SPRING 1997
COMPUTER SCIENCES DEPARTMENT
UNIVERSITY OF WISCONSIN – MADISON
Ph.D. QUALIFYING EXAMINATION

Artificial Intelligence

MONDAY, FEBRUARY 3
3:00 – 7:00 PM
ROOM 2365 CS

GENERAL INSTRUCTIONS:

1. Answer each question in a separate book.
2. Indicate on the cover of *each* book the area of the exam, your code number, and the question answered in that book. On *one* of your books, list the numbers of *all* the questions answered. *Do not write your name on any answer book.*
3. Return all answer books in the folder provided. Additional answer books are available if needed.

SPECIFIC INSTRUCTIONS:

Answer:

- either one (1) of the questions G540-1 or G540-2, *and*
- both (2) questions in the section labeled B760, B761, or B766 corresponding to your chosen focus area, *and*
- any two (2) additional question in the sections B760, B761, and B766, where these two questions need *not* come from the same section, *and*
- both (2) questions in the section labeled A76x that corresponds to your focus area.

Hence, you are to answer a total of *seven* (7) questions.

POLICY ON MISPRINTS AND AMBIGUITIES:

The Exam Committee tries to proofread the exam as carefully as possible. Nevertheless, the exam sometimes contains misprints and ambiguities. If you are convinced that a problem has been stated incorrectly, mention this to the proctor. If necessary, the proctor can contact a representative of the area to resolve problems during the *first hour* of the exam. In any case, you should indicate your interpretation of the problem in your written answer. Your interpretation should be such that the problem is nontrivial.

Answer either one (1) of the questions G540-1 or G540-2.

G540 – Introduction to AI Questions

G540-1. Consider the following two sentences in **first-order logic**:

1. $\forall x\big(\text{boy}(x) \;\Rightarrow\; \exists y(\text{girl}(y) \wedge \text{likes}(x, y))\big)$
2. $\exists y\big(\text{girl}(y) \;\wedge\; \forall x(\text{boy}(x) \;\Rightarrow\; \text{likes}(x, y))\big)$

One of these sentences logically implies the other. Which one? Prove your claim using resolution refutation.

G540-2. Consider the following **game**: Start with a deck of four cards, labeled 1,4,5,6. The computer (i.e, your program) and the user take turns, with the computer going first. At each turn, the player takes a card from the deck. The game ends after precisely two plays by each player, at which point the deck is empty; each player adds up the two cards chosen, and the player whose sum is closest to 10 wins.

computer: 1
user: 5
computer: 4
user: 6

For this play, the computer's sum is 5 and the user's sum is 11, which is closer to 10, so the user wins.

(a) Draw the complete game tree, with the numerically smaller choices on the left. This tree should have 24 leaves.

(b) Assume the computer does an alpha-beta search to determine its best first move. Cross off the nodes in the tree which don't get searched. For *one* of the nodes you crossed off, explain why you crossed it off. Note that the "payoffs" in this game may be considered to be just +1 (computer wins) and −1 (computer loses). Assume that the children of each node are searched left-to-right.

Answer both (2) of the questions in the section labeled B76x that corresponds to your chosen focus area. Also answer any two (2) additional questions in any of the other sections (these two questions need NOT occur in the same section).

B760 – Machine Learning Basic Questions

B760-1. Briefly explain the concept (hypothesis) space and inductive bias of Quinlan's C4.5 *decision tree* algorithm.

Assume that (a) Mitchell's Candidate Elimination algorithm is used to search the same concept space and (b) C4.5 is configured *not* to prune the initial tree it constructs. Would the tree that C4.5 produces be in the resulting version space, assuming they both used the same training examples? If so, would it be a member of the set S? Set G? How would your answer change if C4.5's pruning algorithm was enabled?

B760-2. Imagine you just were hired by a large bank and are given the task of producing a system that detects the fraudulent use of credit cards.

(a) Outline the steps you would take to solve this problem using *machine learning*.

(b) Your new boss asks you to provide an estimate of the accuracy of the system you produce. Explain how you could do so.

B761 – Deductive Problem Solving & Logic Programming Basic Questions

B761-1. Derive the empty clause from the following clauses using binary resolution plus paramodulation. a, b, c are constants and x, y, z are variables.

1. $x = x$
2. $\max(x, y) = x \vee \max(x, y) = y$
3. $\max(x, y) = \max(y, x)$
4. $\max(x, y) \neq y \vee \max(y, z) \neq z \vee \max(x, z) = z$
5. $\max(b, c) = b$
6. $\max(a, \max(b, c)) \neq \max(\max(a, b), c)$

You may use factoring, which is part of binary resolution. Actually, you don't need the $\max(b, c) = b$, but it makes the proof easier. A "human" proof is: We have to show $\max(a, b) = \max(\max(a, b), c)$. If $\max(a, b) = b$, we have $\max(a, b) = b = \max(b, c) = \max(\max(a, b), c)$. If $\max(a, b) = a$, we have $\max(c, b) = b$ and $\max(b, a) = a$, so $\max(c, a) = a$ (by (4)), so $\max(a, b) = a = \max(a, c) = \max(\max(a, b), c)$.

B761-2. Explain the Prolog negation-as-failure, and why it is sound without cut, but not sound if cut is present. Consider ground programs only, and consider programs which contain no negations in the program (but, negations can appear in the queries). A program you might use to illustrate unsoundness is:

```
a :- b(1), !, c(1).
a :- b(2), c(2).
b(1).
b(2).
c(2).
```

Then the query `?- not a.` (written `?- \+ a` in some prologs) will get the answer `yes` even though the program doesn't semantically imply $\neg a$ (in fact, it implies a).

B766 – Computer Vision Basic Questions

B766-1.

(a) Define the cross-correlation match measure for an image **f** and a template **t**. Include a description of how this measure is used for two-dimensional object recognition (i.e., at which points is an object associated with template **t** detected).

(b) Justify cross-correlation as a match measure by showing how it is related to the sum-of-squared-error measure.

(c) Describe two major limitations of using cross-correlation as a match measure for 2D object recognition.

B766-2.

(a) Describe the steps in detecting edges using the Laplacian-of-Gaussian edge detector, including a definition of the two-dimensional operator for a digital image.

(b) Give one main advantage and one main disadvantage in using the Laplacian-of-Gaussian operator instead of the Canny edge operator for detecting edge points.

Answer both (2) of the questions in the section labeled A76x that corresponds to your chosen focus area.

A760 – Machine Learning Advanced Questions

A760-1. One can apply the *reinforcement learning* technique called Q-learning both when one does and does not have a model of the process being controlled.

- **(a)** Explain what it means, in the context of Q-learning, to have a "world model." What still needs to be learned if one has a correct world model? How does the presence or absence of a model alter the way one performs Q-learning?
- **(b)** Explain why a model is required if one simply estimates the utilities of states. How is this need for a model overcome by Q-learning?
- **(c)** Consider the often stated AI claim that "the world is its own best model." Present one (1) argument for and one (1) argument against this claim, with respect to Q-learning.
- **(d)** Imagine one has an *approximate* world model, expressed, say, in propositional logic. Present and critique an approach that would make use of this world model, without necessarily being hindered in the long term by the imperfections in the approximate world model.

A760-2. Define the *overfitting* problem in machine learning and describe its importance. Explain how it can arise and is addressed in each of the following approaches:

- **(a)** instance-based (i.e., nearest-neighbor) methods.
- **(b)** ensemble methods (e.g., training five neural networks and having them vote on the predicted category).
- **(c)** explanation-based learning.