Victoria University of Wellington

DEGREE EXAMINATIONS - 1999

COMP 421

END OF YEAR

COMP 421

ARTIFICIAL INTELLIGENCE

Time allowed: THREE HOURS

Instructions: Attempt all SIX questions.
Each question is worth 20 marks.
Marks for subparts of question are indicated.
There are 120 possible marks on the exam.
Calculators and two sides of personal hand-written notes are permitted.
Keep answers for each question together.

Questions:

1. Search and evolutionary computation
2. Neural networks
3. Information and brains
4. Bayesian inference
5. Belief networks
6. Reinforcement learning

continued…

Question 1. Search, and evolutionary computation.

(a) Describe the Metropolis and Gibbs variants of Markov Chain Monte Carlo, for sampling from probability distributions. How can such methods be used to search for the optimum in some space of parameters? [8 marks]

(b) What guarantees are there on the performance of such methods? [1 mark]

(c) Give a theoretical argument in support of genetic algorithms, and a theoretical reason for caution. [4 marks]

(d) What is "roulette wheel" selection? [3 marks]

(e) Describe the use of momentum and the conjugate gradient method, indicating how these improve on simple gradient descent in searching for a minimum of a cost function. [4 marks]

continued…

Question 2. Neural networks

(a) How does the Hopfield network differ from the Boltzmann machine? [3 marks]

(b) Consider a neural network which is being trained by conventional backpropagation (*i.e.* minimization of the mean of its squared errors on patterns in a training set). Show that this is equivalent to maximizing the likelihood of the network getting the training data correct, under a Gaussian noise model.
[5 marks]

(c) Consider a neural network composed of linear threshold units. The net has three inputs, feeding into three hidden units. These hidden units feed into a single output unit, and there are no direct connections from input to output.
Give weights and biases so that the network implements the parity function: *i.e.*, it outputs a one if exactly 1 or 3 input units are active, and otherwise outputs a zero.
[6 marks]

(d) Describe the "wake-sleep" algorithm for Helmholtz machines. [6 marks]

Question 3. Information and Brains

The Willy Wonka Chocolate Factory is constantly making chocolate bars. The computer in charge of the process codes good bars as a zero. Once in a while the chocolate machine makes an error, which the computer records as a one. Studies show that the probability of a mistake is $1/16$. A sample record is given by

0000100000001010000001000

You suggest that the company encode the data in blocks of two bits, using the following code:

Data block	Code
00	0
01	11
10	100
11	101

(a) Show that this is a Huffman code, and use it to encode the given data record.
[2 marks]

(b) The above code uses more bits to encode the data in the example than an *optimal* code would. How many more bits does it use?
[3 marks]

(c) Give a more efficient code than the one in (a). [5 marks]

(d) "Mammalian brains are modular". Discuss. [10 marks]

continued…

Question 4. Bayesian inference

According to the Bayesian point of view, in making predictions we should integrate over models and parameters.

(a) Explain what is meant by the above statement. [6 marks]

(b) Relate this theory to ML (maximum likelihood), and MAP (maximum *a posteriori*) learning. [6 marks]

Consider the sequence [-1, 3, 7…]. The task is to guess what the next three numbers will be.

Hypothesis *Line* is that the sequence is generated by a process $x_t = c_0 + c_1 t$, which has two parameters. For instance using parameters $c_0 = -1, c_1 = 4$ fits the known data exactly and gives predictions of 11, 15, and 19, for the next three numbers.

Hypothesis *Cubic* is that the sequence is generated by a process $x_t = c_0 + c_1 t + c_2 t^2 + c_3 t^3$, which has four parameters. One choice for these parameters is $c_0 = -1$, $c_1 = 6$, $c_2 = -3$, and $c_3 = 1$. This fits the data exactly, but gives predictions of 17, 39, and 79 for the next three numbers. (Note: many other choices for the parameters will also fit the data given).

(c) Briefly, how would a Bayesian theorist approach the problem of choosing between these two hypotheses? [4 marks]

(d) Describe a *non-Bayesian* technique for comparing the two models on the data given. [2 marks]

(e) Give one major advantage, and one major disadvantage, of using the method given in (d), compared to the Bayesian method. [2 marks]

Question 5. Belief networks

Consider this belief network:

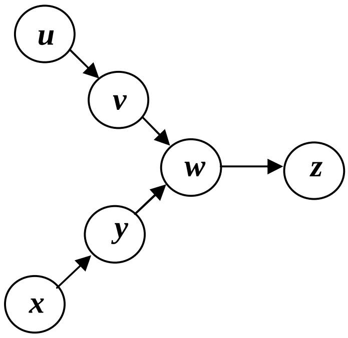

(a) Write down the joint distribution as a product of factors, and draw the factor graph. [3 marks]

(b) Assuming all variables are binary, how many probability values need to be specified? [2 marks]

(c) Are variables ***u*** and ***x*** dependent or independent, given the value of ***z***? [2 marks]

(d) For the network above, give an example of EACH of the following types of inference:
1. causal
2. diagnostic
3. inter-causal ("explaining away") [3 marks]

Consider the belief network shown below, which could be described as a "coupled hidden Markov model":

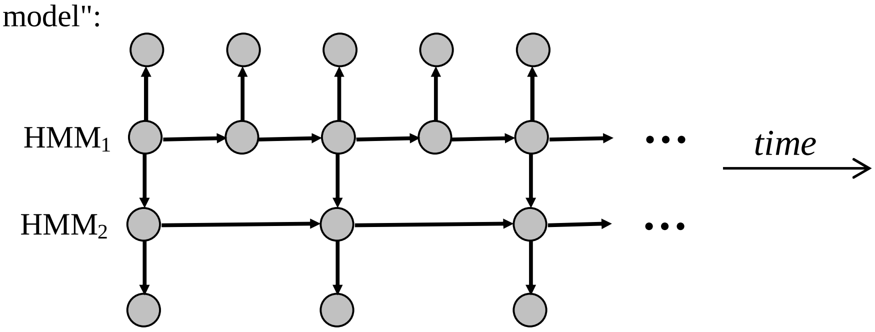

(e) Describe the world-model embodied by this network. [4 marks]

(f) Outline the steps in applying the Junction Tree Algorithm to this network. What does the resulting graph imply about the efficiency of inference in this network? [6 marks]

Question 6. Reinforcement learning

(a) How does one use a decision network to choose between possible actions? [4 marks]

(b) In reinforcement learning, what is the role of the discount factor, γ? [4 marks]

In learning from reinforcement, the 'Q-learning' algorithm only backs up a new value by one state per episode.

(c) Describe how Q(λ) alleviates this problem. [4 marks]

(d) How does Dyna-Q deal with the same problem? [4 marks]

In Markovian tasks, all the information required for an agent to act optimally is contained in the current state. An example of a Markovian task is pole-balancing, in which the agent knows the angle of the pole and the rate-of-change of this angle, and the task is to keep the pole upright for as long as possible. If the agent did *not* know the rate-of-change, the Markov property would not hold.

(e) Explain why TD(λ) is expected to cope with *non*-Markovian tasks better than TD(0). [4 marks]

****