NTNU
Institutt for datateknikk
og informasjonsvitenskap

Norwegian University of Science and Technology
Department of Computer and Information Sciences

Examination paper in

TDT4171 – Artificial Intelligence Methods

May 25th, 2010, 09:00 - 13:00

Contact during examination: Helge Langseth, 735 96488
Quality assurance of examination paper: Shengtong Zhong

Language: English
Examination support: D
No written and handwritten examination support materials are permitted.
A specified, simple calculator is permitted.

Grading done by: June 15th.

Read the questions thoroughly. Make sure to understand exactly what is asked for.

If you think that some information is missing in a question, make a short note about the assumptions you think you need to make to be able to answer the question, make the assumptions, and answer the (modified) problem.

All questions (including all sub-questions) shall be answered. Each question is weighted as indicated in the text. The question-set consists of 5 questions and is on five pages (including this cover-page).

<u>Question 1</u> – Bayesian Networks (25%)

a) Describe the syntax and the semantics of Bayesian Networks.

b) Model the following domain using a Bayesian network. Make your model as simple and easy to understand as possible:

Arne is at the office 50% of his life. When Arne is at the office, he leaves the light turned on 50% of the time. When he is not at the office, he still leaves the light on 10% of the time. Arne is logged on to his computer 100% of the time when he is at the office. He can also log on remotely, and is therefore logged on 10% of the time when he is not at the office.

You are asked to give both the graphical structure ***as well as*** the conditional probability distributions here.

c) A student checks Arne's online status, and finds he is indeed online. Should this change the student's belief regarding whether or not Arne's office light is on? Why/why not?

d) Calculate the probability that Arne is at the office at a given point of time, if we know that he is logged on to his computer.

e) Do you think that Bayesian Networks constitute a *natural* modeling framework for this problem domain? How can one characterize problem domains where Bayesian networks can be used with success? Can you give an example of a problem domain where Bayesian networks do *not* fit well?

<u>Question 2</u> – Case-based reasoning (15%)

Describe the four main steps of the case-based reasoning (CBR) cycle. What is the difference between *instance-based reasoning* (like *k*-nearest neighbor) and "*typical CBR*"?

Question 3 – Markov Processes (20%)

a) Explain the *Markov assumption* using your own words. Give one example of a problem domain where the Markov assumption is (approximately) correct, and one where it is (blatantly) wrong.

b) Let us consider a robot operating in a grid-world with two states denoted *up* and *down* (see Figure).

up
down

Every time the robot is in state *up* it is rewarded one point, while it is rewarded zero points when it is in state *down*, i.e., *R(up)=1*, *R(down)=0.* The robot has two available actions, *move* and *stay*. The actions are correctly performed with probabilities p_{move}=.0.8 and p_{stay}=.0.9, respectively. If the robot performs *move* successfully, it will change its location in the grid. If the action fails, the robot remains at its current position. If the robot performs the action *stay* successfully it will remain at its current location. If the action fails, the location changes.

Describe the problem domain formally as a *Markov decision process*.

Show the first four steps of the *value iteration* algorithm to find the optimal policy for the robot. We use γ =0.5. You should write up both the calculated utilities as well as the corresponding policy for each step of the algorithm.

Explain how you find the optimal policy from the calculated utilities.

Hint: The Bellman-equation looks like this:

$$U_{i+1}(s) \leftarrow R(s) + \gamma \cdot \max_a \sum_{s'} P(s'|s,a) \cdot U_i(s')$$

Question 4 – Conditional probabilities (20%)

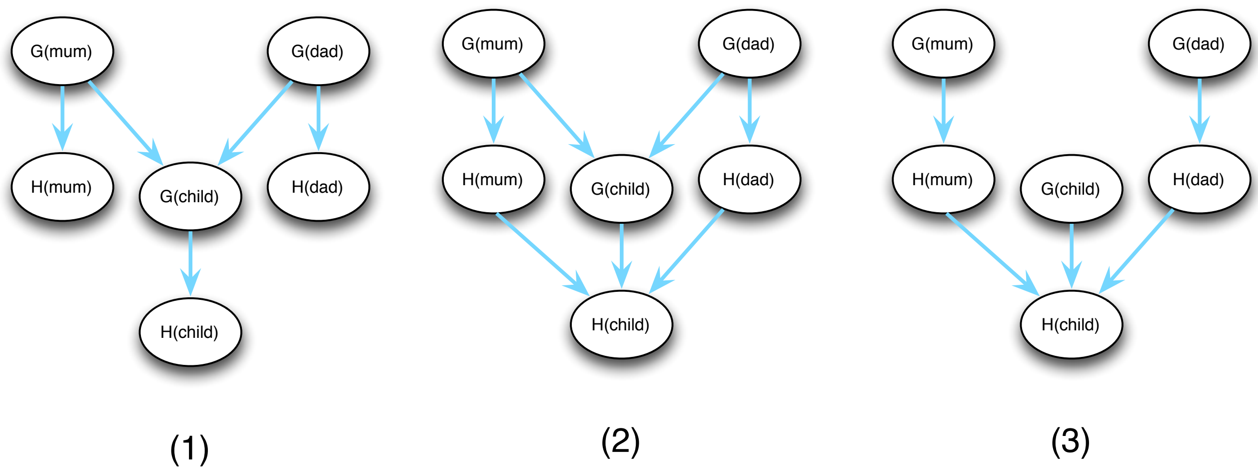

Let *H(x)* be the random variable telling the preferred hand of a person *x*, taking the values *left* or *right*. (So, for instance, *H(dad)* holds information about the preferred hand for the person *dad*.)

One possible explanation model for which hand a person uses is that there is a gene which decides the preferred hand (with some probability). Let *G(x)* represent this gene for person *x*, and assume that also this variable takes on the values *left* and *right*. Assume that *H(x)* takes on the same value as *G(x)* with a given probability. Assume further that a child inherits its gene from either mum or dad, and that it is equally likely that either parent is "chosen" in that respect. Finally, the child's gene can *mutate*, i.e. become the opposite value of the parent with a given probability *m>0*.

a) Which (one or several) of the three networks given above asserts that **P**(*G(mum),G(dad),G(child)*) = **P**(*G(mum)*) * **P**(*G(dad)*) * **P**(*G(child)*) ?

b) Which (one or several) of the three networks given above asserts the same ***conditional independence statements*** as the explanation model described above?

c) Which of the three networks given above is the ***best representation*** of the explanation model described? Give your reasons for the answer.

d) Write down the conditional distribution for *G(child)* given *G(mum)* and *G(dad)*, i.e. **P**(*G(child)* | *G(dad)*, *G(mum)*), based on the structure in network (2).

<u>Question 5</u> – Mixed questions (20%)

a) What does it mean that an agent is *rational*? What is the connection between rational agents and the *maximum expected utility principle*, and what does the maximum expected utility principle say? How can this connection help us when we want to design rational agents?

b) What is *The Strong AI Hypothesis*, and what is *The Weak AI Hypothesis*? Give your view of whether the strong AI hypothesis is fulfilled or not. Refer to acknowledged arguments (e.g., those presented in the syllabus) to strengthen your view.

c) Two of the reasoning methods used with Hidden Markov Models are *filtering* and *smoothing*. Explain what these techniques do, and give practical examples where Hidden Markov Models combined with each of the two reasoning techniques can be useful. What are the differences between the two techniques?

d) We have discussed both neural networks and decision trees during this course. Give a learning situation where you would prefer to work with neural networks to decision trees, and another situation where you find decision trees more suitable than neural networks. Give reasons for your answer; you can, for instance, relate your discussion to which types of data the two work best with, what the properties of the corresponding learning algorithms are, etc.