Victoria University of Wellington
DEGREE EXAMINATIONS - 2001

COMP 421

MID YEAR

COMP 421

ARTIFICIAL INTELLIGENCE

Time allowed: THREE HOURS

Instructions: Attempt ALL questions.
There are 180 marks altogether.
Two sides of personal hand-written notes are permitted.

Sections, and marks included in them, are as follows:

Search and brains	30
Games and maximum likelihood	30
Reinforcement learning	30
Generalization and information theory	30
Hidden Markov models	30
Belief networks	30

Search, and brains. 30 marks

(1) In the ‘No Free Lunch’ theorem, search is treated as the problem of generating a finite number of good sample solutions from an intractably large space of possibilities. Explain what the theorem implies for search algorithms. [5 marks]

(2) As a search algorithm, naïve gradient descent is problematic due to its slowness on long shallow inclines, and oscillation on steep ravines. Explain how ‘momentum’ can be used to avoid both problems [5 marks]

(3) A drawback of using momentum in gradient descent is that the amount of momentum is a free parameter which needs to be set. Line search avoids this but has its own problem. Explain the main problem with line search, and how it is overcome by the conjugate gradient method. [5 marks]

(4) A drawback of all gradient descent methods is that they can become trapped in local minima. Explain how simulated annealing can be used to avoid this problem. [5 marks]

(5) The mammalian cerebral cortex can be loosely thought of as divided into four lobes: occipital, parietal, temporal and frontal. Make brief notes on likely effects of lesions (*i.e.* damage) to each of these regions. [10 marks]

Games and maximum likelihood **30 marks**

(6) Consider a three-person game with players *A*, *B* and *C*. In this game player *A* moves, then *B*, and then *C*. At the leaves of the minimax tree three numbers are kept for each board position, representing the value of the position from the three players' respective points of view. That is, (3 4 2) means a board position is worth 3 points to player *A*, 4 points to player *B*, and 2 points to player *C*. The game tree is shown below:

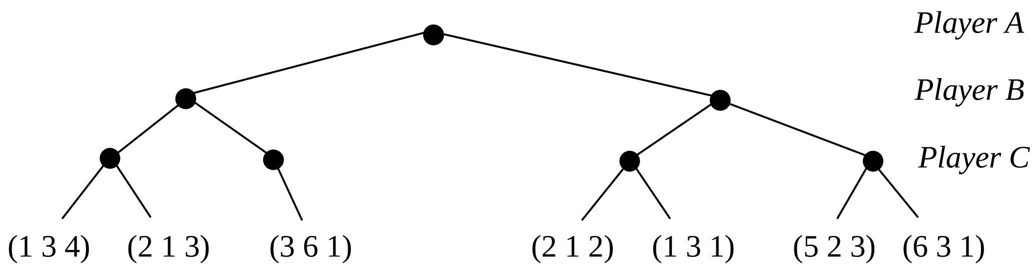

Show how a minimax-like algorithm could handle the 3-player case, by demonstrating it on this example. [10 marks]

(7) "The prisoner's dilemma is not a zero-sum game". Explain. [8 marks]

(8) By assuming one wants to minimize the sum of squared errors, derive the "delta rule" for altering the weights of a single model neuron with a linear activation function. [6 marks]

(9) Justify the above cost from a probabilistic perspective (maximum likelihood). [6 marks]

Reinforcement learning **30 marks**

(10) Consider a 1-dimensional discrete world consisting of a cycle of states 1 to 8, where each state is accessible only from its nearest neighbours, and state 1 is a direct neighbour of state 8. An agent learns in this world, and begins any given episode in a randomly chosen state. It gets a reward of 1 at state 5, which also signifies the end of the current episode. Everywhere else it gets no reward. The agent assumes a discount rate of γ per timestep for delayed rewards. Under an optimal policy, what is the (Q) value of going

- (i) from state 3 to state 4?
- (ii) from state 3 to state 2?
- (iii) from state 8 to state 1? [6 marks]

(11) Explain the role of "eligibilities" in SARSA(λ). [6 marks]

(12) Explain the role of "eligibilities" in Direct Reinforcement Learning. [6 marks]

(13) Describe an advantage **and** a disadvantage of Direct Reinforcement Learning over value-based reinforcement learning algorithms. [8 marks]

(14) What is "perceptual aliasing" in POMDPs (partially observable Markov decision problems)? [4 marks]

Generalization and information theory **30 marks**

(15) Consider a neural network that is being trained to correctly classify hand-written digits. “Early stopping” is one method for complexity control in neural networks. Describe exactly how you could use ten-fold cross-validation to determine an appropriate point to stop training the network. [8 marks]

(16) Bayesians claim to be able to control model complexity *without* sacrificing any of the data. Explain. [12 marks]

(17) “Entropy is average surprise”. Explain. [5 marks]

(18) Suppose we want to transmit a message using an alphabet, where P_i is the true probability of the i^{th} symbol, but we don’t know these probabilities and instead use distribution Q (that is, we believe the probability of the i^{th} symbol to be Q_i). The Kullback-Leibler divergence between probability distributions P and Q is $K = \sum_i P_i \log_2 \frac{P_i}{Q_i}$. Explain why the resulting message, encoded using an optimal coding scheme, will be K bits per symbol longer than it would be if we knew the true probabilities.

[5 marks]

Hidden Markov Models **30 marks**

For the next two questions:
A Hidden Markov Model (HMM) is used to model observable data $x_1, x_2, \ldots, x_T$. It consists of the following two sets of parameters: transition probabilities $a_{ij} = \Pr(s_{t+1} = j | s_t = i)$ between hidden states *s*, together with emission probabilities $e_{ik} = \Pr(x_t = k | s_t = i)$. Useful quantities to calculate are the forward probabilities $\mathrm{f}_i(t) = \Pr(s_t = i | x_1, x_2, \ldots, x_t)$ and backward probabilities $\mathrm{b}_i(t) = \Pr(x_{t+1}, x_{t+2} \ldots x_T | s_t = i)$, for each time step *t*.

(19) Express the posterior probability $\Pr(s_t = i, s_{t+1} = j | x_1, x_2, \ldots, x_T)$ in terms of the above. [6 marks]

(20) How are these probabilities $\Pr(s_t = i, s_{t+1} = j | x_1, x_2, \ldots, x_T)$ used to update parameters *a* in an HMM? [6 marks]

For the next two questions:
Suppose you are given a partially observable Markov prediction problem suitable for modeling with an HMM. Rather than trusting one set of parameters, you decide to take a more Bayesian approach and make predictions using a large number of HMMs, each with different transmission and emission probabilities chosen at random from a maximum entropy prior distribution.

(21) Each HMM makes probabilistic predictions for the next data item. How would you combine these predictions? [9 marks]

(22) Suppose you only have enough memory for one HMM, but have a lot of time in which to perform inference. How could you achieve the same effect as above, by making use of a Markov Chain Monte Carlo method? [9 marks]

Belief networks 30 marks

(23) Suppose that observables *A* and *B* are seen to be correlated but we are not sure whether *A* causes *B*, or *B* causes *A*, or both *A* and *B* share some unknown cause, *X*. Another variable, *C*, is now found to be correlated with *A* but not with *B*. Exactly what new light does this shed about the relationship between *A* and *B*?

[8 marks]

Suppose that your house has a burglar alarm. The alarm (*a*) usually goes off if you are burgled (*b*), but also goes off for some earthquakes (*e*). Your place of work sometimes sways (*s*) in earthquakes, although it often sways in strong winds (*w*) also.

(24) Draw a belief network that captures the structure of relations between *a*, *b*, *e*, *s*, and *w*. [2 marks]

(25) Write out the factorization of the joint probability $P(a,b,e,s,w)$ implied by this network. [2 marks]

(26) Draw the corresponding factor graph. [2 marks]

(27) Are variables *b* and *s* dependent or independent variables, if you know the alarm has gone off? [4 marks]

(28) Briefly describe how Gibbs sampling could be used to estimate the probability that you have been burgled, on a windy day in which you know the alarm has gone off and felt the building sway. [6 marks]

(29) Burglars prefer to work on windy days. Explain the problem this raises for running probability propagation on this network, and suggest ways around the problem that do not require resorting to approximations. [6 marks]

****