

Written by (anonymous ID): _____

Corrected by (anonymous ID): _____

Written examination TIN173/DIT410, Artificial Intelligence

Tuesday 3 May 2016, 8:00–13:00

Examiner: Peter Ljunglöf

This examination consists of six questions. A correctly answered question gives you 2 points, the total number of points are 12.

Grades: To get grade 3/G/pass you need at least 66% correct, i.e., 8 points.

This is only for students from previous years:

To get Chalmers grade 4 you need at least 10 points,
to get GU grade VG/pass with distinction you need at least 11 points,
and to get Chalmers grade 5 you need all 12 points.

Tools: Paper and pencil. *Do not use a red pen when writing the exam!*

No extra books, papers or calculators.

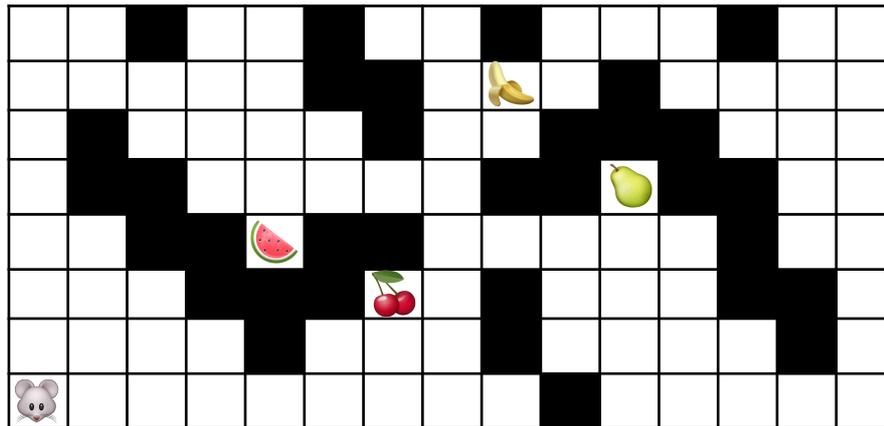
Review: This exam is peer reviewed. This means that you will correct each other's theses. To ensure safety, you should use a red pen while correcting another one's thesis. To ensure anonymity, you should write your ID number at the top of the paper.

Notes: Answer every question directly on the question paper, and write your ID number at the top of every paper.

Write readable, and explain your answers!

1. Fruit-eating mice

A hungry mouse wants to eat all four fruits in a maze such as the one below, in as few moves as possible. At each turn the mouse can move any number of squares in one of the directions up, down, left or right, but it is not allowed to enter (or jump over) any walls (i.e., the black squares). Thus, the mouse moves just like a rook in chess. To eat a fruit, the mouse has to stop at that square.



Assume that the maze has 4 fruits, and a size of $b \times h$ squares.

- a) Give a suitable representation of the states in this searching problem.
Remember to write down the domains of the variables in your representation.

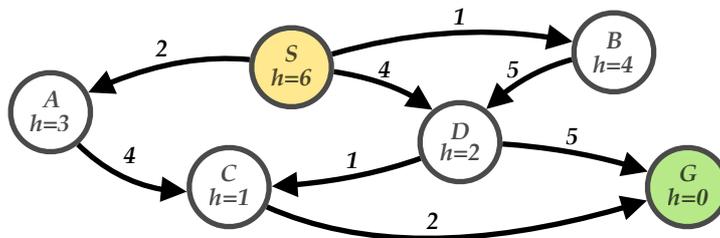
correct?

- b) How many possible actions can the mouse perform at each move?
(I.e., what is the branching factor?)

correct?

2. Search algorithms

The following is a graph of a search problem, where the directed arcs represent the successors of a node. The cost of moving to a node is given by the number on the arc. The value of the heuristic function h is shown inside each node. The start state is S and the goal is G.



For each of the following search strategies, state the order in which states are expanded (i.e., when they are removed from the frontier), as well as the final path returned when the search is finished.

Assume that all ties are resolved in alphabetical order (i.e., the A state is expanded before the B state which is expanded before the C state, etc.).

a) Breadth-first search

correct?

<i>Order in which states are expanded</i>	
<i>Resulting path found by the search algorithm</i>	
<i>Cost of the resulting path</i>	

b) A* search

correct?

<i>Order in which states are expanded</i>	
<i>Resulting path found by the search algorithm</i>	
<i>Cost of the resulting path</i>	

3. Admissible heuristics

- a) Assume that h_1 and h_2 are two different heuristics, that both are admissible.
Which of the following heuristics are admissible? (several might apply)

correct?

- | | |
|--|--|
| <input type="checkbox"/> $h_1 + h_2$ | <input type="checkbox"/> $0.9 \cdot h_1 + 0.5 \cdot h_2$ |
| <input type="checkbox"/> $h_1 - h_2$ | <input type="checkbox"/> $0.7 \cdot h_1 + 0.3 \cdot h_2$ |
| <input type="checkbox"/> $\min(h_1, h_2)$ | <input type="checkbox"/> $\max(h_1, h_2)$ |
| <input type="checkbox"/> None of the above | |

- b) Give one disadvantage and one possible advantage with using a non-admissible non-negative heuristics when performing A* search?

correct?

Advantage:

Disadvantage:

4. Dressed for success

You wake up early in the morning, full of excitement and anticipation – today is *the big day*, it's the written examination for the AI course! For you this is a very special occasion, and you want to show to the world how important it is to you, so you decide to dress up very special. After a while you have narrowed down the choices for each of the four items of clothing, (H)eadwear, (B)odywear, (L)egwear, and (A)ccessory as follows:

$H \in \{ \textit{hat}, \textit{cap} \}$

$B \in \{ \textit{shirt}, \textit{blouse}, \textit{jumper} \}$

$L \in \{ \textit{leggings}, \textit{skirt}, \textit{trousers} \}$

$A \in \{ \textit{scarf}, \textit{tie}, \textit{cravat} \}$

Furthermore, you have derived the following constraints:

- (1) if you choose the *jumper*, the *cap* is the only matching headwear
- (2) the *leggings* do not go together with the *hat*
- (3) if you wear a *shirt* or a *blouse*, then you have to take a *tie* or a *cravat*

And now for the first question:

- a) Draw the constraint graph. Use the variables H, B, L, A as nodes, and label the edges with the numbers of the constraints (1)–(3).

correct?

For question (b), see next page.

4. Dressed for success (continued)

If you have drawn the constraint graph correctly, it should already be arc-consistent. Now, assume that you decide to take the *hat*. What can you deduce about your other clothing?

- b) Perform arc-consistency on the problem after setting $H = \textit{hat}$, using the table below. Use $H \rightarrow B$, $B \rightarrow H$, $L \rightarrow B$, $B \rightarrow L$, etc., to denote edges.

Because the initial problem is arc-consistent, you can start with the queue of only the edges going into H.

You can stop when the graph is arc-consistent, even if you have edges left in the queue (i.e., you don't have to remove the final edges from the queue, when you have reached AC).

correct?

<i>Edge removed from queue</i>	<i>Affected variable</i>	<i>Constraint to be checked</i>	<i>Values removed from domain (if any)</i>	<i>Edge(s) added to queue (if any)</i>

<i>Variable</i>	<i>Final domain</i>
H	
B	
L	
A	

5. Crashing autonomous cars

Volvo, Tesla and Renault have decided to roll-out their autonomous cars in the open market next year. The AI in the testing phase for the three companies shown different degrees of reliability: 98% for Tesla, 97% for Volvo and 93% for Renault.

Market research shows that 5 out of 10 customers choose to buy the Tesla, 3 prefer Volvo and 2 prefer the Renault. Given these statistics, what are the chances of autonomous car malfunction occurring on the roads once these are released?

- a) What prior and conditional probabilities are specified in the text?
Formulate the probabilities in terms of the two random variables
C (the brand of the autonomous car), and M (autonomous car malfunction).

correct?

- b) How do you calculate the chance of malfunction?

correct?

6. Getting wet in Blöteborg

Kim lives in Blöteborg, and every morning, before she goes to their AI lecture, she looks at the sky to see if it is cloudy. If it is, she usually brings an umbrella (about 80% of the time), but if it's not cloudy, there's a 10% chance that she happens to bring the umbrella anyway.

If it is cloudy in the morning, then there's a 70% chance that it will rain in the afternoon. If the morning sky is clear, there's still a 30% possibility of rain in the afternoon. In Blöteborg 60% of the mornings are cloudy.

On the way home in the afternoon, Kim might get wet. If it is raining and she brought an umbrella, there's only a 20% chance that she gets wet, but if she forgot the umbrella, it is absolutely certain that she gets wet. If it's not raining, there's still a 10% chance that Kim gets wet, and this is regardless of whether she has an umbrella or not.

- a) Draw the Bayesian network corresponding to this problem (assuming that the only dependencies are the ones described in the text). Use the variables C, U, R and W: "it was (C)loudy in the morning", "Kim brought an (U)mbrella", "there was (R)ain in the afternoon", and "kim got (W)et on the way home".

correct?

Next to each node in the network, write its associated probability table.

6. Getting wet in Blöteborg (continued)

- b) What is the probability that the sky was cloudy in the morning, provided that Kim got wet in the afternoon?

correct?

Note: You don't have to perform all calculations, but you have to show what you would do if you had a calculator.

Tear-off sheet (you don't have to hand this in!)

Here is the background text for the two questions that are split between pages.
You can tear this page off so that you don't have to flip pages too much.

Question 4. Dressed for success

You wake up early in the morning, full of excitement and anticipation – today is *the big day*, it's the written examination for the AI course! For you this is a very special occasion, and you want to show to the world how important it is to you, so you decide to dress up very special. After a while you have narrowed down the choices for each of the four items of clothing, (H)eadwear, (B)odywear, (L)egwear, and (A)ccessory as follows:

$H \in \{ \textit{hat}, \textit{cap} \}$
 $B \in \{ \textit{shirt}, \textit{blouse}, \textit{jumper} \}$
 $L \in \{ \textit{leggings}, \textit{skirt}, \textit{trousers} \}$
 $A \in \{ \textit{scarf}, \textit{tie}, \textit{cravat} \}$

Furthermore, you have derived the following constraints:

- (1) if you choose the *jumper*, the *cap* is the only matching headwear
- (2) the *leggings* do not go together with the *hat*
- (3) if you wear a *shirt* or a *blouse*, then you have to take a *tie* or a *cravat*

Question 6. Bayesian network

Kim lives in Blöteborg, and every morning, before she goes to their AI lecture, she look at the sky to see if it is cloudy. If it is, she usually brings an umbrella (about 80% of the time), but if it's not cloudy, there's a 10% chance that she happens to bring the umbrella anyway.

If it is cloudy in the morning, then there's a 70% chance that it will rain in the afternoon. If the morning sky is clear, there's still a 30% possibility of rain in the afternoon. In Blöteborg 60% of the mornings are cloudy.

On the way home in the afternoon, Kim might get wet. If it is raining and she brought an umbrella, there's only a 20% chance that she gets wet, but if she forgot the umbrella, it is absolutely certain that she gets wet. If it's not raining, there's still a 10% chance that Kim gets wet, and this is regardless of whether she has an umbrella or not.