ex.)

(.75 points ea)

20: Fitness of Child 1: \_\_\_\_\_

21: Fitness of Child 2: \_\_\_\_\_

22: Fitness of Child 3: \_\_\_\_\_

23: Fitness of Child 4: \_\_\_\_\_

### Q3.3 Optimization

We will now introduce mutations to the children and recheck viability and fitness. The possible mutations are:

Swap mutation: Two genes are swapped at random

Scramble mutation: A subset of genes in the chromosome are scrambled to a random order

Inverse mutation: The order of a random subset of genes in the chromosome is inverted

We have chosen to do a swap mutation next, at indices 2 and 5 (again, zero-based indexing). Swap the genes at these indices and reevaluate the fitness of each child. Please put -1 if the child is not viable.

(.75 points ea)

24: New fitness of Child 1: \_\_\_\_\_

25: New fitness of Child 2: \_\_\_\_\_

26: New fitness of Child 3: \_\_\_\_\_

27: New fitness of Child 4: \_\_\_\_\_

28. Normally we would do these steps for a few generations before selecting the most fit path. But for this question, we are now done. Please note which path was the most fit so far in our search, by ordering each city in terms of what stop it would be from 1-6: (1 point)

Atlanta \_\_\_\_\_

Birmingham \_\_\_\_\_

Charlotte \_\_\_\_\_

Nashville \_\_\_\_\_

Raleigh \_\_\_\_\_

Savannah \_\_\_\_\_

#### Q4 Constraint Satisfaction Problems (CSPs)

Your friends have been training and preparing for **SailGP** all summer and the time has finally come for registration. However they are squabbling over what position each person should be in charge of and so they've turned to you to help them resolve the issue! Recognizing that this is a CSP problem, you've decided to take an algorithmic approach to assign everyone roles.

**Variables:** There are 6 variables, and they are your friends, Alice, Bob, Carol, Dan, Elise, and Fred.

**Domains:** There are 4 values that each variable can be assigned, and they are the positions that one can take on the boat, Driver (D), Flight Controller (FC), Wing Trimmer (WT), and Grinder (G).

It is important to note that there have to be 6 people on the boat at all times, and that each boat may only have exactly 1 Driver, exactly 1 Flight Controller, and exactly 1 Wing Trimmer (there can be more than 1 Grinder). We'll try to capture these requirements as constraints, but you'll see that we won't be able to capture them entirely.

Below in Figure 1. is the starting set of domains for the CSP problem before any domain reduction occurs. We will refer to it as domain state zero.

<b>Alice:</b> {D, FC, WT, G}	<b>Carol:</b> {D, FC, WT, G}	<b>Elise:</b> {D, FC, WT, G}
<b>Bob:</b> {D, FC, WT, G}	<b>Dan:</b> {D, FC, WT, G}	<b>Fred:</b> {D, FC, WT, G}

Figure 1: Domain State Zero

**Constraints:** You've talked to each of your friends individually and collected all the different wishes and grievances they have and reduced them to a set of unary and binary constraints that are easy to implement in the CSP algorithm. They are as follows:

##### Unary Constraints

1. Bob cannot be the FC or a G.
2. Carol cannot be the FC.

## Binary Constraints

1. If a person is assigned D, no other person may be assigned D.
2. If a person is assigned FC, no other person may be assigned FC.
3. If a person is assigned WT, no other person may be assigned WT.
4. If Alice is the FC, neither Bob, Dan, nor Elise can be the WT.
5. If Alice is the FC, neither Bob nor Fred can be the D.
6. If Carol is a G, neither Bob nor Dan can be the G.
7. Elise can be the D or WT if and only if Carol is the FC.
8. Dan or Fred can be the FC if and only if Elise is either the D or the WT.

Constraint Graph: Notice how unary constraints are not reflected in the constraint graph shown in Figure 2. and that all binary constraints between two variables X and Y are expressed in the form of one edge between X and Y. Also note that all edges that only contain the binary constraints 1, 2, or 3 from above are represented using a dashed line and that all edges that contain at least one of the binary constraints 4, 5, 6, 7, 8 are represented using a solid line.




Figure 2: The Constraint Graph

Constraint Table: We also provide you with a nifty table with all of the allowed assignment tuples represented by each edge in the constraint graph (what the algorithm will check when evaluating arc consistency/binary constraints).

Notation: The symbol \* means “any”. The symbol  $\sim$  means “not”.

Edge	Total	Allowed Tuples
(Alice, Bob)	11	(FC, G); (D, $\sim$ D); (WT, $\sim$ WT); (G, *)
(Alice, Carol)	13	(FC, $\sim$ FC); (D, $\sim$ D); (WT, $\sim$ WT); (G, *)
(Alice, Dan)	12	(FC, D); (FC, G); (D, $\sim$ D); (WT, $\sim$ WT); (G, *)
(Alice, Elise)	12	(FC, D); (FC, G); (D, $\sim$ D); (WT, $\sim$ WT); (G, *)
(Alice, Fred)	12	(FC, WT); (FC, G); (D, $\sim$ D); (WT, $\sim$ WT); (G, *)
(Bob, Carol)	12	(FC, $\sim$ FC); (D, $\sim$ D); (WT, $\sim$ WT); (G, $\sim$ G)
(Bob, Dan)	13	(FC, $\sim$ FC); (D, $\sim$ D); (WT, $\sim$ WT); (G, *)
(Bob, Elise)	13	(FC, $\sim$ FC); (D, $\sim$ D); (WT, $\sim$ WT); (G, *)
(Bob, Fred)	13	(FC, $\sim$ FC); (D, $\sim$ D); (WT, $\sim$ WT); (G, *)
(Carol, Dan)	12	(FC, $\sim$ FC); (D, $\sim$ D); (WT, $\sim$ WT); (G, $\sim$ G)
(Carol, Elise)	9	(FC, $\sim$ FC); (D, FC); (D, G); (WT, FC); (WT, G); (G, FC); (G, G)(FC, $\sim$ FC); (D, $\sim$ D); (WT, $\sim$ WT); (G, *)
(Carol, Fred)	13	$\sim$ D); (WT, $\sim$ WT); (G, *)
(Dan, Elise)	12	(FC, D); (FC, WT); (D, $\sim$ D); (WT, $\sim$ WT); (G, *)
(Dan, Fred)	13	(FC, $\sim$ FC); (D, $\sim$ D); (WT, $\sim$ WT); (G, *)
(Elise, Fred)	12	(FC, $\sim$ FC); (D, $\sim$ D); (WT, $\sim$ WT); (G, $\sim$ FC)

Table 1: The Constraint Table

Note\*: This table only applies to binary constraints and does not account for the unary constraints.

29. (2 points)

How many complete assignments are there in total for this CSP problem? For an assignment to be complete, all variables must be assigned a single value and the assignment does not have to be consistent.

30. (2 points)

While using a backtracking search based algorithm, heuristics may be used at each step to choose what variable to assign next and what value to assign it. One such heuristic is the degree heuristic where we select the variable that is involved in the largest number of constraints on other unassigned variables (binary constraints) [1]. Beginning with domain state zero provided above, we will first perform the following steps:

Reduce all domains by applying node consistency. (3 values should get removed across all domains)

Assign the value of FC to the variable Alice. (All remaining values in Alice's domain should get removed and Alice's domain is now fixed)

Apply arc-consistency on Alice only. (8 values should get removed across all domains)

After applying the above steps, we decide that we want to continue with the algorithm. If we use the degree heuristic based only on the number of solid edges in the constraint graph provided above to select the next variable for assignment, which variable(s) may we select for assignment next? (select all that applies)

Hint\*: The node Fred is considered to have a node degree of 2.

1. AIMA, 4th Ed., Russell & Norvig, Chptr.

6.3.1

- Alice
- Bob
- Carol
- Dan
- Elise
- Fred

31. (2 points)

An alternative heuristic that we can use to pick a variable to assign a value to is the minimum remaining value heuristic (MRV). MRV chooses the variable with the fewest “legal” values (values available in its domain) [1]. Beginning with domain state zero provided above, we will first perform the following steps:

Reduce all domains by applying node consistency. (3 values should get removed across all domains)

Assign the value of G to the variable Carol. (All remaining values in Carol’s domain should get removed and Carol’s domain is now fixed)

Make the new state of domains arc-consistent. (5 values should get removed across all domains)

After applying the above steps, we decide that we want to continue with the algorithm. If we use the minimum remaining value heuristic, which variable(s) may we select for assignment next?

1. AIMA, 4th Ed., Russell & Norvig, Chptr. 6.3.1

- Alice
- Bob
- Carol
- Dan
- Elise
- Fred

Alice: {D, FC, WT, G}	Carol: {D, FC, WT, G}	Elise: {D, FC, WT, G}
Bob: {D, FC, WT}	Dan: {D, FC, WT}	Fred: {D, FC, WT, G}

Figure 3: Domain State One

Hint: Note that we are applying LCV to the set of domains without applying node consistency first. Also note that G is missing from some of the domains.

32. (2 points)

The two heuristics mentioned above deal with picking variables to make assignments to at each step. But what about choosing the value to assign to the variable after the variable to be changed has been chosen? A heuristic that comes in handy for selecting a value is the least constraining value heuristic (LCV). LCV chooses the value that rules out the fewest choices for the neighboring variables in the constraint graph[1]. Beginning from the set of domains provided below (domain state one), we choose the variable Dan to assign a value to. If we use the least constraining value heuristic, which value(s) may we assign to variable Dan next?

1. AIMA, 4th Ed., Russell & Norvig, Chptr. 6.3.1

Dan:

- G
- FC
- D
- WT

Q5 Probability

34. Picking Out of a Bag (1 point)

A bag contains 30 pieces of paper marked from 1 to 30. Find the probability of drawing one odd and one even in a single draw of two tickets. \_\_\_\_\_

35. Picking Out of a Bag Part II (1 point)

Bag 1 has four dimes and two nickels. Bag 2 has three dimes and three nickels. A bag is selected at random and a dime is drawn. What is the probability the dime came from bag 2? \_\_\_\_\_

36. Logical Proofs (1 point)

Mark all that are True

- If  $P(A|B) = P(A)$  then  $P(B|A) = P(B)$
- If  $P(A, B|C) = P(C|A, B)$  then  $P(A|C) = P(B|C)$
- If  $P(B|C) = P(B)$  then  $P(A|B, C) = P(A|B) * P(A|C)/P(A)$

37. Entering a Raffle (1 point)

There is a raffle with prizes. To enter the raffle you read and agree to the following rules:

1. Buy a ticket for \$2
2. Each ticket is numbered from 000 - 999, and is assigned at random, and is unique
3. After all of the tickets are sold a winning number from 0 to 999 is selected at random
4. If a ticket contains any of the numbers in the winning numbers in the same ordered sequence (ones, tens, hundreds) as the winning number a monetary prize is awarded
5. The possible prizes are shown below
  - a. All 3 numbers the same: \$1200
  - b. Two numbers the same: \$100
  - c. One number the same: \$10

For example, if the winning number was 692 and you got ticket with the number 632 you would get a prize of \$100 for having two of the numbers in the same ordered sequence as the winning number

Note: You can only win the highest possible prize that applies to your ticket. If you get 2 numbers correct you only receive \$100 and not any of the \$10 prize for getting 1 number correct.

What is the expected profit (or loss) you would expect from a purchase of 3 tickets? If a loss please enter a negative number. (2 points)

37.) \_\_\_\_\_





Q6. Bayes Net

Sometimes, the outputs from Bayes Networks can seem counter-intuitive. To understand this, consider the following example: Earthquakes and burglaries are independent events. Either of them happens with a probability of 0.01, as shown in Table 1 and 2. Either can possibly cause an alarm to go off. The probability graph is shown in Figure 1. The joint probability distribution should be calculated as:

def

$$\Pr(B = b, E = e, A = a) = p(b) \cdot p(e) \cdot p(a|b, e)$$

e	p(e)
1	0.01
0	0.99

Table 1: The probability of an earthquake happened. The value of 1 represents the occurrence of an event, and the value of 0 represents the non-occurrence.

b	p(b)
1	0.01
0	0.99

Table 2: The probability of a burglary happened. The value of 1 represents the occurrence of an event, and the value of 0 represents the non-occurrence.

The conditional probability of  $P(a | b, e)$  is provided in the Table 3. For example,  $P(a = 1 | b = 0, e = 1) = 1.0000$  denotes 'the probability that if there is an earthquake and no burglary, the alarm will go off'.




Figure 1: The probability graph of an alarm (A) goes off.

$b$	$e$	$a$	$P(a   b, e)$
0	0	0	1.0000
0	0	1	0.0000
0	1	0	0.0000
0	1	1	1.0000
1	0	0	0.0000
1	0	1	1.0000
1	1	0	0.0000
1	1	1	1.0000

Table 3: Conditional probability table for  $P(a | b, e)$

38. Given all the information above, please fill in the joint probability distribution table. Round your answers to 4 decimal places (2.5 points)

$b$	$e$	$a$	$P(A = a, B = b, E = e)$
0	0	0	
0	0	1	
0	1	0	
0	1	1	
1	0	0	
1	0	1	
1	1	0	
1	1	1	

39. (2 points)

Given that the alarm has sounded, what is the probability of a burglary having occurred, without knowing about an earthquake?

Probability: \_\_\_\_\_

40. (2 points)

Calculate the probability  $P(B = 1 | A = 1, E = 1)$ . Round your answer to 2 decimal places.

Probability: \_\_\_\_\_

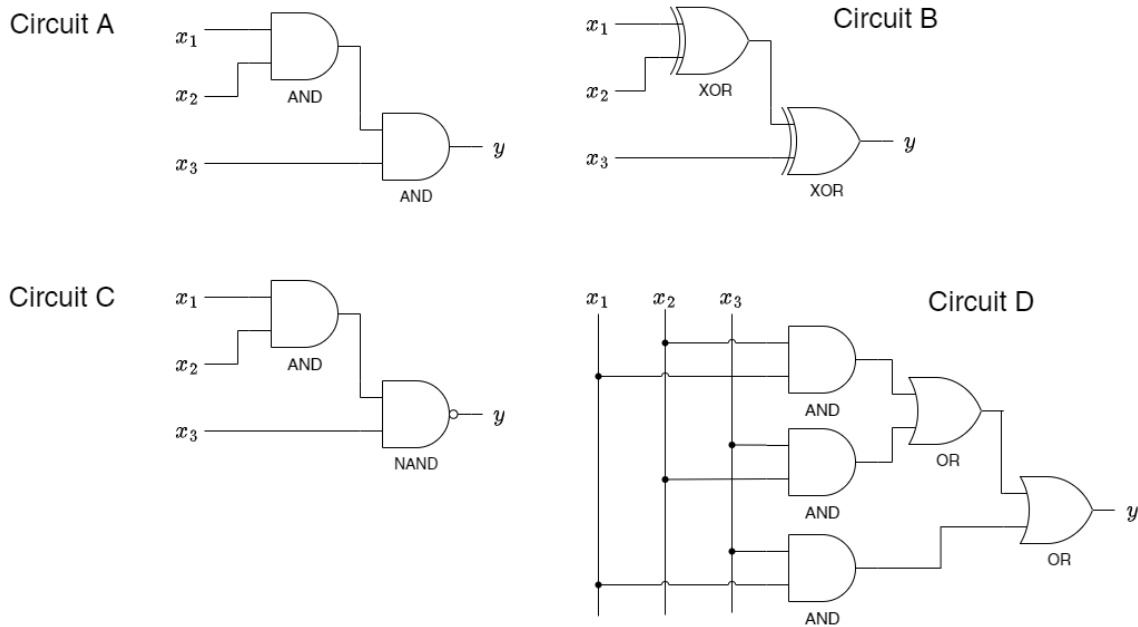
41. Does hearing that there's an earthquake increase, decrease, or keep constant the probability of a burglary? (1.5 points)

- Increase
- Decrease
- Keep constant

Q7. Machine Learning

Model complexity

In the realm of machine learning, the choice of an appropriate level of model complexity is crucial for achieving the best performance on a task. While a complex model is prone to overfitting, an overly simplistic model may not have enough representational power to capture the relationships between the input variables and desired outputs. Let us study this concept in the context of binary classification tasks based on the logical circuits shown below.



Each circuit has three binary inputs  $x_1$ ,  $x_2$  and  $x_3$ , which can be 0 or 1. Let us train one of the simplest models, a linear classifier, to predict the outputs of logical circuits (separately). The linear classifier is given by,

$$F(x_1, x_2, x_3) = w_1x_1 + w_2x_2 + w_3x_3 + b$$

Where,  $w_1$ ,  $w_2$  and  $w_3$  are weights and  $b$  is the bias.  $w_1$ ,  $w_2$ ,  $w_3$  and  $b$  are real numbers. The model predicts 1 if  $F(x_1, x_2, x_3) \geq 0$  and 0 otherwise.

42. On which of the provided logic circuits, is it possible for the linear classifier F to attain 100% test accuracy? (Hint: you may find building truth tables and the notion of separability useful.) (1 point)

- o Circuit A
- o Circuit B
- o Circuit C
- o Circuit D

## 7.2 Machine Learning

### Decision Trees

You are an ML engineer at a major music streaming service. To analyze trends in music and improving the quality of music recommendations, you collected a dataset analyzing musical characteristics of 10 songs, shown in Table 1. The features you focused on were, Tempo, which is measured in beats per minute; Instrumentalness (0 to 1), which measures how instrumental (1) or vocal (0) the song content is; Energy (0 to 1), which is the average energy level of the song based on its frequency spectrum and Danceability (0 to 1), which measure of favorability of a song for dancing.

Song	Tempo	Instrumentality	Energy	Danceability	Genre
1	120	0.05	0.85	0.75	Rock
2	140	0.10	0.80	0.70	Pop
3	100	0.80	0.40	0.60	Jazz
4	90	0.95	0.30	0.50	Jazz
5	160	0.20	0.90	0.80	Rock
6	110	0.70	0.50	0.65	Jazz
7	130	0.15	0.75	0.70	Pop
8	105	0.90	0.35	0.55	Jazz
9	150	0.30	0.95	0.85	Rock
10	125	0.60	0.55	0.65	Pop

Table 1: Song dataset

What is the Gini gain corresponding to the following features in the dataset, assuming that the split is performed at the mean of the feature values? (Report 6 digits) (1.5 points ea)

43: Tempo: \_\_\_\_\_

44: Instrumentality: \_\_\_\_\_

## 7.3 Machine Learning

### KNNs and Feature pre-processing

J3PO, an upcoming artist, just released a song on your streaming platform. You wanted to add it to a genre relevant playlist. However, it is not possible for you to manually listen to the song and determine the genre, as you have thousands of other songs to classify. Having written code to analyze musical features, you determine that J3PO's song has a tempo of 125 bpm, instrumentality of 0.85, energy of 0.4 and danceability of 0.5.

45: Determine the 5 most similar songs to J3PO's song based on the euclidean distance metric. The euclidean distance metric is given as,

$$d(\mathbf{p}, \mathbf{q}) = \sqrt{\sum_{i=0}^n (p_i - q_i)^2}$$

Where,  $p, q$  are two vectors in  $n$ -dimensional space. (2 points)

- Song 1       Song 2       Song 3       Song 4       Song 5  
 Song 6       Song 7       Song 8       Song 9       Song 10

46: What is genre of J3PO's song based on  $k$ -nearest neighbors, with  $k = 5$ . (Choose the majority genre)(2 points)

Genre: \_\_\_\_\_

47: It was lunch break and you decide to relax to J3PO's song. Something definitely felt wrong. You found out that your  $k$ NN model had wrongly predicted its genre! You rush to your computer terminal to investigate what's wrong. Turns out, using raw non-normalized tempo values had biased your classifier significantly. Apply min-max normalization (between 0 and 1) to the tempo feature and determine the 5 most similar songs to J3PO's song. (Use the euclidean distance metric) (2 points)

- Song 1       Song 2       Song 3       Song 4       Song 5  
 Song 6       Song 7       Song 8       Song 9       Song 10

48: Now that tempo values are min-max normalized, what is genre of J3PO's song based on  $k$ -nearest neighbors with  $k = 5$ . (Choose the majority genre) (2 points)

Genre: \_\_\_\_\_

## 7.4 Machine Learning

### Neural nets



49: Which of the following statements are true? (2 points)

- Gradient descent is guaranteed to converge to the global minimum on a convex loss surface.
- Precision is a more suitable metric than recall when the cost of false positives is significantly higher than the cost of false negatives.
- Choosing  $k = 1$  for a KNN model is guaranteed to cause underfitting.
- A multi-layer perceptron with an activation function  $f(x) = 1.5x$  can learn any decision boundary, given an arbitrary number of layers. Assume that the same activation function is used in every layer.

Q8 Patterns Recognition Through Time: Dynamic Time Warping (DTW)

Dynamic Time Warping (DTW) - Tutorial:

Dynamic Time Warping (DTW) stands as a robust algorithm utilized to assess the similarity between two time-series sequences. Its applications span a wide range of domains, encompassing shape recognition, handwriting recognition, and speech recognition. This algorithm offers a higher degree of flexibility when compared to the Euclidean distance metric, making it particularly valuable for data with significant variance. To illustrate the distinction, consider a visual representation of the way these two algorithms compare two similar time series by evaluating individual data points. The Dynamic Time Warping algorithm aligns data points in a manner that optimizes their comparability, in contrast to the somewhat more simplistic approach of the Euclidean matching algorithm.



We can perform DTW with two signals and a grid:

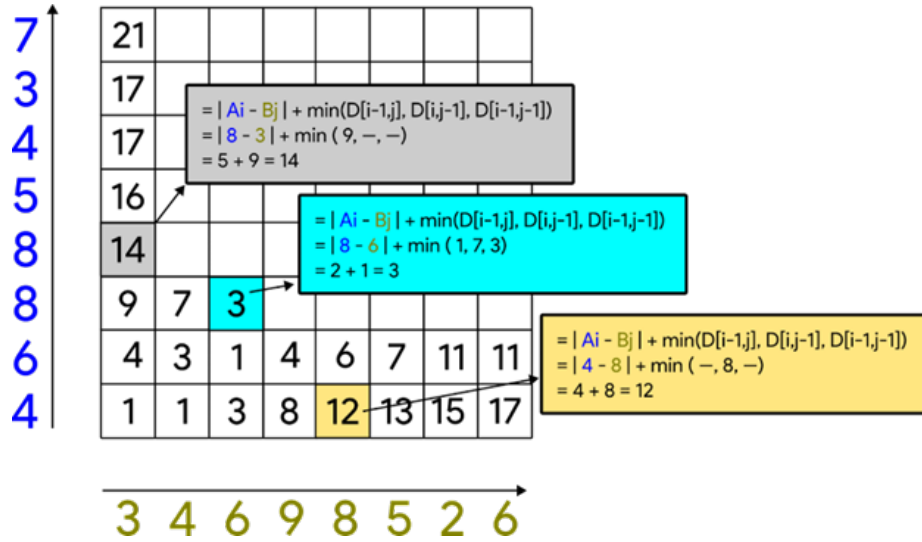
$$A = [4,6,8,8,5,4,3,7] \quad B = [3,4,6,9,8,5,2,6]$$

where each square is computed with the following formulas:

$$D_{i,j} = |A_i - B_j| + \min(D_{i-1,j}, D_{i,j-1}, D_{i-1,j-1}) \tag{1}$$

resulting in the following.






**Grid Filling:** The task begins with filling a grid, starting from the bottom left corner and progressing upward. The filling is done using a specific formula as in (1), which calculates values for each cell in the grid. This formula in (1) takes into account various factors, and the grid is constructed from bottom to top as shown in the figure above.

**Path Finding:** Once the grid is filled, the objective shifts to finding the lowest-cost path within this grid. The path is portrayed as it begins at the bottom left corner and concludes at the top right corner but to identify the lowest-cost path, a backtracking process is employed, moving from the top to the bottom as shown in the figure below..

**Movement Choices:** At each step during the path-finding process, you have three choices for movement: diagonal, left, or down. These choices dictate how you traverse from one grid cell to the next.

**Handling Equal Values:** It's possible to encounter multiple paths with the same cost values. In such situations, prioritization is given to diagonal move-ments. This means that if there are ties in cost, the diagonal path is preferred.

**Objective:** The primary goal of this entire process is to pinpoint the path that incurs the least overall cost while navigating from the starting point (bot-tom left corner) to the destination (top right corner) within the grid. This cost-minimization approach ensures an efficient and cost-effective route, but the path is obtained by backtracking.



Our path from our example is 1-1-1-2-2-2-3-4-5, going from the start (bottom left) to the end (top right).

Dr. Emily is developing a music recommendation system and employs Dynamic Time Warping (DTW) to align and compare three music sequences, SongA, SongB, and Song C. Below are the three different music sequences, each represented as a series of pitch values:

SongA = [59, 64, 65, 67, 69, 71, 72, 71, 76, 69, 71]

SongB = [60, 62, 63, 65, 67, 69, 71, 69, 74, 65, 73]

SongC = [66, 65, 64, 69, 71, 75, 74, 76, 78, 79, 81]

These three songs have different sequences of pitch values, allowing for variations in melody and musical structure. When using Dynamic Time Warping (DTW) or other similarity metrics, you can analyze how these songs compare in terms of musical similarities and differences. The DTW distance between these sequences will reflect their musical relationships, with closer distances indicating greater similarity and more distant distances suggesting differences in the melodic content. The time series data, in this case, reflects the pitch progression over time in the respective songs, allowing for comparative analysis using Dynamic Time Warping (DTW) or other similarity metrics.

Q8.A

Calculating the Dynamic Time Warping (DTW) distance between the Songs is essential because it provides a measure of their musical similarity while accounting for timing variations and structural nuances. The DTW path, which represents the optimal alignment between the two songs, reveals where and how the sequences match, helping identify shared melodic patterns. The DTW distance quantifies the overall dissimilarity, with a lower distance indicating a stronger musical resemblance. In essence, DTW unveils the hidden musical connections and differences between the songs, enhancing our ability to assess their similarity in a more musically meaningful way than simple Euclidean distance.

Perform DTW on the signals obtained by Dr. Emily. Enter the path between the three pairs. Make sure you get full credit by writing your path with values separated by hyphens (e.g., 1-2-3-4-5-6). Grids are provided for your convenience and will not be graded.

SongA, SongB (2 points)

50: -----

73		45		27		17		16		17	
65											19
74											
69											12
71											
69											20
67											
65											50
63											
62											76
60											
	59	64	65	67	69	71	72	71	76	69	71

Note: The above grid is only for your calculation. It is not needed for Grading.

SongA, SongC (2 points)

51: -----

81		99		76		48		43		32	
79											32
78											
76											23
74											
75											23
71											
69											27
64											
65											49
66											
	59	64	65	67	69	71	72	71	76	69	71

Note: The above grid is only for your calculation.  
It is not needed for Grading.

SongB, SongC (2 points)

52: -----


81		118		93		64		52		43	
79											37
78											
76											28
74											
75											26
71											
69											28
64											
65											42
66											
	60	62	63	65	67	69	71	69	74	65	73

Note: The above grid is only for your calculation. It is not needed for Grading.

53. Which song is the most likely outlier or anomaly in the set-in terms of musical similarity to the other songs? (1 point)

- SongA
- SongB
- SongC
- They are all similar
- They are all too far apart to be similar





Visualization of Sakoe-Chiba bands

In analyzing her songs SongD and SongE, Emily encounters time and space constraints. To efficiently measure their similarity while respecting these limitations, she employs Sakoe-Chiba bands within the Dynamic Time Warping (DTW) framework. This approach enables a meaningful comparison, capturing the musical essence of both compositions despite variations in timing and tempo.

SongD = [45, 48, 50, 52, 55, 57, 59, 60, 62, 64, 67]  
 SongE = [42, 45, 47, 49, 52, 54, 56, 57, 59, 61, 64]

Based on the understanding from Sakoe-Chiba Bands find the path for the following cases.

**54:** (2.5 points)  
 Consider our Sakoe-Chiba Band has a bound of  $N=0$ . What is the path between SongD and SongE?  
 Make sure you get full credit by writing your path with values separated by hyphens (e.g., 1-2-3-4-5-6). Grids will not be graded.

Path: -----

**55:** (2.5 points)  
 Consider our Sakoe-Chiba Band has a bound of  $N=2$ . What is the path between SongD and SongE?  
 Make sure you get full credit by writing your path with values separated by hyphens (e.g., 1-2-3-4-5-6). Grids will not be graded.

Path: -----

Note: The Grid depiction is not given, you have to figure out the Grid based on the concepts explained above.



## Q9 Logic and Planning

The rise of smart homes has led to an increased demand for autonomous household robots that can assist customer's daily tasks. Your team at the Innovative Robotics Company (IRC) is in the process of designing a new household robot named "RoboHelper."

### 9.1 Task 1: Decide Your Robot

As a technological lead in the AI team at IRC, you are assigned the task of creating a data-driven solution for decision-making. There are several potential features that could be considered. However, the team must prioritize for which features to develop for the first release due to budget and time constraints. The following features have been identified as critical to the success of the robot's development:

- is Energy Efficient (EE): The robot's capability in optimizing power usage during tasks, including intelligent charging and power management tailored to user routines
- has Emotional Intelligence (EI): The robot's ability to recognize and respond to human emotions, providing more personalized services.
- has Health Monitoring (HM): The integration of sensors to track health-related variables such as air quality and remind users of medication schedules.
- has User Preferences Learning (UPL): The capability to learn and adapt to individual customer preferences.

The success of RoboHelper (isSuccessful) lies on the logic circuit depicted in 1. Your task is to analyze this circuit and construct a corresponding truth table.

Given the market study collected by IRC, you decide which robot functionalities IRC should build in the future. The decision must be made based on historical data, which is encoded by the generated truth table. Note that T = True and F = False.

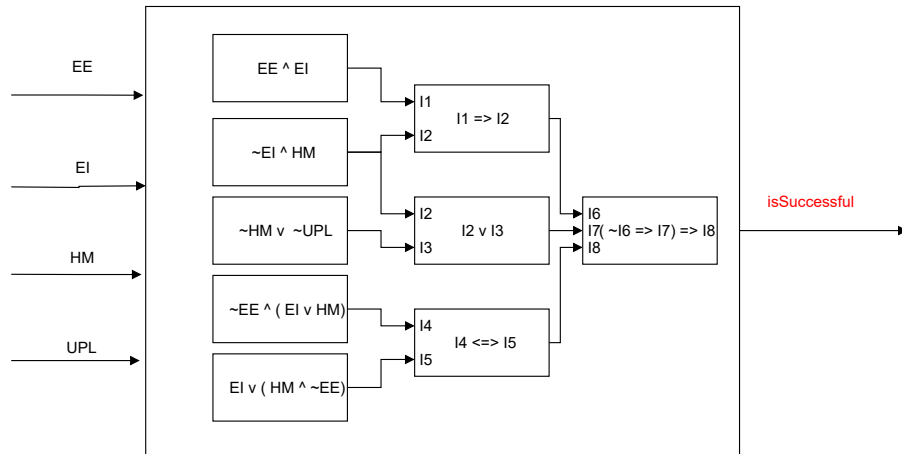


Figure 1: Logic circuits whether or not the robot is successful based on the four factors

- 56 Find the outputs for all combinations of inputs in the truth table, i.e. type T or F. (4 points)

EE	EI	HM	UPL	isSuccessful
F	F	F	F	
F	F	F	T	
F	F	T	F	
F	F	T	T	
F	T	F	F	
F	T	F	T	
F	T	T	F	
F	T	T	T	
T	F	F	F	
T	F	F	T	
T	F	T	F	
T	F	T	T	
T	T	F	F	
T	T	F	T	
T	T	T	F	
T	T	T	T	

Your manager leaves comments on the future robot platform of IRC to you. The executives are debating between two types of potential robots

1. Drone: Compact and agile drone is energy efficient, good at learning user preferences, and can understand emotional subtlety. The drone is perfect for personalized aerial tasks within the home, but it does not offer health monitoring.
2. Wheeled robot: Robust and reliable wheeled robot is energy efficient, has a function of health monitoring, and can user preference. The wheeled robot is ideal for daily ground-based assistance in a health-focused household. However, it lacks emotional intelligence.
3. Humanoid robot: Advanced and interactive humanoid robot has a full suite of features: energy efficiency, emotional intelligence, health monitoring, and user preference learning. The humanoid robot is a versatile all-rounder for comprehensive household tasks.

- 57 Based on the truth table and considering your manager's insights on potential robot platforms, determine which robot type, Drone, Wheeled robot, or Humanoid robot, IRC should focus on developing first. Select all it's applicable. (3 points)

- Drone
- Wheeled robot
- Humanoid robot

## 9 Logic and planning Part B

### Planning for Robot

IRC has decided to attach a vacuum function for the household robot. Your team has been assigned a new AI task for a smart vacuum-cleaning robot. IRC assumes that the customer's home has three rooms: a Living Room (LR), a Kitchen (K), and a Bedroom (BR). IRC also assumes that each room can be either clean (C) or dirty (D). The robot can move between the rooms, clean the rooms, and have a battery level indicator that can be low (L) or full (F). The robot must return to its charging dock in the living room before the battery runs out. Also, the robot should not clean a room that is already clean. The task of the team is to create a plan for the vacuum robot that ensures all rooms are clean and makes the robot return to the charging dock. The plan should minimize the number of moves and cleaning operations to conserve battery life.

#### Initial Conditions:

- The living room (LR) is dirty.
- The kitchen (K) is clean.
- The bedroom (BR) is dirty.
- The robot is located in the living room (LR).
- The robot's battery is full (F).

#### Actions that Robot Can Take:

- Move from one room to another using the battery ( Move( $\cdot$ ) )
- Clean a room using the battery ( Clean( $\cdot$ ) )
- Charge the battery in the living room ( Charge( $\cdot$ ) )

#### Constraints:

- The robot cannot take action if the battery is low.
- The robot can clean the room where the robot is currently at.
- The robot can charge its battery only in the living room.

#### Propositional Logic Symbols:

- $LRC, LC, BRC$  indicate that the living room, kitchen, and bedroom are clean, respectively.
- $LRD, LD, BRD$  indicate that the living room, kitchen, and bedroom are dirty, respectively.
- $R_{LR}, R_L, R_{BR}$  indicate the robot's location, respectively.
- $B_F, B_L$  indicate the battery is full or low, respectively.

58 If the initial state is written as  $LRD \wedge KC \wedge BRD \wedge RLR \wedge BF$  (i.e.,  $LRD \wedge KC \wedge BRD \wedge RLR \wedge BF$ ), what is the goal state you want to achieve? (2 points)

---

59 What are the possible actions that can result in a state change that does not include battery low? Select the applicable clauses. (2 points)

- Move(LR)
- Move(K)
- Move(BR)
- Clean(LR)
- Clean(K)
- Clean(BR)
- Charge()

**Q10** Planning under uncertainty

Bob runs a multifamily short term rental property who charges different rates for a room as the day progresses, depending on whether he has many or few vacancies, to ultimately maximize his daily income. Customers will accept a rate  $r_i$  with probability  $p_i$ , reject with  $1 - p_i$ , in which case the customer leaves, and does not return that day. Help Bob maximize his income.

Assumptions

- Bob's income is \$0 in the case that there are no customers,
- Bob's income is \$0 in the case that he has no rooms,
- Bob's rates and their probabilities are illustrated in Table 1,
- Discount factor is 1 ( $\gamma = 1$ ).

Rate (r)	\$100	\$200	\$500	\$800
Probability customer accepts rate (p)	0.80	0.35	0.12	0.06

Table 1: Rates per day, and probability customer will accept.

60. Fill in the utilities of each state (4 points)

Hint: Utilize the utility function below to compute your expected utilities. Be sure to keep a record of the rate you chose for each state.

$$U(i, j) = \max_{n=1,2,3,4} [p_n(r_n + U(i-1, j-1)) + (1-p_n)U(i, j-1)]$$

	1 Customer	2 Customers	3 Customers	4 Customers	5 Customers	6 Customers	7 Customers
1 Vacancy	80.000000			213.781760			313.096509
2 Vacancies							
3 Vacancies			240.000000				
4 Vacancies					386.969600		
5 Vacancies		160.000000				469.575680	

**Q10.2** Known number of customers

Suppose Bob knows the number of customers that will ask for a room.

Bob has 5 vacancies and 7 customers will ask for a room.

**61.** What rate should Bob charge for this room? (2 points)

**62.** Suppose Bob fills one of his rooms with his first customer. Given Bob's computed utilities. What rate should Bob charge for the next room? (Case where 4 vacancies, and 6 customers) (2 points)

**63.** Bob has 1 vacancy and 3 customers will ask for a room.

What rate should Bob charge for the room? (1 point)



Q10.3 Unknown number of customers

Bob has lost his ability to know how many customers will want to book a room. Bob decides to employ a policy assuming that the current customer is the last customer of the day.

Bob has 1 vacancy. A customer enters and requests a room.

64. What rate should Bob charge for the room? (2 points)

65. Bob's first customer has decided to not book the room and leaves. A second customer enters. What rate should Bob charge for the room? (2 points)

66. Select all that are true in the case of unknown number of customers (1 point)

- Bob is assuming his state estimates are exact
- Bob is employing a one-step lookahead policy
- Bob's policy is optimal
- Bob is selecting a rate that maximizes  $p_n(rn + U(i-1, j-1) - U(i, j-1))$