

# CS 540: Introduction to Artificial Intelligence

*Final Exam: 1:00-3:30 pm, August 8, 2003  
Room 265 Materials Sciences Building*

CLOSED BOOK

(two-sided sheet of handwritten 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 you need to make in order to solve the problem. You may use the backs of these sheets for scratch work. Budget your time wisely.

Before you begin, write your name on every page of the exam and read through all the questions (as some have multiple parts and are more involved than others). Make sure your exam contains *seven (7)* problems on *eleven (11)* pages.

Name \_\_\_\_\_

Student ID \_\_\_\_\_

<u>Problem</u>	<u>Score</u>	<u>Max Score</u>
1	_____	20
2	_____	15
3	_____	25
4	_____	15
5	_____	30
6	_____	10
7	_____	10
TOTAL	_____	125

## Problem 1: General Search Questions (20 points)

a) Answer each of the following questions *true* or *false*:

- i. Tabu search with a horizon of 1 behaves the same as a greedy hill-climbing search.
- ii. Simulated annealing with a temperature  $T = 0$  also behaves identically to a greedy hill-climbing search.
- iii. Breadth-first search is always a complete search method, even if all of the actions have different costs.
- iv. When hill-climbing and greedy best first search use the exact same admissible heuristic function, they will expand the same set of search nodes.
- v. If two admissible heuristic functions evaluate the same search node  $n$  as  $h_1(n) = 6$  and  $h_2(n) = 8$ , we say  $h_1$  dominates  $h_2$ , because it is less likely to overestimate the actual cost.

b) Provide a short answer to the following questions:

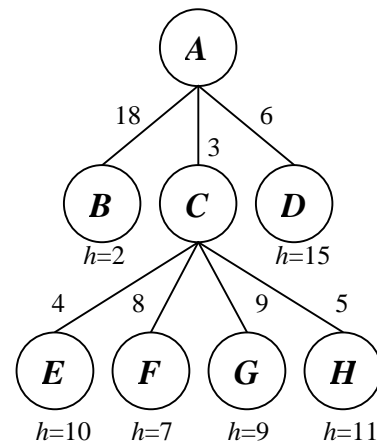
- i. What are the three basic components of any genetic algorithm?
- ii. Compare and contrast genetic algorithms to beam search.

- c) Consider the following *partial search tree*, where edges are labeled with actual costs of the associated action, and each node is labeled with its heuristic evaluation. Which node will be expanded next by each of the following search methods?

i. Greedy best-first search:

ii. Uniform cost search:

iii. A\* search:



- d) Your good friend is the groundskeeper at the mansion of mean old Mr. Mathis. In the backyard, there is a huge fountain with a complex network of pipes controlled by over 100 valves. One weekend, Mr. Mathis announces he's going on vacation, and when he returns he wants the fountain to spray as high as it can... but the plans for the pipe network have been lost! Plus, since she only has the weekend, your friend can't possibly try *all* of the valve combinations to find the optimal setting. Which local optimization search method might she want to use *in real life* (since the fountain can't be simulated on a computer) to maximize the height of the fountain? You may assume that a valve is either *on* or *off*, and the water height is easily measured. If you need to make any other assumptions, state them clearly.

**Problem 2: Logic and Planning Questions (15 points)**

e) Provide a short answer to each of the following questions:

- i. Are the literals  $P(F(y), y, x)$  and  $P(x, F(A), F(v))$  unifiable? If so, show the unifying substitution  $\theta$ . If not, explain why.
- ii. Convert this FOL sentence to conjunctive normal form:  $\forall x \exists y \text{ Owns}(x, y) \Rightarrow \neg \text{Like}(x, y)$
- iii. What is the frame problem in situation calculus? How do we deal with it?

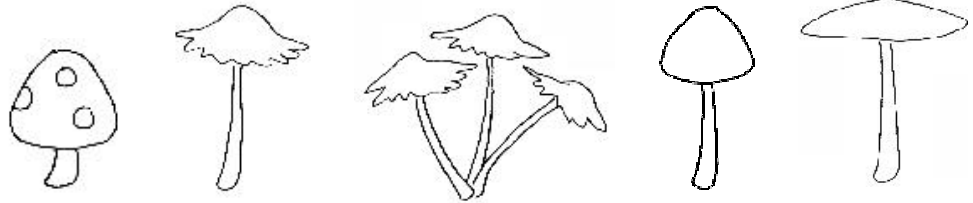
f) Use forward chaining to show that  $\text{Bigger}(\text{Smaug}, \text{Bilbo})$  is entailed by the following KB. Indicate which sentences were used with GMP, and show each substitution  $\theta$ .

- |                                    |   |
|------------------------------------|---|
| 1. $\text{Hobbit}(\text{Bilbo})$   | 4. $\text{Dragon}(x) \wedge \text{Wizard}(y) \Rightarrow \text{Bigger}(x, y)$       |
| 2. $\text{Wizard}(\text{Gandalf})$ | 5. $\text{Wizard}(x) \wedge \text{Hobbit}(y) \Rightarrow \text{Bigger}(x, y)$       |
| 3. $\text{Dragon}(\text{Smaug})$   | 6. $\text{Bigger}(x, y) \wedge \text{Bigger}(y, z) \Rightarrow \text{Bigger}(x, z)$ |

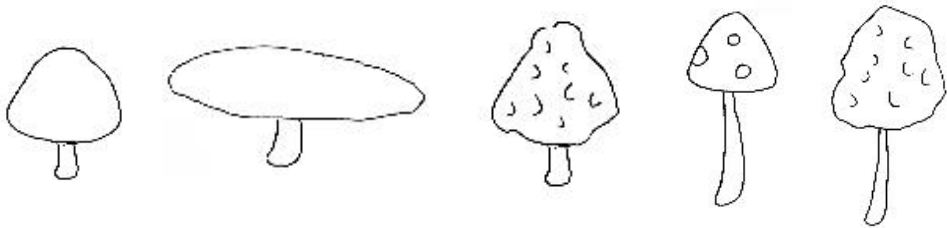
### Problem 3: General Learning Algorithms (25 points)

Consider the task of learning to identify mushrooms that are SAFE or POISONOUS to eat based on a set of physical features. Four Boolean and discrete valued features that you could use are: STEM = {short, long}, BELL = {rounded, flat}, TEXTURE = {plain, spots, bumpy, ruffles}, and NUMBER = {single, multiple}. Consider using these features on the following training data:

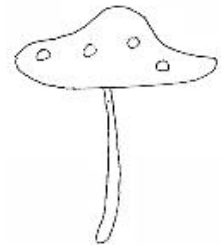
**SAFE:**



**POISONOUS:**



- a) How would the naïve Bayes classifier label the following example:

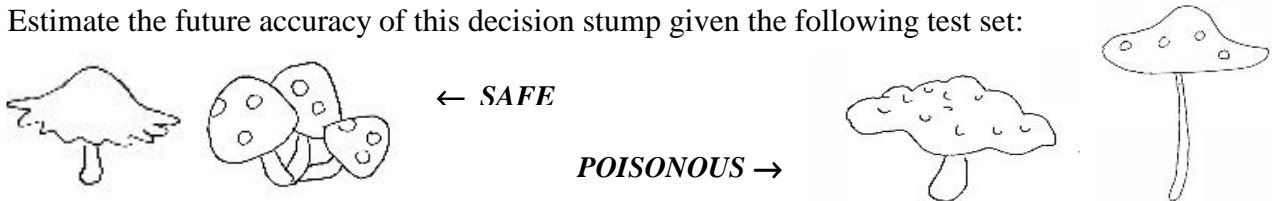


- b) How would 3-nearest neighbors, using hamming distance and unweighted voting, classify the same example from part (a)?

Name: \_\_\_\_\_

- c) Use information gain to choose between TEXTURE and NUMBER to use in a *decision stump* (a decision tree with only one internal node at the root). Draw a diagram of the learned stump, and break ties at classification nodes by labeling them as *POISONOUS* (just to be on the safe side). Show all your work for partial credit.

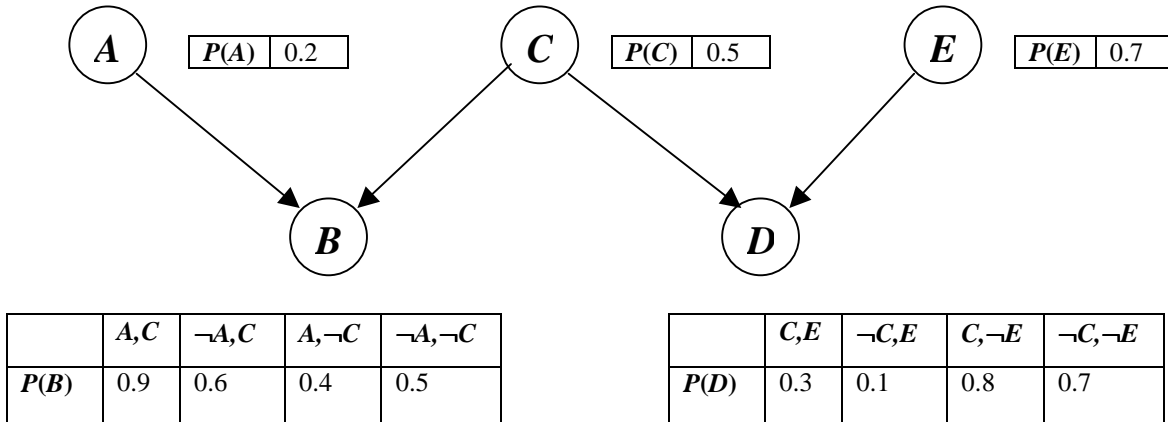
- d) Estimate the future accuracy of this decision stump given the following test set:



- e) Draw a perceptron that might be used to learn this problem (just the structure is fine, you do not need to label with weights).
- f) What is the dimensionality of a perceptron's hypothesis space for this problem?

### Problem 4: Bayesian Networks (15 points)

Consider the following Bayesian network where all variables are Boolean:



- How large would the full-joint probability table (FJPT) have to be for this problem?
- What is the joint probability that all five variables are simultaneously true?
- Compute the probability  $P(\neg C \mid A, \neg B, D, \neg E)$

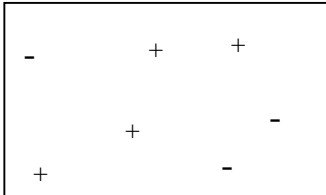


Name: \_\_\_\_\_

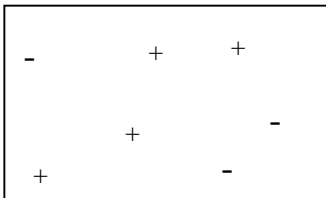
- e) What is the *curse of dimensionality*? Which learning methods are most affected by it? Describe one approach to dealing with this problem.
  
  
  
  
  
  
  
  
  
  
- f) What relationship does the *minimum description length* principle have with the ID3 learning algorithm? What about backpropagation learning for neural networks?
  
  
  
  
  
  
  
  
  
  
- g) Ensemble learning algorithms often benefit from the combined knowledge of its constituent agents, which learn from *slightly* different subsets of examples. Imagine that we want to create an ensemble learner where each constituent agent learns from subsets that are *very* different. How might we do this? How might the ensemble then classify new examples? Clearly state any assumptions that you need to.

## Problem 6: Feature Space (10 points)

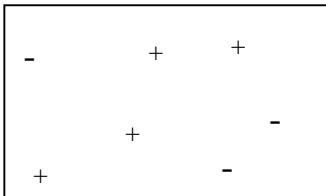
The following Venn diagrams represent training examples for a Boolean-valued concept function plotted in *feature space*. Show how each of the following machine learning algorithms *might* partition the space based on these examples. Briefly explain to the right of each diagram why that algorithm would partition the data that particular way. (No need for calculations... just give a qualitative answer.)



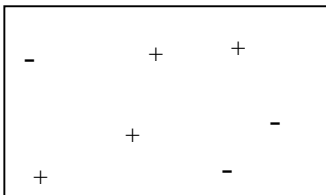
Naïve Bayes



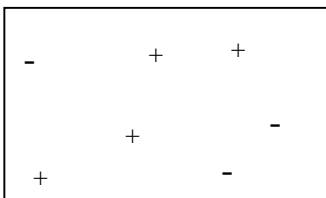
1-Nearest Neighbor



ID3



Neural Network with 1 hidden layer of 2 units



Neural Network with unlimited hidden layers/units

Name: \_\_\_\_\_

## Problem 7: Miscellaneous Definitions (10 points)

Briefly define the following concepts and explain the significance of each to A.I.  
(Write your answer *below* the phrase.)

---

*$\alpha$ - $\beta$  Search*

---

*Occam's Razor*

---

*Tuning Sets*

---

*Promotion / Demotion*

---

*The Markov Assumption*