CS 540: Introduction to Artificial Intelligence

Final Exam: 2:45-4:45pm, December 18, 2000
Room 1221 CS & Stats

CLOSED BOOK
(two sheets of notes and a calculator allowed)

Write your answers on these pages and show your work. If you feel that a question is not fully specified, state any assumptions that you need to make in order to solve the problem. You may use the backs of these sheets for scratch work.

Write your name on this and all other pages of this exam. Make sure your exam contains seven problems on nine pages.

Name _____

Student ID _____

Problem	**Score**	**Max Score**
1	_____	**15**
2	_____	**15**
3	_____	**15**
4	_____	**10**
5	_____	**15**
6	_____	**12**
7	_____	**18**
TOTAL	_____	**100**

Name: _____

Problem 1 – First-Order Predicate Calculus (15 points)

Convert each of the following English sentences into First-Order Predicate Calculus, using reasonably named predicates, functions, and constants. If you feel a sentence is ambiguous, clarify which meaning you're representing in logic.

Every election has a winner.

Only those trees that are tall have long roots.

Eating dinner does not impact one's height. [You must use *situation calculus* here.]

Sledding down a hill lowers one's altitude. [You must use *situation calculus* here.]

Name: _____

Problem 2 – Miscellaneous Questions (15 points)

(a) Do these two Wff's *unify*? Justify your answer.

P(?x, ?y, ?x) *P(f(?a, ?b), ?a, f(?b, ?b))*

(b) Imagine you are assigned the task of creating a *production system* that helps people buy birthday presents. Present (in a formal notation) one production rule that would be useful in such a system. Include some brief comments about that it does.

(c) What kind of *search strategy* (e.g, best-first, depth-first, hill-climbing, beam, etc) do productions systems employ? Justify your answer.

(d) What do you feel is the most significant difference between the machine learning task that *decision-tree induction* ("supervised learning") addresses and that addressed by *Q-learning* ("reinforcement learning")? Explain your answer.

Name: _____

Problem 3 – Neural Networks (15 points)

Consider a *perceptron* that has two real-valued inputs and an output unit with a sigmoidal activation function. All the initial weights and the bias ("threshold") equal 0.5. Assume the teacher has said that the output should be 1 for the input *in1 = 0.7* and *in2 = -0.3*.

Show how the perceptron learning rule would alter this neural network upon processing this training example. Let η (the learning rate) be 0.1, and be sure to adjust the output unit's bias during training.

Perceptron BEFORE Training

Perceptron AFTER Training

Briefly explain why a perceptron cannot learn the "exclusive OR" function?

Name: _____

Problem 4 – Genetic Algorithms (10 points)

Consider the following *fitness function*:

Fitness(<bitstring>) = number of 1's in the bitstring where ***both*** *adjacent bits are 0's*

For example, fitness("010110100") = 2, fitness("100011011") = 0, and fitness("010101010") = 4. (Notice that 1's in the first or last position in the string are not counted in the fitness function, even if adjacent to a 0.)

Assume the design of our genetic algorithm is:

(a) Create an initial population containing 4 random 9-bit strings.
(b) Discard the 2 least-fit ones (break ties randomly).
(c) Do a *cross-over* using the 2 most fit.
The 2 children that results and their parents constitute the next generation.
(d) Randomly *mutate* 1 bit in 1 string in the population.
(e) Go to step (b)

Start with the initial population below and show what the next two (2) generations might look like. Explain your reasoning.

Generation 0

011110110
011001011
101101110
000010101

Generation 1 ***Explanation***

Generation 2 ***Explanation***

Name: _____

Problem 5 – Important AI Concepts (15 points)

Describe each of the following AI concepts and briefly explain its most significant aspect.

Inference Rules

Searle's Chinese-Room Story

Fuzzy Logic

Vector-Space Model

Weight Space

Name: _____

Problem 6 – Bayesian Networks (12 points)

Consider the following Bayesian Network, where variables **A-E** are all Boolean-valued:

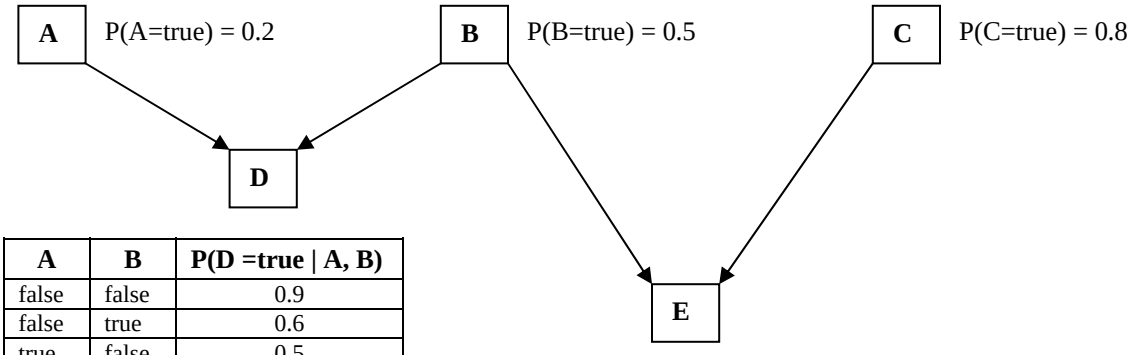

A	B	P(D =true | A, B)
false	false	0.9
false	true	0.6
true	false	0.5
true	true	0.1

B	C	P(E=true | B, C)
false	false	0.2
false	true	0.4
true	false	0.8
true	true	0.3

(a) What is the probability that all five of these Boolean variables are simultaneously true?

(b) What is the probability that **A** is false given that the four other variables are all know to be true?

Name: _____

Problem 7 – Bayes' Rule (18 points)

(a) First-grader Maggie has divided her books into two groups, those she likes and those she doesn't.

The five (5) books that Maggie likes contain (only) the following words:

animal (5 times), *mineral* (15 times), *vegetable* (1 time), *see*(1 time)

The ten (10) books that Maggie does *not* like contain (only) the following words:

animal (5 times), *mineral* (10 times), *vegetable* (30 times), *spot*(1 time)

Using the Naïve Bayes assumption, determine whether it is more probable that Maggie likes the following book than that she dislikes it. Show and explain your work.

see mineral vegetable // These three words are the entire contents of this new book.

Name: _____

(b) In the general population, one in a million people have the dreaded *xyz* disease. Fortunately, there is a test (*test4xyz)* for this disease. Unfortunately, it is only 99% accurate. That is, if you have the disease, 99 times out of 100 *test4xyz* will turn out positive; if you do *not* have the disease, 99 times out of 100 the test will turn out negative.

You take this test and the results indicate you have the disease. Use Bayesian reasoning to calculate the probability that you actually have the disease. That is, compute:

$$Prob(haveXYZ = true \mid test4xyz = true)$$