UNIVERSITY OF MANITOBA

Winter 2003

COMPUTER SCIENCE

Machine Learning

Paper No.: **800**
Examiners: **Jacky Baltes**
Date: **24th April 2003**
Time: **18:00**
Room: **Frank Keneddy Center, Gold Gym**

(Time allowed: 180 Minutes)

NOTE: Attempt all questions.
This is a *closed* book examination.
Use of calculators is *permitted*.
Show your work to receive full marks.

SURNAME:

FORENAME(S):

STUDENT ID:

A	B	C	D	E	Total
20	20	20	20	20	100

Surname: _____

Forename(s): _____

Section A: Candidate Elimination

1. Given the hypothesis space below, the candidate elimination algorithm is trained on a sequence of four instances (Training data D).

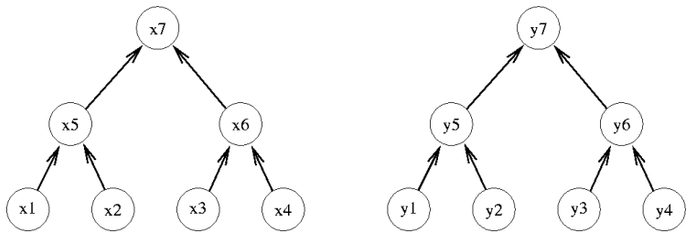

After training the candidate elimination algorithm returns the following version space V:

$$V = \begin{array}{ll} \text{S-Set:} & \texttt{<x2,y5>} \\ \text{G-Set:} & \texttt{<x7,y7>} \end{array}$$

Given the generalization hierarchy in question 1, show all concepts that are consistent with the training data (i.e. are included in the version space V).

[3 marks]

<x2,y5>
<x5,y5>
<x7,y5>
<x2,y7>
<x5,y7>
<x7,y7>

2. How many concepts consistent with V classify the instance <x1,y1> as positive and how many classify it as negative.

[2 marks]

<x1,y1> is classified as positive by 4 concepts and negative by 2 concepts.

3. Given the generalization hierarchy in question 1, what instance should the user classify next to reduce the number of questions by the system (i.e., the sample complexity).

   How many consistent concepts are classified as positive/negative respectively by the instance you selected.

[5 marks]

The system should ask next about the instance x1,y3. This instance is classified as positive by 3 concepts and as negative by 3 concepts.

Surname: _____

Forename(s): _____

4. Below is a trace of the candidate elimination algorithm.

Trace of Candidate Elimination Algorithm

```
x1      +       S-Set: <u2,v1>
                G-Set: <u7,v3>
x2      +       S-Set: <u5,v1>, <u6,v1>
                G-Set: <u7,v3>
x3      -       S-Set: <u5,v1>
                G-Set: <u5,v3>, <u7,v1>
x4      -       S-Set: <u5,v1>
                G-Set: <u5,v1>
```

The version space consists of two attributes. The set of instances is given by `x1`, `x2`, `x3`, and `x4`.

There are two attributes `U` and `V`. The generalization hierarchy for attribute `U` includes the symbols `u1`, `u2`, `u3`, `u4`, `u5`, `u6`, and `u7`. The generalization hierarchy for attribute `V` includes the symbols `v1`, `v2`, `v3`.

Show the conjunctive generalization hierarchy that would result in the trace shown above.

If no such conjunctive generalization hierarchy exists then say so and explain why.

[10 marks]

x1= x2= x3= x4=

Version Spaces U V

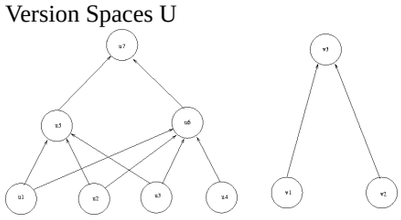

Section B: Decision Trees

The information gain Gain(S,A) of an attribute A for a sample set S is defined as

$$\text{Gain}(S, A) = \text{Entropy}(S) - \sum_{v \in \text{Values}(A)} \frac{|S_v|}{|S|} \text{Entropy}(S_v)$$

A graph of the entropy function is shown in Fig. 1 below. You can use this graph when answering the following questions.

5. Assume a domain with two attributes `A`, and `B`. Attribute `A` has two possible values (`a1` and `a2`); attribute `B` has three possible values (`b1`, `b2`, and `b3`).

Given below is a set of instances from the domain in question 1.

Calculate the information gain (Gain$(S, ?)$) for the attributes `A` and `B`. Which attribute would be selected by the standard ID3 algorithm?

If it is impossible to calculate the information gain from the given information, then specify so in your answer and explain why.

[10 marks]

Surname: _____

Forename(s): _____

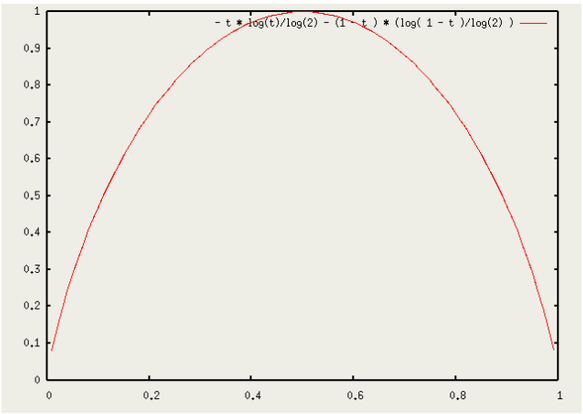

Figure 1: Graph of the Entropy Function

A	B	Target
a1	b1	No
a2	b1	Yes
a1	b2	No
a2	b2	No
a1	b1	Yes
a2	b3	Yes
a1	b3	No
a2	b2	Yes

Table 1: Training set for the A, B domain

Gain$(S, \mathrm{A}) = 0.18$
Gain$(S, \mathrm{B}) = 0.06$
ID3 would select the attribute A
e1 = myEntropy(4/8) = 1
eA1 = 4/8 * myEntropy(3/4) = 0.4056391
eA2 = 4/8 * myEntropy(3/4) = 0.4056391
g1 = e1 - eA1 - eA2 = 0.1887219
eB1 = 3/8 * myEntropy(2/3) = 0.3443609
eB2 = 3/8 * myEntropy(2/3) = 0.3443609
eB3 = 2/8 * myEntropy(0.5) = 0.25
g2 = e1 - eB1 - eB2 - eB3 = 0.0612781

6. Given the instances below, assign values for the attribute `B` which has three values (`b1`, `b2`, and `b3`), in such a way that the ID3 algorithm will select attribute `B` instead of attribute `A`.

   If it is impossible to assign values to attribute `B` given the constraints above, then say so in your answer and explain why.

   [10 marks]

A	B	Target
a1	b1	No
a2	b2	Yes
a1	b1	No
a2	b1	No
a1	b2	Yes
a2	b3	Yes
a1	b1	No
a2	b3	Yes

Surname: _____

Forename(s): _____

Section C: Bayesian Learning

7. Given the following data set (Table 2, what is the naive Bayesian classification of the new instance `<L,white>`.

Size	Color	Target 2
XS	green	Yes
L	green	Yes
XS	white	No
M	black	No
XL	green	Yes
XS	white	Yes
L	black	No
M	green	Yes

Table 2: Training set for the Size, Color domain

[5 marks]

```
P(Yes) = 5/8 = 0.625
P(No) = 3/8 0.375
P(L|Yes) = 1/5 = 0.2
P(L|No) = 1/3 = 0.33
P(white|Yes) = 1/5 = 0.2
P(white|No) = 1/3 = 0.33
P(yes) = P(Yes) * P(L|Yes) * P(White|Yes) = 5/8 * 1/5 * 1/5 = 0.025
P(no) = P(No) * P(L|No) * P(White|No) = 3/8 * 1/3 * 1/3 = 0.041
```

8. You are given the following information about a new computer. The name of the associated random variables is given in brackets.

   - (Power) The power in the building is on with 40% probability.
   - (Switch) The power switch is in the on position with 80% probability.
   - (Light) If the power switch is on and there is power in the building, then the probability of the light in the front panel of the computer being on is 90%. If the power in the building is off or the power switch is off, then the probability of the light in the front panel being on is 0%.
   - (Main-board) If the power switch is on and there is power in the building, then the probability of the main-board having power is 80%. If the power in the building is off or the power switch is off, then the probability of the light in the front panel being on is 0%.
   - (Hard-disk) The hard-disk is working correctly with 50% probability.
   - (Boot) If the main board has power and the hard disk is working correctly, then the computer will boot successfully with 70% probability. If the main-board is connected and the hard disk is not working correctly, then the computer will boot with 30% probability. In all other cases, the computer will not boot.

   Describe this information in the form of a bayesian belief network. Show the graph of the Bayesian belief network as well as all the conditional probabilities for all random variables.

   [10 marks]

Surname: _____

Forename(s): _____

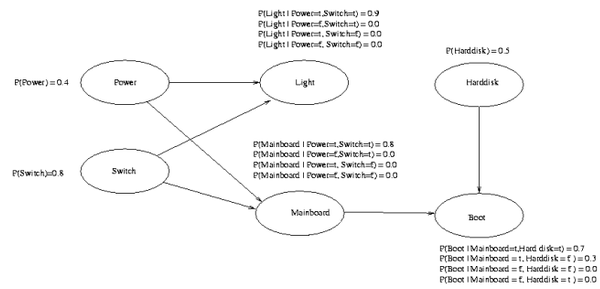

9. What is the probability of your computer booting successfully when you turn on the power switch? In other words, calculate $P(\text{Boot} = t|\text{Switch} = t)$.

[2 marks]

$P(Boot|Switch) =$
$P(Power) * P(MB|Power = t, Switch = t) = 0.24$
$P(HD) * P(Boot|MB = t, HD = t) * 0.24 = 0.5 * 0.7 * 0.24 = 0.0.084$
$P(\ HD) * P(Boot|MB = t, HD = f) * 0.24 = 0.5 * 0.3 * 0.24 = 0.036$
$P(Boot|Switch) = 0.084 + 0.036 = 0.12$

10. Assume that the computer does not boot, what is most likely explanation (Power is off, Switch is off, or the Hard disk is not working) for this failue assuming that only one part is broken. In other words, calculate the maximum of $P(\text{Power=f, Switch=t, Harddisk=t}|\text{Boot=f})$, $P(\text{Power=t, Switch=f, Harddisk=t}|\text{Boot=f})$, and $P(\text{Power=t, Switch=t, Harddisk=f}|\text{Boot=f})$.

[3 marks]

Surname: _____

Forename(s): _____

Section D: Reinforcement Learning

11. Zaku is a mobile robot that has learned the following Q-table for the 4 by 3 environment. B represents an obstacle and X represents the goal.

```
  00   00   00   00
 ___  ___  BBB  XXX
0___01___00BBB00XXX0
0___00___00BBB00XXX0
  13   00   00   00
  10   00   00   75
 ___  BBB  BBB  ___
0___00BBB00BBB00___0
0___00BBB00BBB00___0
  17   00   00   42
  13   00   00   56
 ___  ___  ___  ___
0___21___32___43___0
0___37___13___21___0
  00   00   00   00
```

There are two types of robots.

- Robots of type A have a global positioning system which allows them to determine which square they are in at any given moment in time.
- Robots of type B only have a local positioning system which allows them to sense whether or not there are obstacles to the north, east, south and west of the robot's current position.

Which type of robot (A or B) is Zaku? Explain your answer?

If the Q-table does not include sufficient information for you to identify the type of robot, then say so and explain why.

[4 marks]

12. Below is a Q-table learned by robot 5DPO. 5DPO has a global positioning system (i.e., it can accurately determine which square it is in) and highly accurate actuators (i.e. it moves exactly one square in any direction).

```
  00   00   00   00   00   00
 ___  ___  ___  ___  BBB  XXX
0___07___08___00___00BBB00XXX0
0___00___00___00___00BBB00XXX0
  06   00   00   00   00   00
  00   00   00   00   00   99
 ___  BBB  BBB  BBB  BBB  ___
0___00BBB00BBB00BBB00BBB00___0
0___00BBB00BBB00BBB00BBB00___0
  00   00   00   00   00   00
  00   00   00   01   00   41
 ___  ___  ___  ___  ___  ___
```

Surname: _____

Forename(s): _____

```
0___20___32___40___50___40___0
0___30___43___31___30___80___0
  00   00   00   00   00   00
```

How would the learned Q-table change if 5DPO's actuator control fails. After the failure 5DPO moves by one 1 square 66% of the time, but with probability 33% it will move two squares unless it is stopped by an obstacle.

Show the new Q-values for 5DPO after its actuator control has failed.

[8 marks]

13. You apply the TD(λ) algorithm to a navigation problem. Not sure what values to select for the parameters, you read on the Internet that $\alpha = 0.1, \lambda = 0.3, \gamma = 0.9$ are reasonable values.

    The navigation task has the following features:

    - Your environment is modeled after an office domain, with long corridors connecting a number of different rooms.
    - Your robot only receives a reward (+100) if it moves into the goal square. All other transitions (even if the robot attempts to move onto a blocked square) receive an immediate reward of 0.

    You use an ϵ-greedy action selection policy and you notice that your robot covers all parts of the domain.

    In spite of this, your robot is unable to learn the correct Q-values to navigate the environment. Most of the Q-values are still 0. Which parameter would you change and why?

    [3 marks]

    I would increase/decrease parameter λ because the robot will only rarely reach the goal location and it seems as if the corresponding rewards are not propagated to the previous states.

14. Given below is the elgibility table for an agent. The information shows the current elgibility value for the states. For example, the eligbility of action A_1 in state S_0 is 0.25.

    The reinforcement learner used the following parameters: $\alpha = 0.1$, $\gamma = 0.9$, $\epsilon = 0.2$, and $\lambda = 0.7$.

Surname: _____

Forename(s): _____

S_0	A_0	0.78	S_4	A_0	0.00
	A_1	0.25		A_1	0.00
	A_2	0.00		A_2	0.00
	A_3	0.00		A_3	0.00
S_1	A_0	0.00	S_5	A_0	0.00
	A_1	0.15		A_1	0.00
	A_2	0.00		A_2	0.39
	A_3	0.00		A_3	0.00
S_2	A_0	0.00	S_6	A_0	0.00
	A_1	0.00		A_1	0.00
	A_2	0.00		A_2	0.00
	A_3	0.00		A_3	0.00
S_3	A_0	0.00	S_7	A_0	0.00
	A_1	0.00		A_1	1.00
	A_2	0.00		A_2	0.00
	A_3	0.00		A_3	0.00

The robot started in state S_0 and executed A_0 first.

[5 marks]

S0/A0 -¿ A1/A1 -¿ S0/A1 -¿ S5/A2 -¿ S0/A0 -¿ S7/A1

Section E: Instance-Based Learning

15. Given a domain with two attributes `M` and `N`. The domain for each attribute are integers in the interval 0 to 10.

    Assume that an instance based learning algorithm uses the *Manhattan distance* to compute the similarity between two instances. The Manhattan distance between two instances $I_1 =$<n_1,m_1> and $I_2 =$<n_2,m_2> is defined as

    $$d_{Man}(I_1, I_2) = abs(n_1 - n_2) + abs(m_1 - m_2)$$

    For example, the Manhattan distance between <6,3> and <4,4> is 3.

    Show the estimate of the target function of the `distance weighted 3-Nearest Neighbor` algorithm. The target function estimate $f'(x_q)$ of the distance weighted 3-Nearest Neighbor algorithm where C_1, C_2, and C_3 are the three nearest neighbors is computed as follows:

$$f'(x_q) = \begin{cases} \dfrac{\sum_{i=1}^{3} f(C_i) K(d_{Man}(x_q, C_i))}{\sum_{i=1}^{3} K(d_{Man}(x_q, C_i))} & \text{if } \sum_{i=1}^{3} K(d_{Man}(x_q, C_i)) \ != 0 \\ 0 & \text{else} \end{cases}$$

Surname: _____

Forename(s): _____

Use the kernel function K

$$K(d) = \begin{cases} 2 & \text{if } d = 0 \\ 1 & \text{if } d = 1 \\ 0.5 & \text{if } d = 2 \\ 0.25 & \text{if } d = 3 \\ 0.1 & \text{if } d = 4 \\ 0 & \text{else} \end{cases}$$

The training data consists of the instances C_1 =<E,4>, C_2 =<H,1>, and C_3 =<E,7>. The target function of the three instances are $f(C_1) = 4$, $f(C_2) = 1$, and $f(C_3) = 8$.

[10 marks]

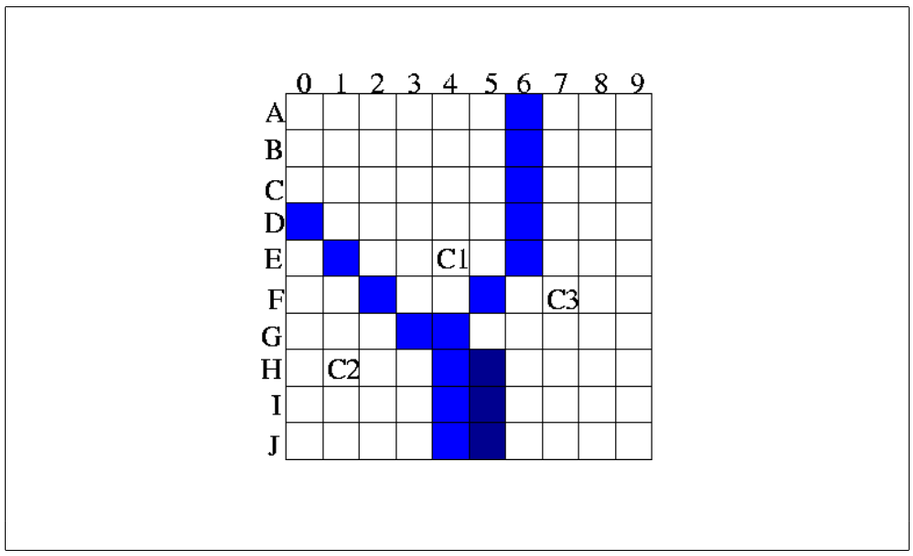

16. Given below is the target function for an instance based learning algorithm. The classification uses the 1-nearest neighbor algorithm with the Manhattan distance as distance metric.

    If a square is equi-distant to two or more cases, then a classification is calculated as the average of all equi-distant cases.

	0	1	2	3	4	5	6	7	8
A	1	1	4	4	4	4	4	8	8
B	1	1	4	4	4	4	4	8	8
C	1	1	4	4	4	4	4	8	8
D	1	1	4	4	4	4	4	8	8
E	1	1	4	4	4	4	4	8	8
F	1	1	4	4	4	4	4	8	8

What is the minimum number of cases necessary to represent the target function. Give a set of cases that can represent the target function correctly.

If it is impossible to find a set of cases to represent the target function, then say so and explain why in your answer.

[10 marks]

Surname: _____

Forename(s): _____

The minimum number of cases to represent the target function correctly is 4.

Cases: C_1 =<A,0>, C_2 =<A,4>, C_3 =<A,5>, C_4 =<A,8>

Surname: _____

Forename(s): _____

Additional work pages