University of Wisconsin-Madison
Computer Sciences Department

CS 760 — Machine Learning

Spring 1994

Midterm Exam

(calculator and one page of notes allowed)

100 points, 90 minutes

April 25, 1994

Write your answers on these pages and show your work. If you feel that a question is not fully specified, state any assumptions you need to make in order to solve the problem. You may use the backs of these sheets for scratch work. Notice that all questions do not have the same point-value. Divide your time appropriately.

Before starting, write your name on this and all other pages of this exam. Also, make sure your exam contains four (4) problems on seven (7) pages.

Problem	Score	Max Score
1	_____	30
2	_____	25
3	_____	30
4	_____	15
Total	_____	100

1. Decision Trees (30 points)

i) Assume you are given the following three nominal features with the possible values shown.

```
F1   ∈  {v1, v2}
F2   ∈  {v3, v4, v5, v6}
F3   ∈  {v7, v8, v9}
```

Using Quinlan's ID3 and its max-gain formula, produce a decision tree that accounts for the following classified examples. ***Show all your work.***

```
F1 = v1    F2 = v3    F3 = v8    +
F1 = v1    F2 = v5    F3 = v9    -
F1 = v2    F2 = v4    F3 = v9    -
F1 = v2    F2 = v3    F3 = v8    -
F1 = v2    F2 = v6    F3 = v9    +
```

ii) Propose and justify a GA-based approach for learning decision trees. *Unlike* in the lecture notes, where entities in the population were represented as trees, use *fixed-length vectors* to represent your trees. For simplicity and concreteness, you may assume that you have *three* binary-valued features.

iii) Discuss *one* strength and *one* weaknesses of your ‘‘fixed-length vectors’’ approach for learning decision trees compared to:

a) representing the entities as arbitrary-sized trees (e.g., linked lists)

b) Quinlan’s ID3 approach

2. Artificial Neural Networks (25 points)

Consider using Bayes' rule to derive a function, which will be optimized by backpropagation, that measures the quality of a point in a given neural network's weight space. For simplicity, assume that there is only one output unit for the task at hand and that the correct output is either 0 or 1 (e.g., the task might be one of recognizing members of a category).

When the network outputs x ($x \in [0,1]$), interpret this as *the square root of* the probability that the correct output is 1.

Assume that the prior probability of a point W in weight space is given by

$$P(W) = \alpha \prod_{w \in W} \frac{1}{w^2}$$

where α is a normalization constant. (To be technically correct, we will assume that weight space is discrete and finite; hence, you need not deal with probability densities. Also, you may assume that no weight exactly equals zero.)

Show the function that results from applying Bayes' rule under these assumptions. Explain your derivation and state any additional assumptions you had to make.

3. Reinforcement Learning (30 points)

i) Consider the state-space drawn below; states are represented as nodes and actions by directed arrows. The number near an arrow is the current estimate of the *Q*-value for that state transition. Unless indicated, the reward (R) received upon entering a state is zero.

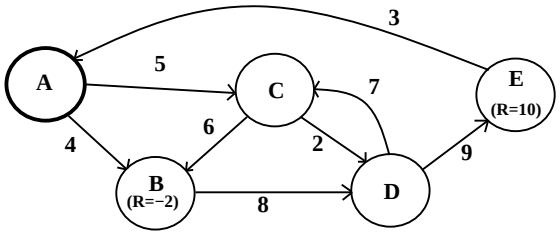

Assume an agent starts in state *A*, deterministically follows the policy specified by the current *Q*-function, and wishes to maximize its discounted cumulative reward (with $\gamma = 0.9$). Also assume this agent is performing *one-step Q-learning* (with $\eta = 0.25$).

In the drawings below, circle the arcs the agent will take and show the updated *Q*-function following each of its first four steps. *Briefly* explain your calculations in the space to the right.

---

t=1

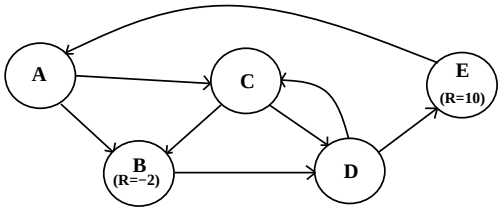

---

t=2

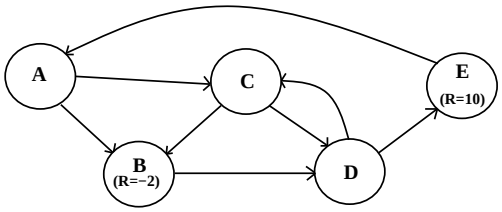

---

t=3

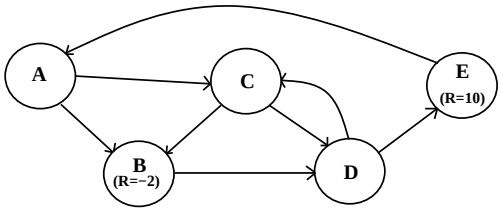

---

t=4

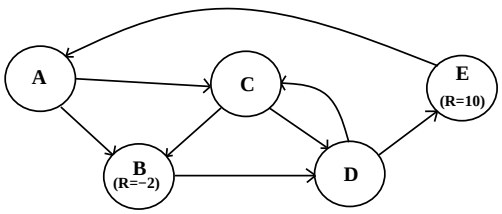

ii) *Briefly* discuss the impact of using a neural network in part (i) to represent the *Q*-function. Is there any additional information you would need to do this? Assuming you were given any missing information, qualitatively describe how your solution to part (i) might change.

iii) Why in reinforcement learning does it make sense to occasionally *not* follow the current policy? Demonstrate with a simple, concrete example.

4. Short Essays (15 points)

i) *Briefly* discuss the significance of the following ML topics.

---

inductive bias

---

overfitting avoidance

---

t-tests

---

ii) Describe the differences between *feature space* and *concept space*.
Is there any relationship between them?