FALL 1999
COMPUTER SCIENCES DEPARTMENT
UNIVERSITY OF WISCONSIN – MADISON
PH.D. QUALIFYING EXAMINATION

Artificial Intelligence

September 13, 1999
3:00 – 7:00 PM
2345 Engineering Hall

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, B766, or B780 corresponding to your chosen focus area, *and*
- any two (2) additional question in the sections B760, B766, and B780, where these two questions need *not* come from the same section, *and*
- both (2) questions in the section labeled A7xx 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. Convert the following English sentences into first-order predicate calculus.

a. Buying a painting changes its owner but not its artistic value.

b. Everyone who currently can be seen by someone else has a unique costume.

c. If some number is a perfect square then that number is an acceptable key for Mary's code.

G540-2. Imagine you want to cast as a search problem the task of route planning between any two cities in a map. Assume that your map contains the cities A, B, C, D, and E and that the following (direct connection) roads exist:

- between A and B (27 miles)
- between A and C (10 miles)
- between A and E (92 miles)
- between B and C (12 miles)
- between B and D (15 miles)
- between D and E (12 miles)

All roads are bidirectional (i.e., two-way streets). The current task is to find a route between cities A and E. In addition, the "as the bird flies" (i.e., Euclidean) distances to city E are:

- 31 miles from A
- 23 miles from B
- 10 miles from C
- 11 miles from D

a. Represent the above information as a search space. What would be an appropriate heuristic? Is it an *admissible* one? Explain your answer.

b. Show the steps in applying the A^* algorithm to the task of finding a route between cities A and E. What guarantees, if any, does this algorithm offer?

c. Show the steps in applying the *iterative-deepening* algorithm to the same task. What guarantees, if any, does this algorithm offer?

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. An important issue in machine learning is choosing a good set of features for describing examples. Explain how this task could be cast as a heuristic-search problem and briefly illustrate how your feature-selection algorithm works by using a small testbed and a simple nearest-neighbor learner.

B760-2. Describe the tradeoffs in machine learning between using simple and complex "languages" for the concepts being learned. How does this issue arise within the field of neural networks? Sketch and justify a neural-network training approach that estimates an appropriate tradeoff between the simplicity and complexity of induced concepts.

B766 – Computer Vision Basic Questions

B766-1. Consider the task of edge detection.

a. What are three major criteria for an "optimal" edge detector?

b. What edge model is the Canny operator optimized to detect?

c. Suppose we are only interested in detecting edges oriented at multiples of 45 degrees (relative to the x-axis in an image). Define a set of filters and how they would be used/combined in order to perform "edge enhancement," i.e., estimate the edge strength and edge orientation at a pixel (i, j).

B766-2. Consider performing 2D object recognition using interpretation trees.

a. If an image contains 3 features and we are interested in detecting instances of a 2D object model containing 3 features, show the interpretation tree representing the complete search space for this problem. Assume each model feature matches at most one image feature. What does a leaf node represent in this tree?

b. In general, model features may not be visible (because they are occluded by another object), and image features may be noise or from different objects. Explain briefly how each of these two situations could be dealt using the interpretation tree approach.

B780 – Robot Motion Planning Basic Questions

B780-1. Consider the $Bug1$ and $Bug2$ sensor-based, motion-planning algorithms. What is the worst-case length of the path generated by those two procedures? (note: you do not need to prove the limits, just list them.) Explain the symbols used in your expressions. Give one major advantage that $Bug1$ has over $Bug2$ and one major advantage $Bug2$ has over $Bug1$. Explain your answer in detail.

B780-2. Consider a two-dimensional, two-link RR arm manipulator (R = "revolute"). Write the inverse kinematics equations for this arm. How many solutions for the two joint angles are possible? What is the shape of this arm's work space? What about its configuration space? Draw a sketch *or* explain.

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. What is the relation between the fundamental matrix, F, and a pair of corresponding points, p and p', in two images? Give your answer as a formula or as quantitatively as possible.

b. What is the size of F, rank of F, and number of degrees of freedom in F? Explain briefly.

c. What is the relation between the epipole in one image, say el in the left image, and the fundamental matrix, F? Give an equation and explain what it means intuitively.

d. What information is needed in order to do 3D scene reconstruction of the following types:

   i. projective

   ii. affine (i.e., up to a scale factor)

   iii. Euclidean (i.e., metric)

A766-2. As an alternative to the voxel coloring method for 3D reconstruction, consider the following volume intersection method. Assume you have a set of calibrated images, each containing a view of a single object against a black background. For each image, the object (silhouette) is segmented from the background and then back-projected into a cone in 3D scene space. Next, all of the cones are intersected resulting in a volumetric representation of the object.

a. How is the 3D object reconstruction produced by this algorithm likely to be better, worse or the same as the result using voxel coloring and the same set of input images?

b. One problem with voxel coloring is that the correlation test is based on comparing individual pixels, resulting in noisy assessments of whether or not a voxel is consistently projected into the set of images. Briefly describe a generalization of the original test which could be applied to small texture patches instead of single pixels.