FALL 1993
COMPUTER SCIENCES DEPARTMENT
UNIVERSITY OF WISCONSIN—MADISON
PH. D. QUALIFYING EXAMINATION

Artificial Intelligence
Breadth Examination

Monday September 20, 1993
3:00 – 5:00 PM
2365 Computer Sciences

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

- any two (2) of the questions B540-1, B540-2, and B540-3, *and*
- both questions in any *one* of the sections labeled B760, B761, B762, or B766.

Hence, you are to answer a total of *four* (4) questions.

POLICY ON MISPRINTS AND AMBIGUITIES:

The Exam Committee tires to proofread the exam as carefully as possible. Nevertheless, the exam sometimes contains misprints and ambiguities. If you are convinced 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 any two (2) of the three questions labeled B540-x.

B540 – Introduction to AI Questions

B540-1. Explain **alpha-beta search** and give an explicit example. Try to keep your example simple, but it should demonstrate how cutoffs work to prune the tree three levels down as well as one level down.

B540-2. Explain syntactic ambiguity as occurs in **natural language parsing**, using as a specific example the sentence, "I killed the mosquito with a fly swatter." In particular, describe a context-free grammar that produces two parses for this sentence. Then, give another example of syntactic ambiguity that does not involve prepositional phrases.

B540-3. Define the **frame problem** and explain its significance to the AI task of planning. Describe how the frame problem is addressed by (a) STRIPS and (b) situation calculus. Illustrate your answer by showing how a concrete, real-world action could be formalized in each of these two formalisms.

Answer BOTH of the questions in any ONE of the sections labeled B76x.

B760 – Machine Learning Breadth Questions

B760-1. Suppose a machine learning researcher is given a new stock-market dataset and is asked to produce an accurate stock-market advisor. As a first step, this person divides the data into disjoint training and testing sets. Next, the researcher uses backpropagation to independently train 25 **neural networks**, each with one layer of hidden units. The same training examples are always used, but in each run the initial weights, number of hidden units, learning rate, etc. are randomly chosen. Each network is trained until the accuracy on the training set reaches 95%.

Following training, this researcher measures the correctness of each of the 25 networks on the test set. The researcher discards all but the network with the best test-set accuracy and then reports the remaining network's test-set accuracy as the prediction of its future performance.

Critique this *experimentation methodology*. Describe and justify any modifications you would make.

B760-2. State, define, and motivate the two major types of **hypothesis-space biases** used in inductive learning. Give the bias of (a) ID3 and (b) backpropagation. Justify your answers.

B761 – Deductive Problem Solving & Logic Programming Breadth Questions

B761-1. Most **Prolog** systems do not do an *occur check*. Explain what this means, and give an example where it can lead to nonsense results. Also explain why this is usually not a problem in practical Prolog AI programming.

B761-2. If **R** is a binary relation that is irreflexive (x **R** x is always false) and transitive, then it can never have a 3-element cycle (x **R** y and y **R** z and z **R** x). Express this inference in Resolution format, and then **derive the empty clause** using paramodulation plus binary resolution.

B762 – Knowledge Representation & Expert Systems Breadth Questions

B762-1. One popular architecture for expert systems emphasizes an inference engine and a knowledge base as the two primary components.

(a) Portray the structure and function of each in the case where the knowledge of an expert system is represented as **predicate calculus** expressions.

(b) Portray the structure and function of each in the case where the knowledge of an expert system is represented as **frames** (or objects).

B762-2. The GOFAI (good old-fashioned AI) approach is perhaps best summed up in Newell and Simon's **physical symbol system hypothesis**.

(a) Describe the hypothesis.

(b) Explain how the connectionist approach either fits within the GOFAI approach, changes part of it, or changes all of it.

B766 – Computer Vision Breadth Questions

B766-1. It is often recommended that **edge detectors** be applied at multiple scales to an image and these outputs then be "integrated" in some way.

(a) Give an example of a single image where detecting edges at two different scales is necessary to detect different salient structural features.

(b) Given a single image, how are multiple resolution images usually produced? Be as specific as possible. Is it appropriate or inappropriate for this step to use isotropic (i.e., rotationally-invariant) operators? Explain briefly.

(c) Briefly describe a method for combining the results of edge detection at multiple scales into a single representation.

B766-2. Labeling a line drawing of a polyhedral scene can be formulated as a "consistent labeling" problem, assigning a label to each edge subject to constraints at the vertices.

(a) Describe how the **Waltz filtering** algorithm works by showing the initial configuration and several snapshots of the algorithm as it is applied to a simple scene of your choosing.

(b) Characterize how the Waltz filtering algorithm differs from (i) hypothesize-and-test, and (ii) backtracking tree search procedures for solving this problem.