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

Friday May 29th, 2009, 09:00 - 13:00

Contact during examination: Helge Langseth, 735 96488
Quality assurance of examination paper: Tor Gunnar Houeland

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

Grading done by: Friday June 19th.

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 three 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 consider to take the following subjects at the university:

- ***Artificial Intelligence Methods*** *(requires ability to do* ***probability calculus****,* ***logic reasoning*** *and* ***programming****)*
- ***Logic*** *(requires* ***logic thinking****,* ***linguistic instinct****,* ***programming****)*
- ***Latin*** *(requires* ***linguistic instinct*** *and* ***logic thinking****).*

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

c) Extend the model from question (b) with the following information:

- *Peter has good results in*
  - ***Introduction to Programming*** *(requires* ***programming****)*
  - ***Mathematics*** *(requires* ***logic thinking****)*
  - ***French*** *(requires* ***linguistic instinct****).*
- *He failed* ***Introduction to Statistics*** *(requires* ***probability calculus****).*

How can Peter utilise this extra information to say something about the probability to pass each of the three subjects he considers taking? (Algorithms for the calculations are not expected here, only a qualitative discussion about how the evidence influences the query variables.)

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 (15%)

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

Question 3 - Learning (20%)

a) Give an example of a decision tree with two internal nodes (including the root of the tree). Explain the syntax. Explain how the decision tree will classify a new example.
b) Describe the learning algorithm for decision trees from the curriculum (ID3) informally. What is *Occam's razor*, and how does it come in play when learning decision trees?
c) *Gradient descent* is a powerful general-purpose algorithm, which (among other things) can be employed to learn the weights of a perceptron network. Describe the main steps in the algorithm, and discuss the algorithm's main strengths and weaknesses.

Question 4 – Markov Processes (30%)

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) Hidden Markov Models define a special class of Markov models. Which assumptions are made when working with Hidden Markov Models? Explain why each of the assumptions is needed.
c) What is a Markov Decision Process? What are these models used to model, and what assumptions are used when building these models? Briefly describe an algorithm to solve a Markov Decision Process.

Question 5 – Mixed questions (15%)

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 "*Turing test*", and how does this test relate to the concept "*Artificial intelligence*"? Discuss to what extent the Turing test is relevant if one wants to claim that a system is intelligent.
c) What does an *Information Retrieval System* do, and what are the main components of such a system? Give an example of an information retrieval system that is in daily use by "ordinary people". Information retrieval systems are typically evaluated using two criteria. Define these two criteria. Why is it not sufficient to use only one of these criteria to evaluate an information retrieval system?