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

Tuesday May 25th, 2010, 15:00 - 19:00

Contact during examination: Helge Langseth, 735 96488
Quality assurance of examination paper: Agnar Aamodt

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

Grading done by: Tuesday 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 four pages (including this cover-page).

Question 1 – Bayesian Networks (20%)

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:

Peter wakes up in the morning with a sore throat. It can be either due to a common cold or due to an angina. Peter wants to know what is wrong with him, as he will stay at home if he has the angina. To better understand his problem, he can:

- *Measure his temperature to detect if he has a fever*
- *Look for yellow spots in his throat*

Both the common cold as well as the angina can result in a fever, but only angina gives rise to the yellow spots in the throat.

What modelling assumptions about independences and conditional independences can be read off the model?

You are not asked to give the conditional probability distributions here, only the qualitative of the Bayesian network (the graph).

c) Which conditional probability tables are required for the modelling of the domain to be complete? Note that you are not asked to come up with the numbers of these conditional distributions, only list which conditional distributions are required in your Bayesian network.

d) Do you think that Bayesian Networks constitute a *natural* modelling framework for this problem domain? How can one characterise 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?

Question 2 – Case-based reasoning (20%)

a) How does the instance-based reasoning method *k-nearest neighbour* solve problems? What are the strengths and weaknesses of this method?

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

Question 3 - Learning (25%)

a) Research on Artificial Neural Networks is partly motivated by knowledge about how the brain works. Explain this relation.

b) Artificial Neural Networks are often trained using the "Backpropagation Algorithm". Describe the strengths and weaknesses of Artificial Neural Networks trained in this way. Your discussion should, for instance, cover topics like what types of data the networks can handle, which assumptions – if any – one has to make regarding the shape of the target function, properties of the learning algorithm, etc.

c) Create a multilayer artificial neural network, and find weights to ensure that the network represents the function

$$\textbf{Exact2}(x_1, x_2, x_3),$$

where x_1, x_2 and x_3 are binary inputs; that is, the output is true if exactly two of x_1, x_2 or x_3 is true, and false otherwise. You should make the network as simple as possible (i.e., containing as few nodes as possible).

Can the function **Exact2**(x_1, x_2, x_3) be represented exactly by a *perceptron* (a neural network *without* a hidden layer)? Explain your answer.

***Hint*: In question (c) you can let the transfer function in the nodes be the *step-function*.**

Question 4 – Markov Processes (25%)

a) Explain the *Markov assumption* using your own words. Give an example of a problem domain where the Markov assumption appears to be (approximately) correct, and an example of a domain where it does *not* fit.

b) *Kalman Filters* define a special class of Markov models. Which assumptions are made when working with Kalman Filters? Explain **why** each of the assumptions is needed, and what the result would be if each of the assumptions was lifted.

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

(It may be easier to explain this if you use a drawing of a Kalman filter in your description.)

<u>Question 5</u> – Mixed questions (10%)

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.