

Victoria University of Wellington
DEGREE EXAMINATIONS - 2000
MID YEAR

COMP 421

COMP 421 ARTIFICIAL INTELLIGENCE

Time allowed: **THREE HOURS**

Instructions: Attempt ALL questions.
 There are 100 marks altogether.
 Two sides of personal hand-written notes are permitted.

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

Search	20
Reinforcement learning	10
Maximum likelihood and cost functions	15
Bayesian inference and information	25
Belief networks	20
Hidden Markov models	10

Search**20 marks**

(1) Explain how "roulette wheel" selection is invariant to multiplication of all fitness values by an arbitrary positive constant. [3 marks]

(2) Explain how "Boltzmann" selection is invariant to addition of an arbitrary constant to all fitness values. [3 marks]

(3) What is "elitist" selection? [2 marks]

(4) The Metropolis-Hastings algorithm for sampling from a distribution p accepts a transition from state μ to state ν with probability

$$P_{\mu \rightarrow \nu} = \min \left(1, \frac{p_{\nu}}{p_{\mu}} \right)$$

Show that this system obeys detailed balance. [4 marks]

(5) Describe the relationship between Metropolis-Hastings sampling and the search algorithm, Simulated Annealing. [2 marks]

(6) Explain how each of the following work to accelerate convergence of gradient descent:

- i. momentum
- ii. line search
- iii. conjugate gradient

[6 marks]

Reinforcement learning**10 marks**

- (7) Briefly describe the dynamic programming technique of "value iteration".
[3 marks]
- (8) How does Q -learning differ from Sarsa?
[3 marks]
- (9) Describe the use of eligibilities in "direct" reinforcement learning (DRL).
[4 marks]

Maximum likelihood and cost functions**15 marks**

- (10) Show that Maximum Likelihood learning with a Gaussian noise model is equivalent to minimizing the mean squared error (that is, LMS) on a training set.
[5 marks]
- (11) Why is LMS an inappropriate cost function for use in learning "inverse problems"?
[2 marks]
- (12) What is over-fitting, and how is it avoided by adding terms to the usual cost function (*ie.* regularization)?
[4 marks]
- (13) Describe how cross-validation might be used to set the amount of weight decay.
[4 marks]

Bayesian inference and information**25 marks**

(14) Outline a hierarchy of capitulations from "ideal" Bayesian inference, mentioning the assumption being made with each capitulation. [6 marks]

(15) What is the corollary of regularization in the Bayesian view of inference? (cf. Question 12) [2 marks]

Suppose that an unknown variable can take one of n discrete states:

(16) What is the Maximum Entropy prior for the state probabilities? [1 mark]

(17) Show that this entropy is $\log_2 n$. [1 mark]

(18) "Information refers to how events change our uncertainty". Discuss, invoking the concept of entropy. [6 marks]

(19) The Kullback-Leiber (KL) divergence between a true probability distribution P , and an approximation to it Q , is

$$P\|Q = \sum_i p_i \log \frac{p_i}{q_i}$$

Interpret this measure in terms of entropies, and describe its main properties.


[6 marks]

(20) What is the corollary of regularization from the minimum description length (MDL) perspective? [3 marks]

Belief networks**20 marks**

- (21) In which circumstances could "explaining away" inference be performed by ancestral simulation in belief networks? [4 marks]
- (22) Describe the use of Gibbs sampling for inference in belief networks. [4 marks]

The remainder of this section concerns the following belief network:



- (23) Write out the joint probability as a product of local factors. [2 marks]
- (24) Assuming each variables can take one of THREE values, how many probability values need to be specified
- (a) in the full joint distribution? [1 marks]
- (b) in the belief net? [1 marks]
- (25) Are variables a and b dependent or independent, given the value of e ? [2 marks]
- (26) Are variables a and b dependent or independent, given the value of s ? [2 marks]
- (27) In what sense does inference by probability propagation in so-called "Bayesian" belief nets *fail* to be Bayesian? Very briefly suggest how this might be addressed. [4 marks]

Hidden Markov models**10 marks**

- (28) What is *perceptual aliasing* in POMDPs (partially observable Markov decision problems), and how is it dealt with in HMMs (hidden Markov models)? [5 marks]
- (29) Give a strength, and a weakness, of Hidden Markov models for modelling structure in time series. [2 marks]
- (30) Describe how actions are selected in dynamic decision networks. [3 marks]