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

Artificial Intelligence

MONDAY, SEPTEMBER 15
3:00 – 7:00 PM
ROOM 3345 ENGINEERING

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. Represent the following English sentences in the *situation calculus*:

Getting a speeding ticket increases one's insurance rates.

Painting a house changes its value but not who owns it.

Today, apples cost more per pound than bananas.

Be sure to explain the semantics of the predicates and functions you use.

G540-2. Consider the following list of words:

CAB CAD CAT COD CUD DAB HAT ROB ROD ROT

and consider the puzzle: Change "cat" to "rob" by changing only one letter at a time, and *using only words contained in the list.* You may consider this as a graph problem with "cat" as the start node and "rob" as the goal node. Draw the graph. One may solve puzzles like this by using best-first search, with the heuristic being the number of wrong letters. So, "cat" has heuristic value 3, while "cab" has heuristic value 2. Explain why this heuristic is admissible. Then, trace best-first search in solving the puzzle. In the trace, if you have a choice of two nodes to explore with the same estimated cost, take the alphabetically first one.

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. Assume you wish to create a system that classified WWW pages as either *interesting* or *uninteresting*. Sketch the steps you would need to go through in order to apply Quinlan's *decision-tree algorithm* to this task. Be sure to concretely describe how you would represent training examples. Describe one (1) significant strength and one (1) significant weakness of applying decision-tree induction, as opposed to another machine learning method, to this task. Finally, name another machine learning approach that addresses the weakness you described; justify your answer.

B760-2. Explain the *exploration/exploitation* tradeoff in *reinforcement learning* (RL). Devise and describe a simple, concrete RL environment where a non-optimal policy can be learned by Q learning if the system always chooses the action with the highest current Q-value. It is fine, but not necessary, to assume absorbing states, repeated restarts at a specific or random state, non-deterministic actions, function approximation by a neural network, etc. Do explain any assumptions you need to make.

Explicitly show the calculations for the first three Q-learning steps in your environment.

B761 – Deductive Problem Solving & Logic Programming Basic Questions

B761-1. Convert the following logical sentences to clauses, and then derive the empty clause using binary resolution plus paramodulation.

1. $\forall x(x = x)$
2. $\forall x \exists y(x = g(y))$
3. $\forall x \forall y(f(g(x)) = f(g(y)) \;\Rightarrow\; x = y)$
4. $\neg \forall x \forall y(f(x) = f(y) \;\Rightarrow\; x = y)$

B761-2. Consider the database:

$$\forall X(person(X) \;\Rightarrow\; (m(X) \vee f(X)))$$
$$person(pat)$$

(a) Explain why neither $m(pat)$ nor $f(pat)$ are logical consequences of the database, but $m(pat) \vee f(pat)$ is a logical consequence of the database.

(b) The *closed world assumption* is the informal idea that if a statement φ is not a logical consequence of the database, then we conclude $\neg\varphi$. In this example, we conclude $\neg m(pat)$ and $\neg f(pat)$, which, together with $m(pat) \vee f(pat)$, is contradictory. Explain how the closed world assumption is handled in the formal semantics of logic programming languages, such as Prolog, and explain why such contradictions are avoided.

B766 – Computer Vision Basic Questions

B766-1. Briefly describe *thresholding-based methods for image segmentation* for each of the following situations:

(a) A scene containing non-overlapping, dark objects on a light background, but the illumination is very uneven so that one corner of the image is very dark and the opposite corner is very light.

(b) A color image of an outdoor scene containing regions that each have a distinct, but unknown, mean color, and the sizes of the regions vary significantly from very large to very small.

(c) An image containing three regions, all with the same mean gray level, but distinguishable by their differing textures: background (fine texture), region1 (uniform intensity), and region2 (coarse texture).

B766-2.

(a) Define the *Laplacian operator*, both in its continuous form and as a discrete 3x3 operator.

(b) Is the Laplacian (i) rotation invariant, (ii) linear, (iii) shift invariant, (iv) implementable as a convolution?

(c) You are given an image, g, that is a blurred (i.e., averaged) version of an ideal, unblurred image, f, where

$$g(i,j) = 1/5[f(i,j) + f(i+1,j) + f(i-1,j) + f(i,j+1) + f(i,j-1)].$$

Show how the original image f (or an image that "close to" f) can be recovered from g and the Laplacian of g.

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

A766 – Computer Vision Advanced Questions

A766-1.

(a) Briefly describe the *appearance-based representation* method (as used by Turk and Pentland, and Murase and Nayar) as it would be used to represent a set of frontal, mug-shot-type images of human faces.

(b) What are two main advantages of this approach over the more traditional shape-based approaches to recognition?

(c) Describe two kinds of image normalization and/or preprocessing steps that must be done for each image in order to use this method? Briefly describe each and why it is necessary.

(d) How might the method be extended to allow for partial occlusion?

A766-2.

(a) One of the most important issues in creating *image mosaics* is aligning a set of images into a single common coordinate frame. Say we have a set of images of a small area of the Earth captured by a camera on a satellite.

(i) What are the transformational degrees of freedom needed for this task in order to align all of the satellite images?

(ii) Define an error measure that is appropriate for determining the best alignment between two images.

(b) Argue for or against the following: Mosaics are only good for static scenes viewed from a distance and could not be adapted for representing the appearance of a 3D object over a large range of poses (i.e., 3D rotations).