
UNIVERSITY OF MANITOBA

Final Examination

Winter 2005

COMPUTER SCIENCE

Machine Learning

Paper No.: 444
Examiners: Jacky Baltes
Date: 18 April 2006
Time: 9:00
Room: University College Great Hall (25 - 48)

(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

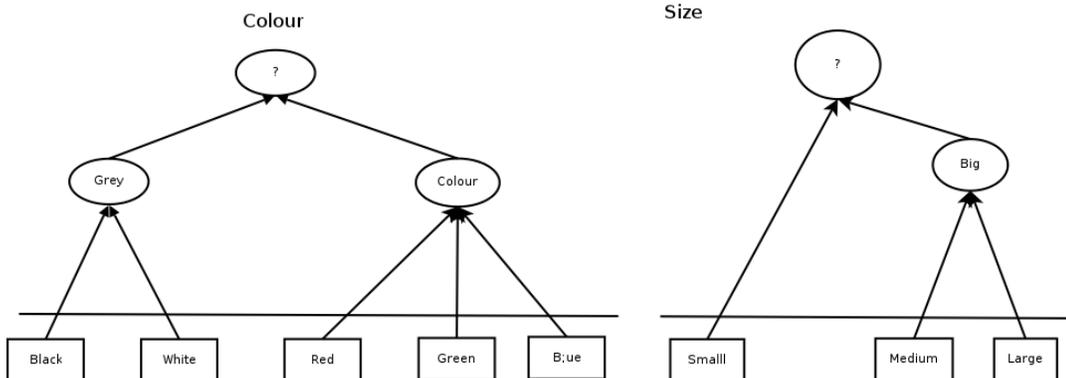
CONTINUED

Surname: _____

Forename(s): _____

Section A: Candidate Elimination

1. Given the generalization hierarchy H shown below. The target concept is $\langle \text{Colour}, \text{Medium} \rangle$.



Show the minimum training set such that the candidate elimination algorithm will learn the target concept. For each sample in the training set show the classification, the resulting S -set and G -set. One entry of the training set is already given in the answer box below.

If it is impossible to determine a minimum training set D such that the candidate elimination algorithm is able to learn the concept $\langle \text{Colour}, \text{Medium} \rangle$, then say so in your answer and explain why.

[7 marks]

The minimum size of the training set is _____ samples.

$D =$	Instance	Classification	S/G-set
	$\langle \text{Red}, \text{Medium} \rangle$	+	S -set: $\langle \text{Red}, \text{Medium} \rangle$ G -set: $\langle ?, ? \rangle$
	$\langle \text{Blue}, \text{Medium} \rangle$	+	S -set: $\langle \text{Colour}, \text{Medium} \rangle$ G -set: $\langle ?, ? \rangle$
	$\langle \text{Black}, \text{Medium} \rangle$	+	S -set: $\langle \text{Colour}, \text{Medium} \rangle$ G -set: $\langle \text{Colour}, ? \rangle$
	$\langle \text{Red}, \text{Large} \rangle$	+	S -set: $\langle \text{Colour}, \text{Medium} \rangle$ G -set: $\langle \text{Colour}, \text{Medium} \rangle$

CONTINUED

Surname: _____

Forename(s): _____

2. Given the generalization hierarchy H in question 1, assume that after a training sequence D , the candidate elimination algorithm returns the following version space V :

S-set : <Red, Large>
 G-set : <Colour, ?>, <?, Big>

List all remaining concepts that are still consistent with the training data.

[7 marks]

There are 8 concepts remaining that are consistent with the training data.

- 1 <Red, Large>
- 2 <Colour, Large>
- 3 <?, Large>
- 4 <Red, Big>
- 5 <Red, ?>
- 6 <Colour, Big>
- 7 <Colour, ?>
- 8 <?, Big>

3. Programmer Joe shows you the output of his implementation of the Candidate Elimination algorithm after training it with a sample sequence D of five instances.

Output of Joe's program: S-set : <Colour, Small>, <Colour, Medium>
 G-set : <?, Big>, <Black, Medium>, <Colour, Big>

In spite of the fact that you do not know the training sequence D , you tell Joe that there must be a bug in his program.

Show all errors and explain why they can not possibly be the result of training the candidate elimination algorithm on the hypothesis space H .

If the output of Joe's program is the correct output of some training set D then say so in your answer and show the training set D that would result in the output of Joe's program.

[6 marks]

- Error 1: S-set must include both <Colour, Small> and <Colour, medium>
 Error 2: <Black, Medium> must cover at least one element of the S -set.
 Error 3: <Colour, Big> is more specific than <?, Big> and only the most general items must be in the G -set.

Surname: _____

Forename(s): _____

Section B: Decision Trees

The information gain $\text{Gain}(S,A)$ 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.

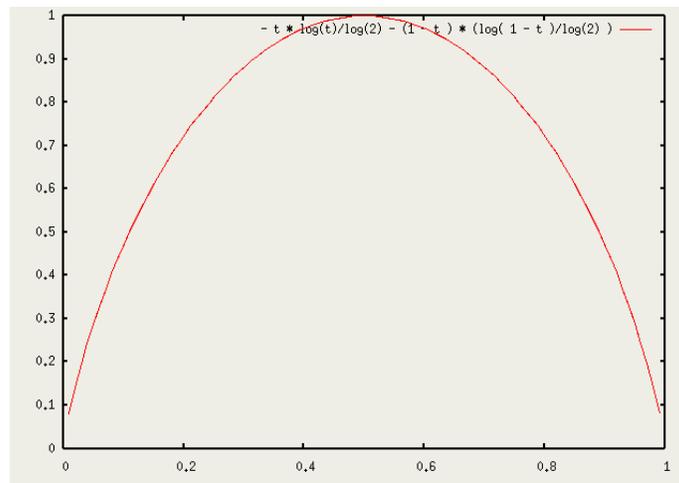


Figure 1: Graph of the Entropy Function

4. Assume a domain with three attributes A, B, and C. Each attribute has two possible values T and F. Given below is a set of instances.

A	B	C	Target
T	T	T	Yes
F	F	T	No
F	T	F	No
T	F	T	No
T	F	F	No

Calculate the information gain ($\text{Gain}(S, ?)$) for the attributes A, B, and C. 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.

[7 marks]

CONTINUED

Surname: _____

Forename(s): _____

$$\text{Gain}(S, A) = -0.08$$

$$\text{Gain}(S, B) = 0.07$$

$$\text{Gain}(S, C) = -0.08$$

ID3 would select any attribute

$$\text{ent}(2,3)=0.97$$

$$\text{IG}(A)=\text{ent}(1,4) - 3/5 * \text{ent}(1,2) - 2/5 * \text{ent}(0,2) = 0.17$$

$$\text{IG}(B)=\text{ent}(1,4) - 2/5 * \text{ent}(1,1) - 3/5 * \text{ent}(3,0) = 0.32$$

$$\text{IG}(C)=\text{ent}(1,4) - 3/5 * \text{ent}(1,2) - 2/5 * \text{ent}(0,2) = 0.17$$

5. A colleague ran an experiment and learned an important decision tree from a sample set. You want to repeat your colleague's experiment for independent verification. Unfortunately, your colleague is on a mountaineering trip in the Himalayas and can not be reached.

From your colleague's notes it is clear that:

- The initial entropy of the dataset was 1.0.
- The domain included two attributes Q and R with three and two values respectively.
- The decision tree has four internal nodes (that is nodes in the tree not counting the actual classification).
- The decision tree classified all instances in the dataset correctly.

You know that the training set was one of the following

Training set	positive instances	negative instances	Information Gain Q	Information Gain R
Set 1	3	6	0.3	0.1
Set 2	5	5	0.5	0.4
Set 3	6	6	0.3	0.4
Set 4	2	3	0.7	0.3
Set 5	4	2	0.5	0.6
Set 6	3	4	0.6	0.0
Set 7	5	4	0.3	-0.1
Set 8	3	3	0.7	0.4

If you can determine uniquely which dataset was used by your colleague then say so in your answer and explain why. If there is more than one possible candidate then show all possible candidates and explain your answer. If it is impossible that the decision tree was learned with any of these training sets then say so in your answer and explain why.

[8 marks]

The decision tree was learned with dataset _____

There are two possibilities, either Set 2 or Set 8, since both of them have an initial entropy of 1.0 and the IG of Q is higher than that of R resulting in a four node tree.

CONTINUED

Surname: _____

Forename(s): _____

6. This question uses a domain with three attributes X, Y, Z , that all have two possible values (T and F). So there are at most $2^8 = 256$ (the power set of 8 different concepts in the domain).

How many syntactically different **decision trees with exactly three internal nodes** are possible in this domain? Do **not** count the actual classification (T or F) as part of the tree. You can also ignore the fact that syntactically different decision trees may result in semantically the same concept.

[5 marks]

The are at most $6*8+3*4=60$ syntactically different decision trees with exactly three nodes in this domain.

Surname: _____

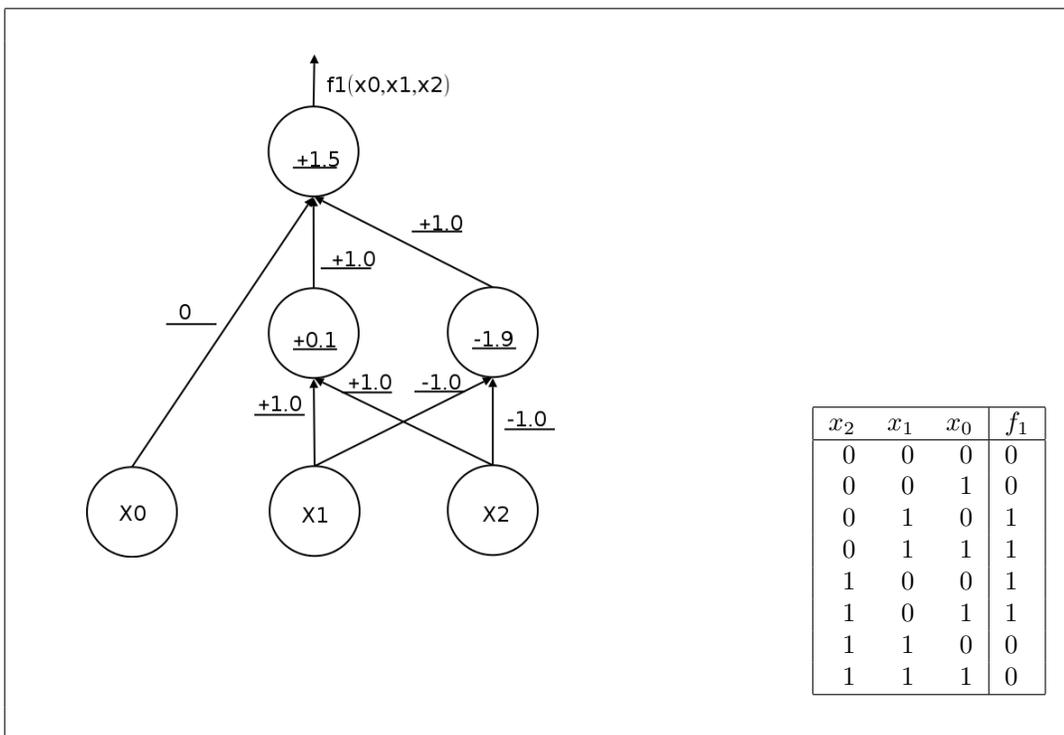
Forename(s): _____

Section C: Neural Nets

7. Given below is an artificial neural network (ANN) with three input nodes (X0,X1,X2), two hidden nodes, and one output node. The network uses **simple threshold nodes** (i.e., the node will output 1.0 if the sum of the weighted inputs is greater than or equal to the threshold, 0 otherwise).

Show a set of weights and thresholds for all nodes that implement the boolean function f_1 . If it is impossible to represent the boolean function f_1 with the given neural network, then state this in your answer and explain why this is impossible.

[10 marks]



The output y of a sigmoid activation unit j is given by

$$a_i = \sum_{j=0}^n w_{ij}x_j$$

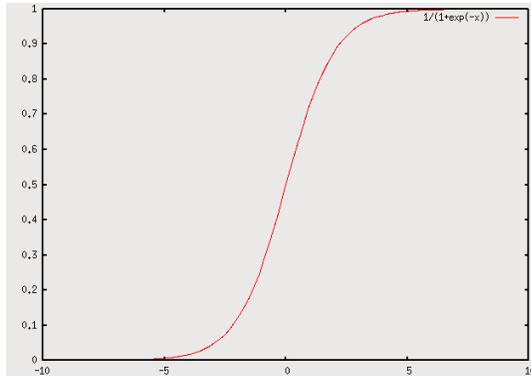
$$y_j = \frac{1}{1+e^{-a_j}}$$

The following graph shows the activation function for sigmoid activation units.

CONTINUED

Surname: _____

Forename(s): _____



The backpropagation algorithm shown below is one of the most popular training algorithms for artificial neural networks (ANNs) with sigmoid activation functions.

Backpropagation(Output nodes y , Hidden nodes h , Weights w_{ij} , Training Set D)

$\forall w_{ij} :=$ Initialize to small random value

do Until the termination condition has been met

do $\forall \{ \langle x_1, x_2, \dots, x_n \rangle, \langle t_1, t_2, \dots, t_k \rangle \} \in D$

Apply $\langle x_1, x_2, \dots, x_n \rangle$ to the network and compute outputs y_1, \dots, y_k

do $\forall y \in$ Output nodes

$$\delta_y = y(1 - y)(t - y)$$

od

do $\forall h \in$ Hidden nodes

$$\delta_h = y_h(1 - y_h) \sum_k w_{hk} \delta_k$$

od

do $\forall w_{ij} \in$ Weights

$$w_{ij} = w_{ij} + \alpha \delta_j x_{ij}$$

od

od

od

8. Given below is a small neural network with four **sigmoid**

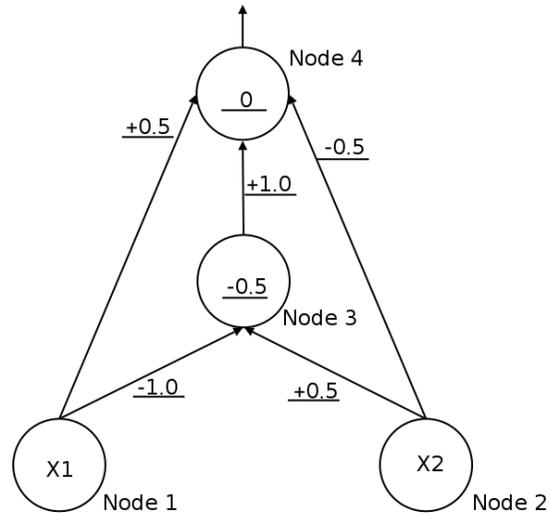
activation units: one output node, one hidden node, and two input

nodes. The current weights and thresholds of the network are shown in the figure. Note that thresholds are implemented as weights to a constant -1 input.

CONTINUED

Surname: _____

Forename(s): _____



After applying a training instance $\langle x, t \rangle$, the backpropagation algorithm is run with a learning rate of α and some of the weights are updated.

Given the information in the answerbox below, compute the target value for the training instance $\langle x, t \rangle$, so that the application will result in the described weight change(s).

Learning rate	α	0.5
Training instance	$\langle x_1, x_2 \rangle$	$\langle 0, 1 \rangle$
Weight changes	Weight Node 3 to Node 4	changed to 0.975

[10 marks]

The target value t for the given training instance was 0.1

```

Activation of node 3:  $a_3 = 1/(1+\exp(-1)) = 0.73$ 
Activation of node 4:  $a_4 = 1/(1+\exp(0.23)) = 0.44$ 
 $dw_{43} = 0.025 = \alpha * \delta_4 * x_{43}$ 
 $\Rightarrow \delta_4 = 0.025/(0.5*0.73) = 0.07$ 
 $\delta_4 = 0.07 = 0.44*(1-0.44)(t - 0.44)$ 
 $0.27 = t - 0.44$ 
 $t = 0.71$ 
  
```

CONTINUED

Surname: _____

Forename(s): _____

Section D: Bayesian Learning

9. Given the following data set, what is the naive Bayesian classification of the new instance $\langle M, \text{black} \rangle$. Show your work for full marks.

Size	Color	Target
S	black	No
M	white	Yes
L	green	No
S	black	No
L	black	Yes
M	green	Yes
S	white	No
M	black	No
S	green	No

Table 1: Training set for the Size, Color domain

[5 marks]

Naive Bayesian classification of $\langle M, \text{black} \rangle$ is: Yes

$$P(\text{Yes}) = 1/3$$

$$P(\text{No}) = 2/3$$

$$P(M|\text{Yes}) = 2/3$$

$$P(M|\text{No}) = 1/6$$

$$P(\text{black}|\text{Yes}) = 1/3$$

$$P(\text{black}|\text{No}) = 1/2$$

$$P(\text{yes}|M \ \& \ \text{black}) = P(\text{Yes}) * P(M|\text{Yes}) * P(\text{black}|\text{Yes}) = 1/3 * 2/3 * 1/3 = 2/27 = 0.07$$

$$P(\text{no}|M \ \& \ \text{black}) = P(\text{No}) * P(M|\text{No}) * P(\text{black}|\text{No}) = 2/3 * 1/6 * 1/2 = 1/18 = 0.05$$

-> Yes

CONTINUED

Surname: _____

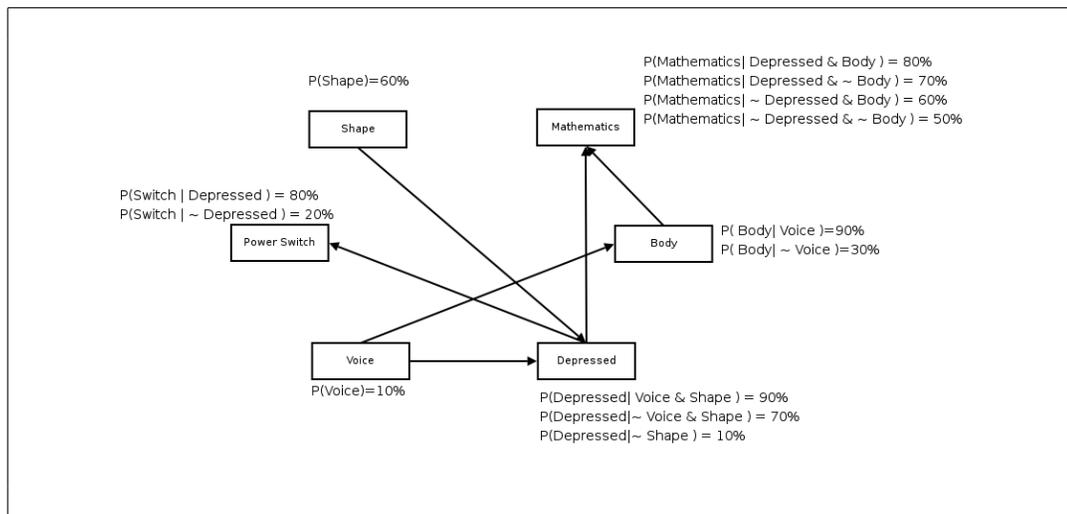
Forename(s): _____

10. You are given the following information about the robot domain. The name of the associated random variables is given in brackets.

- (Shape) The robot is a humanoid (60%) or non-humanoid (40%) robot.
- (Voice) With 30% probability, the robot will have a metallic robot voice. With 70% probability, the robot has a human-like voice.
- (Depressed) The probability that a humanoid robot with a metallic voice is depressed is 90%. The probability that a humanoid robot with a normal voice is 70%. All non-humanoid robots are depressed with a probability of 10%.
- (Body) The probability of a robot with a metallic voice having a metallic body is 90%. The probability of a robot with non-metallic voice having a metallic body is 30%.
- (Mathematics) A depressed robot with a metallic body has an 80% probability of being good in Mathematics. A depressed robot with a non-metallic body has a 70% probability of being good in Mathematics. A non-depressed robot with a metallic body has a 60% probability of being good in Mathematics. A non-depressed robot with a non-metallic body has a 50% probability of being good in math.
- (Switch) A depressed robot has a 80% probability of having a power switch. A non-depressed robot has an 20% chance of having a power switch.

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.

[5 marks]



Surname: _____

Forename(s): _____

11. Calculate $P(\text{Mathematics}|\text{Voice \& Shape})$, that is the probability that a robot is good in mathematics given that it has a humanoid shape and robotic voice.

[5 marks]

$$P(\text{Mathematics}|\text{Voice \& Shape}) = 77\%$$

$$P(\text{Depressed}|\text{Voice \& Shape}) = 90\% \text{ \# Given voice and shape}$$

$$P(\text{Body}|\text{Voice}) = 90\% \text{ \# Given voice}$$

$$P(\text{Mathematics}|\text{Voice \& Shape}) =$$

$$P(\text{Math}|\text{Dep \& Body \& Voice \& Shape}) * P(\text{Dep}|\text{Voice \& Shape}) * P(\text{Body}|\text{Voice \& Shape}) +$$

$$P(\text{Math}|\text{Dep \& } \sim \text{Body \& Voice \& Shape}) * P(\text{Dep}|\text{Voice \& Shape}) * P(\sim \text{Body}|\text{Voice \& Shape}) +$$

$$P(\text{Math}|\sim \text{Dep \& Body \& Voice \& Shape}) * P(\sim \text{Dep}|\text{Voice \& Shape}) * P(\text{Body}|\text{Voice \& Shape}) +$$

$$P(\text{Math}|\sim \text{Dep \& } \sim \text{Body \& Voice \& Shape}) * P(\sim \text{Dep}|\text{Voice \& Shape}) * P(\sim \text{Body}|\text{Voice \& Shape})$$

=

$$0.8 * 0.9 * 0.9 +$$

$$0.7 * 0.9 * 0.1 +$$

$$0.6 * 0.1 * 0.9 +$$

$$0.5 * 0.1 * 0.1 = 0.77$$

12. Without any other information, what is the probability of a robot having a metallic body. In other words calculate $P(\text{Body})$.

[5 marks]

$$P(\text{Body}) = 36\%$$

$$P(\text{Body}) = P(\text{Body}|\text{Voice}) * P(\text{Voice}) + P(\text{Body}|\sim \text{Voice}) * P(\sim \text{Voice}) =$$

$$0.9 * 0.1 + 0.3 * 0.9 = 0.36$$

13. Calculate the probability that a robot has a humanoid shape, given that a robot has a power switch, but does not have a metallic body, that is calculate the following probability:

[5 marks]

$$P(\text{Shape}|\text{Switch} \& \sim \text{Body}) =$$

$$0.1956/0.25384 = 77\%$$

CONTINUED

Surname: _____

Forename(s): _____

Section E: Reinforcement Learning

The online TD(λ) reinforcement learning algorithm using an ϵ -greedy policy is shown below.

```

 $\forall s, a \quad Q(s, a) := \text{Initialize to random value}$ 
do Repeat for each episode
   $\forall s, a \quad e(s, a) := 0$ 
   $s := \text{Start state}, a := \text{First action}$ 
  do Repeat for each state in the episode
    Take action  $a$ , observe  $r$ , and  $s'$ 
    Choose  $a'$  from  $s'$  using  $\epsilon$ -greedy policy
     $\delta := r + \gamma Q(s', a') - Q(s, a)$ 
     $e(s, a) := e(s, a) + \delta$ 
    do  $\forall s, a$ 
       $Q(s, a) := Q(s, a) + \alpha \delta e(s, a)$ 
       $e(s, a) := \lambda \gamma e(s, a)$ 
    od
   $s := s', a := a'$ 
od

```

14. Below is a figure for a mobile robot environment. The robot T1 can move North, South, East and West only. The robot has a global positioning system; in other words, it knows which state $A0..D3$, it is in.

The robot also includes a precise sonar sensor that can sense obstacles in all four directions. So it will never try to move into a target square occupied by an obstacle.

The actuators of the robot are very accurate; the robot will always move exactly one square, unless the target square is occupied by an obstacle, in which case the robot will not move.

Once the robot reaches $A3$, the episode is terminated.

Some of the cells have intermediate rewards associated with them (These appear as circles in the cell). A robot receives this reward for any state-action pair that results in the robot moving into the square associated with the reward. For example, the reward for moving into square $B0$ is -5 .

The values in the boxes beside the actions are the Q-values for performing this action in the respective state.

The action selection policy π is the optimal policy, that is the robot will always choose the action with the best Q-value from the current state.

The discount factor γ and some of the reward values are missing. Determine the discount factor γ and the missing reward values.

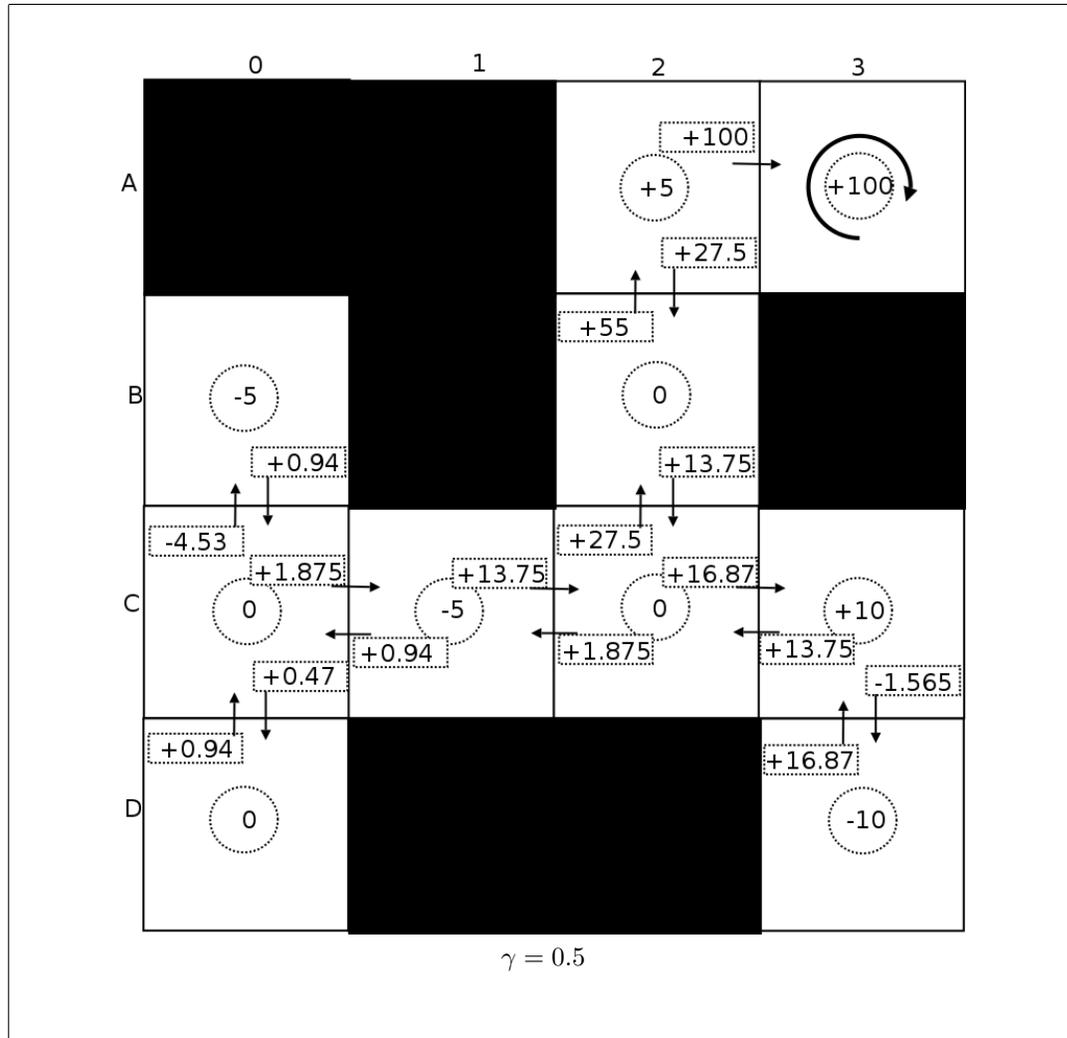
If it is impossible to determine the missing values then say so in your answer and explain why.

[5 marks]

CONTINUED

Surname: _____

Forename(s): _____



CONTINUED

Surname: _____

Forename(s): _____

15. Assume that robot T1's (as described in question 14) global positioning system fails. Now the only feedback that the damaged T1 receives about its current state is from the sonar array. It can only sense obstacles (or the edge of the playing field) to the north, south, east, or west.

Assume that the damaged robot now learns the Q-values shown for the problem in question 14. Which Q-values will be different from the ones learned in question 14 (i.e., when T2's global positioning system was working). Show the new Q-values, assuming that T2 will visit each state with uniform probability.

[5 marks]

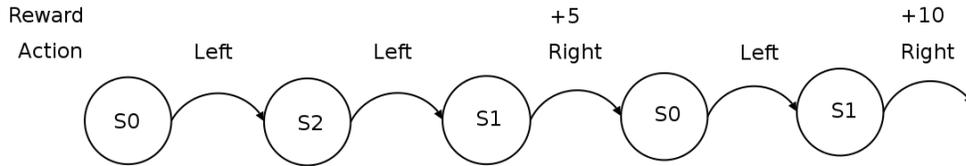
The following Q values will be different from the ones in question 14: $Q(D0,D3)$. $Q(D0,N)$
 $= Q(D3,N) = (16.87 + 0.94) / 2 = 8.905$.

Surname: _____

Forename(s): _____

16. Below is a trace of the TD(λ) algorithm for a short episode. Show the Q-table and eligibility values for each state action pair after the robot reaches the end of the sequence. The TD(λ) algorithm uses the following parameters: $\alpha = 0.3$, $\lambda = 0.9$, $\gamma = 0.8$.

If it is impossible to determine the Q values and eligibilities, then say so in your answer and explain why.



[10 marks]

Q-Table

$Q(s, a)$	S0	S1	S2
Left	0.951	0.00	1.15
Right	0.00	1.59	0.0

Eligibility

$e(s, a)$	S0	S1	S2
Left	0.99	0.0	0.37
Right	0.000	0.52	0.0

Surname: _____

Forename(s): _____

Additional work pages

Surname: _____

Forename(s): _____

Additional work pages
