VICTORIA UNIVERSITY OF WELLINGTON
Te Whare Wananga o te Upoko o te Ika a Maui

EXAMINATIONS — 2002

MID-YEAR

COMP 421

Artificial Intelligence

Time Allowed: 3 Hours

Instructions: There are 180 marks in total in this exam.
There are 12 questions: each question is worth 15 marks.
Answer all questions.
TWO pages (i.e. two sides only) of hand-written notes are permitted.

Question 1. No Free Lunch [15 marks]

The No Free Lunch theorem claims that all search algorithms for finding maxima of a surface perform the same, when averaged over all surfaces. However, an algorithm that 'searches' by taking uphill steps will inevitably end up at a higher point than one that searches by taking downhill steps. Explain this disagreement by referring to the particularly general notion of search used in the theorem.

Question 2. Markov Chain Monte Carlo [15 marks]

The Metropolis-Hastings algorithm for sampling from a probability distribution p accepts a transition from state i to state j with probability

$$\Pr_{i \to j} = min\left(1,\ \frac{p_j}{p_i}\frac{Q_{j|i}}{Q_{i|j}}\right)$$

where Q is the proposal distribution and $Q_{j|i}$ is the probability that state j is proposed from state i. Note that for a transition to occur it must be proposed AND accepted. Show that this dynamics will indeed sample states with their correct probabilities (p_i for the i-th state), by showing it obeys detailed balance.

Question 3. Searching for optima on surfaces [15 marks]

(a) [5 marks] Briefly explain how the introduction of 'momentum' can enhance the performance of gradient-based search algorithms.

(b) [5 marks] Describe Brent's algorithm for line search.

(c) [5 marks] What problem with naive line-search is solved by the method of conjugate gradients?

Question 4. Learning in a deterministic neuron [15 marks]

Consider a single model neuron producing real-valued outputs. It has learnable weights $\mathbf{w}$ and produces output $y = f(\phi)$ where

$$\phi = \sum_i w_i x_i$$

for input pattern $\mathbf{x}$, and f is the sigmoid function,

$$\frac{1}{1 + e^{-\phi}}$$

Suppose each input pattern $\mathbf{x}$ has an associated target output, Y. Showing each step, derive the learning rule for altering the weights so as to minimize the sum of squared errors over a training set of input-output pairs:

$$C = \frac{1}{2} \sum_\mu (Y_\mu - y_\mu)^2$$

Question 5. Batch *vs* online learning [15 marks]

So-called 'supervised' learning of neural networks, in which target outputs are readily available, can be trained *on-line* or in *batch* mode.

(a) [7 marks] Explain the distinction between batch and online learning.

(b) [8 marks] The US postal service uses neural networks to recognise hand-written digits in the area codes (zip codes) of mail addresses. For this (as well as other large training sets) it was observed that online learning can arrive at good weights substantially faster, in terms of passes through the training set, than batch learning. Suggest an explanation for this effect.

Question 6. Value-based RL [15 marks]

Q learning is a value-based approach to reinforcement learning, in which a lookup table is built up of Q values, namely the expected discounted future return following from carrying out each action in each state.

(a) [5 marks] Given a table of Q values, what is 'ϵ-greedy' action selection?

(b) [5 marks] Given a table of Q values, what is 'soft-max' action selection?

(c) [5 marks] Under what circumstances would it be possible to learn and use values of V, the expected value of the discounted future return from being in a given *state*, as opposed to Q?

Question 7. Function approximation in RL [15 marks]

Function approximation can be used in place of a simple lookup table of values in value-based reinforcement learning algorithms such as Q learning. Describe the advantages and disadvantages of this approach.

Question 8. Controlling model complexity [15 marks]

Suppose we are faced with a classification task: we need to classify new inputs into class A and class B (say), based on a fairly modest amount of known data in the form of input-output pairs.

(a) [7 marks] Describe what the terms 'under-fitting' and 'over-fitting' mean in this context.

(b) [8 marks] Outline TWO methods for using a validation set to control model complexity when training this classifier on the known data.

Question 9. Bayesian inference [15 marks]

In contrast to methods such as 'hold-out' validation, Bayesians claim to be able to control model complexity *without* sacrificing any of the data. Explain how they are able to do this. You may find it useful to refer to a particular case such as the setting of the amount of weight-decay in a neural network. You may also find it useful to invoke the Bayesian concept of 'evidence' - the probability of the data, integrated over the parameter space of the model class.

Question 10. Hidden Markov models [15 marks]

Consider taking the standard HMM and adding a second hidden variable, as shown in the figure below. In such a model of the world there are multiple underlying states that co-exist and combine to generate the observable data. Do you anticipate any problems generalising the forward-backward algorithm to this case? Explain your answer.

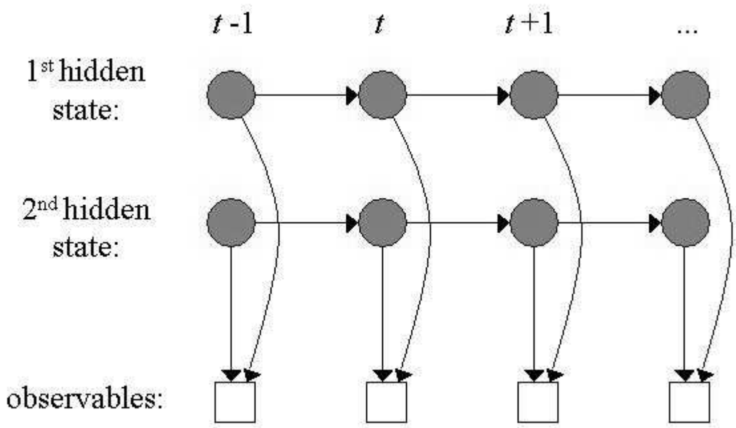

Question 11. Belief networks [15 marks]

Consider a large Belief Network for which most of the values are hidden (unknown). In which circumstances could "explaining away" inference be performed by ancestral simulation? (Include a description or sketch an example of 'explaining away' in your answer).

Question 12. Dynamic decision networks [15 marks]

(a) [6 marks] Describe the relationship between Dynamic Decision Networks and Belief Networks.

(b) [9 marks] How are sequences of actions chosen using Dynamic Decision Networks? (Note: you may like to consider action selection in 'non-dynamic' decision nets first).

****