CS472
Foundations of Artificial Intelligence

Final Exam

December 14, 2000
12-2:30pm

Name:

Instructions: You have 2.5 hours to complete this exam. The exam is a closed-book, closed-notes exam.

#	description	score	max score
1	search I	____	15
2	search II	----	15
3	machine learning I	____	20
4	machine learning II	____	25
5	planning	____	20
6	logic and knowledge-based systems	____	25
7	learning and search	____	15
Total score:		________ /	135

Search: (15 points total)

To celebrate the end of the semester, Claire (C) and the four CS472 TA's (Eric (E), Jason (J), Kamen (A), and Kevin (K)) hope to eat out for dinner at Just a Taste, a tapas restaurant in downtown Ithaca. They want to order five dishes from the menu to share. The dishes must satisfy the following contraints:

- Claire is a quasi-vegetarian, so at most one dish can use beef or pork.
- Eric doesn't like hot foods, so at most one dish can be spicy.
- Kevin and Kamen appreciate variety, so no two dishes can have the same main ingredient.
- In honor of Jason and the capture-the-flag assignment, one dish must have some sort of flag-like garnish.

Be sure to read all parts of the question before starting.

1. (5 pts) Specify this task as a state-space search problem. What is the initial state? The goal state? The operators that transform one state into another?

2. (3 pts) How does the **state space** change if you are to find a legal combination of five dishes with the least total cost? How must the **search** change?

3. (3 pts) Design a heuristic function for the problem of part 2, suitable for use in the A^* algorithm, that can improve performance over blind search. Will your heuristic allow A* to obtain an optimal solution? Why or why not?

4. (4 pts) Describe the original task as a constraint satisfaction problem making any additional definitions or assumptions as needed.

Heuristic Search: (15 points total)

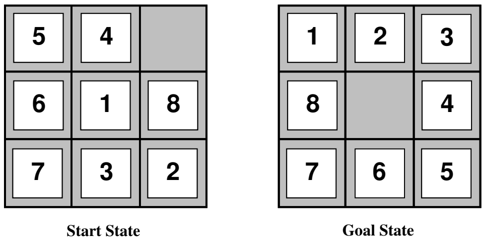

Figure 1: 8-puzzle Example.

Consider the problem of solving the 8-puzzle as depicted in Figure 1 and as described in class. Determining a set of moves that will take you from the start board to the goal board can be viewed as a problem in state-space search. In addition, there may be several possible heuristic functions that estimate the distance (i.e., the number of moves) from the current board to the goal board. One such heuristic function, h, can be computed by counting the number of tiles that are out of place with respect to the goal board.

For each of the following search strategies, indicate whether the search **will succeed**, **will succeed optimally**, or **will not succeed** given a valid arbitrary starting board and goal board. Give a short (one sentence) justification of your answer. (No credit will be awarded without the justification.)

1. (3pts.) Hill-climbing with h.

2. (3pts.) Simulated annealing with h.

3. (3pts.) A* search with h.

4. (3pts.) Alpha-beta search.

5. (3pts.) A* search with $2h$.

Machine Learning I (20 points total)

The CS472 course staff is interested in developing a machine learning classification system that predicts the performance of a student on the CS472 final exam. To make its decisions, the classifier has access to information on a small number of previous CS472 exam-takers. Each of the training cases (shown below) identifies: (1) the student's interest in AI as either high or low; (2) whether or not the student has attended most of the CS472 lectures; (3) the amount of sleep obtained the night before the exam; and (4) whether or not the student passed the CS472 final exam.

	Interest in AI	Attended Most Lectures	Sleep	Exam Performance
1	high	yes	little	passed
2	high	no	lots	failed
3	low	no	lots	failed
4	high	no	lots	failed
5	low	no	lots	failed
6	low	no	lots	failed
7	high	yes	lots	passed
8	low	no	lots	failed
9	low	no	lots	failed
10	low	no	little	passed
11	low	yes	little	passed
12	high	yes	lots	passed
13	high	yes	little	passed

1. (5 pts.) Given the above data, show the decision tree that would be created by a system that uses information gain to determine branching points (i.e. by an ID3-style decision tree induction system). You don't/shouldn't have to show any calculations as long as you can create the tree without them.

2. (6 pts.) There are many variations of the k-nearest neighbor (k-nn) algorithm (which is also known as instance-based learning or case-based learning). Provide commented, high-level pseudo-code for the training and test phases of any such variation. Then use your k-nn algorithm to classify the test case below given the CS472 final exam data set as training.

   - Interest in AI: high; Attended Most Lectures: no; Sleep: lots.

3. (4 pts.) In general, which of the following inductive learning algorithms will provide the best generalization accuracy across a variety of data sets: (1) a decision tree system, (2) a k-nearest-neighbor (case-based) system, or (3) feedforward neural networks trained via backpropagation. Briefly explain your answer.

4. (5 pts.) Choose **one** of the following algorithms: k-nn, decision tree induction, backpropagation on neural nets. Then, given the task of classifying mushrooms as edible or poisonous and a training set of instances of edible and poisonous mushrooms, provide commented high-level pseudo-code for a variation of your selected algorithm that tries to minimize misclassification cost. The misclassification costs for the mushroom task are given in the table below.

	Predicted Class	
True Class	Edible	Poisonous
Edible	0	1
Poisonous	5	0

Machine Learning II (25 points total)

1. (8 pts.) Briefly (e.g. in two sentences) describe the general, two-step procedure for showing that a concept class of a particular form (e.g. decision lists, decision trees, k-CNF) is PAC-learnable.

2. (7 pts.) Show that k-CNF formulas (k literals per clause, k fixed) are PAC-learnable.

3. (5 pts.) Identify two differences between reinforcement learning and supervised learning from examples.

4. (5 pts.) List one merit and one drawback of self-play in learning an evaluation function for game playing.

Planning (25 points total)

1. (4 pts.) List two problems with a search-based problem-solving approach to planning.

2. (4 pts.) List two problems encountered when formulating planning as a logical inference problem, i.e. when formulating planning as deduction within standard first-order logic.

3. (4 pts.) List and briefly describe the four components of a **plan** data structure as employed in a classical partial-order planner.

4. Consider the initial and goal states for the TOWERS-OF-HANOI problem of Figure 2. Each disk must be on some peg and several disks can be on the same peg as long as no disk is on a disk of smaller size. A disk can be moved to another peg if it is clear of all disks and all the disks on the destination peg are larger than the disk we are trying to move.

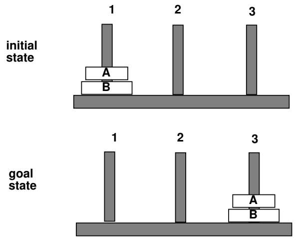

Figure 2: Towers of Hanoi Example.

We want to solve the towers-of-Hanoi problem of Figure 2 using a partial order planner. Assume that the initial state is described as follows: *(on A B) (on B 1) (smaller A B) (disk A) (disk B) (peg 1) (peg 2) (peg 3) (clear A) (clear 2) (clear 3)*. In addition, assume that we have three operators available (each of which conforms to the rules of the TOWERS-OF-HANOI problem as stated above):

(a) **move-d-d-p**: Move a disk from a disk to a peg.

(b) **move-d-p-p**: Move a disk from a peg to a peg.

(c) **move-d-p-d**: Move a disk from a peg to a disk.

(a) (4 pts.) Show the initial partial order plan (either pictorially or as a plan data structure) for the towers-of-Hanoi problem of Figure 2.

(b) (4 pts.) Show one possible non-initial partial plan for the towers-of-Hanoi problem of Figure 2.

Logic and Knowledge-based Systems (25 points total)

1. (6 pts.) Represent the following statements using *first-order logic*.

   If a day is sunny, some people are cheery. If a day has a thunderstorm, no people are cheery.

2. (7 pts.) Prove by resolution that, *if a day has a thunderstorm, the day isn't sunny.*

3. (4 pts.) Define the *frame problem* and provide a concrete example of it.

4. (4 pts.) Specify the components of a forward-chaining production system.

5. (4 pts.) Describe one major problem and one advantage of using traditional forward-chaining rule-based systems for problem-solving.

Learning and Search (15 points total)

Given a set of labelled training examples, E, and a specific network structure, the backpropagation algorithm tries to learn a set of weights that minimizes the classification error across all examples in E. Consider the problem of using a genetic algorithm to learn the optimal number of hidden units for the learning task defined by E. Specify each component of the genetic algorithm, e.g. the individuals, the fitness function, crossover and mutation operators. You should assume that the network is a fully connected, feedforward network with one layer of hidden units. You can also assume a particular maximum number, n, of hidden units (e.g. $n = 32$).