Written examination
TIN172/DIT410, Artificial Intelligence

Contact person: Peter Ljunglöf, 0736–24 24 76

27 May 2014, 14:00–18:00

This examination consists of eight basic questions (numbered 1–8) and three advanced (numbered A–C).
There are no points awarded for the questions, but you can either give a correct answer, or fail.

Grading

The number of questions (basic plus advanced) that you need to answer correctly in order to get a certain grade is shown in the following table:

Basic questions	Advanced questions	Final grade
≥ 6	—	3/G
≥ 7	≥ 1	4/VG
≥ 8	≥ 2	5

Your result

Accessories

- Paper and pencil.
- Crayons, paper glue, scissors.
- One A4 cheat sheet with any information you want on it.
- No books or calculators.

Notes

- Answer directly on the question page,
  but you can also use empty papers if you run out of space.
- ***Write readable, and explain your answers!***

1 Treasure hunt (state space)

Imagine a treasure hunter agent ( ) who wishes to collect gold coins (③) in a maze like the one shown below. The agent is not directional and can move one step in any direction (**n, s, e, w**) at any time step, as long as there is no wall in the way. The gold coins do not move. The agent's goal is to find a plan for collecting all coins using as few moves as possible. Assume that the grid has size $M \times N$ and there are 4 coins in the maze.

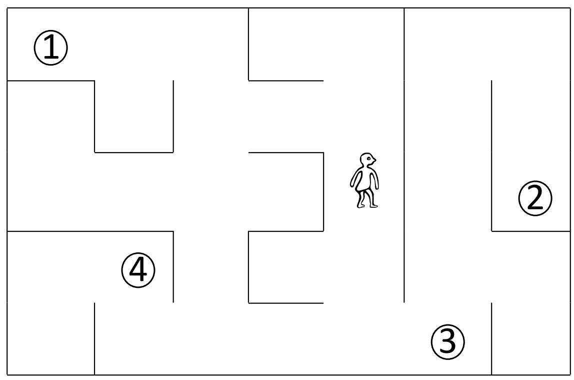

Give a suitable representation of the states in this searching problem.

What is the size of the state space?

2 Treasure hunt (heuristics)

This is a continuation of question 1.

Below are seven possible heuristics for the same searching problem in the previous question, where a denotes the agent, G is a set consisting of the gold coins that have not been collected yet, and $d(x, y)$ is the Manhattan distance between x and y, i.e., the sum of the horizontal and the vertical distance between x and y.

- $h_1 = \min_{g \in G} d(a, g) =$ the minimum Manhattan distance from the agent to any remaining gold coin
- $h_2 = \max_{g \in G} d(a, g) =$ the maximum Manhattan distance from the agent to any remaining gold coin
- $h_3 = \sum_{g \in G} d(a, g) =$ the sum of all Manhattan distances from the agent to the remaining gold coins
- $h_4 = \max_{g, g' \in G} d(g, g') =$ the maximum Manhattan distance between any two remaining gold coins
- $h_5 = h_1 + h_4$
- $h_6 = h_2 + h_4$
- $h_7 = h_3 + h_4$

Which ones of these heuristics are admissible? (Hint: it's more than one, but not all)

3 Cost-based search

The following is a representation of a search problem, where S is the start node and G is the goal. There is also a heuristics h which is defined in the table.

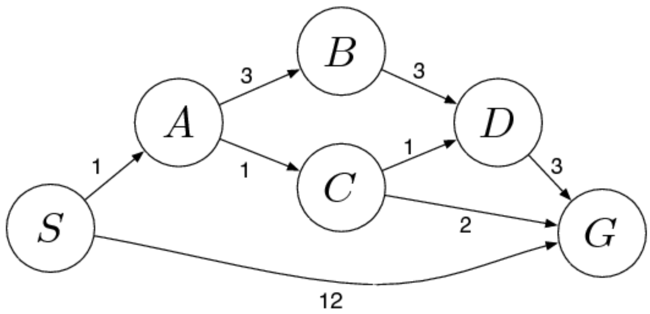

n	$h(n)$
S	3
A	2
B	1
C	4
D	2
G	0

Assume that you are in the middle of a search and the current frontier is $\mathbf{F} = \{B, C, G\}$, and you are about to select the next node from $\mathbf{F}$ to expand.

Which node will be expanded next, assuming that you are using…

- **…lowest-cost-first search (also known as uniform-cost search)?**

- **…greedy best-first search?**

- **…A* search?**

4 Rolling the dice

Assume that you have two variables, A and B, both with domain $\{1, 2, 3, 4, 5, 6\}$. The constraints on the values are that $A < B$ and $A + B = 5$.

Draw the constraint graph.

Perform arc consistency on the graph.

5 Proving in a knowledge base

Consider the following knowledge base:

$a \leftarrow b \wedge c \wedge f$	$d \leftarrow c \wedge f$
$b \leftarrow c \wedge e$	$e \leftarrow d \wedge g$
$b \leftarrow d \wedge c$	f
c	$g \leftarrow e \wedge c$

Is a a logical consequence?

Either prove it using the top-down or bottom-up proof procedure,
or explain why it is not a consequence.

6 Abdul Alhazred the psychic: Bayesian network

Abdul Alhazred (AA) claims that he is psychic and can always predict a coin toss. Let $P(A) = p_0 = 0.1$ be your prior belief that AA is a psychic, and let B_k denote the event that AA predicts the k-th coin toss correctly. In all questions about Abdul the psychic, make the following assumptions:

1. All experiments are conducted with a fair coin, whose probability of coming heads at the k-th toss is $P(H_k) = \frac{1}{2}$ and where H_k is independent of $H_{k-1}, \dots, H_1$.
2. Your utility for money is linear, i.e. $U(x) = x$ for any amount of money x.

What are the dependencies between A, H_1, B_1, H_2 and B_2? Draw a Bayesian network to represent them.

7 Abdul the psychic: Marginal probabilities

This is a continuation of question 6. In short: $P(A) = p_0 = 0.1$ is your prior belief that AA is a psychic. B_k is the event that AA predicts the k-th coin toss correctly. The probability of the coin coming heads at the k-th toss is $P(H_k) = \frac{1}{2}$, where H_k is independent of $H_{k-1}, \dots, H_1$. Your utility for money is linear, $U(x) = x$.

What is the marginal probability $P(B_1)$?

What is the marginal probability $P(B_2)$?

8 Abdul the psychic: Conditional probabilities

This is a continuation of question 6. In short: $P(A) = p_0 = 0.1$ is your prior belief that AA is a psychic. B_k is the event that AA predicts the k-th coin toss correctly. The probability of the coin coming heads at the k-th toss is $P(H_k) = \frac{1}{2}$, where H_k is independent of $H_{k-1}, \ldots, H_1$. Your utility for money is linear, $U(x) = x$.

Now assume that AA predicts the first coin toss correctly, i.e. B_1 holds.

What is the marginal probability $P(B_2|B_1)$?

A [Advanced] Abdul the psychic: Betting

This is the final continuation of question 6. In short: $P(A) = p_0 = 0.1$ is your prior belief that AA is a psychic. B_k is the event that AA predicts the k-th coin toss correctly. The probability of the coin coming heads at the k-th toss is $P(H_k) = \frac{1}{2}$, where H_k is independent of $H_{k-1}, \ldots, H_1$. Your utility for money is linear, $U(x) = x$.

At the beginning of the experiment, AA bets you 100€ that he can predict the next four coin tosses, i.e. he is willing to give you 100€ if he doesn't predict all four tosses.

How much are you willing to bet against, i.e., how much are you willing to pay if he predicts them correctly?

B [Advanced] Rob, the coffee delivery robot

Rob is a coffee delivery robot who lives in a world with three locations: a coffee shop (*cs*), a laboratory (*lab*), and Sam's office (*off*). Rob can pick up coffee (*puc*) at the coffee shop, move (*mc*/*mcc*), and deliver coffee (*dc*). Delivering the coffee to Sam's office will stop Sam from wanting coffee. Assume that we represent the planning problem with two boolean features (*rhc* – "rob has coffee", and *swc* – "Sam wants coffee"), and one three-valued feature (*rloc* – "Rob's location"). Then the actions can be defined like this:

Action	*mc* (move clockwise)		
Precondition	[*rloc=cs*]	[*rloc=off*]	[*rloc=lab*]
Effect	[*rloc=off*]	[*rloc=lab*]	[*rloc=cs*]

Action	*puc* (pick up coffee)
Precondition	[*rloc=cs, ¬rhc*]
Effect	[*rhc*]

Action	*mcc* (move counterclockwise)		
Precondition	[*rloc=off*]	[*rloc=lab*]	[*rloc=cs*]
Effect	[*rloc=cs*]	[*rloc=off*]	[*rloc=lab*]

Action	*dc* (deliver coffee)
Precondition	[*rloc=off, rhc*]
Effect	[*¬rhc*, *¬swc*]

Assume that the start state is [*rloc=off, ¬rhc, swc*]. How does a regression planner solve the goal [*¬swc*]?

You do not have to draw the entire search tree (which is infinite anyway), but you should at least show all possible choice points that the planner has to go through and what choices it has to make in order to find a solution. Also, describe the search state at each choice point.

C [Advanced] A chain of variables

Assume that you have a chain of n variables $X_1, \ldots, X_n$, each with the domain $\{1, \ldots, m\}$, where $n \leq m$. The chain should be in ascending order, meaning that $X_i < X_{i+1}$ for all $1 \leq i < n$.

What will the resulting domains of variables $X_1, \ldots, X_n$ be, after arc consistency has been enforced?