Artificial Intelligence

If you are taking the breadth exam,

answer any *four (4) of* the six breadth (basic) questions *B1 through B6*.

If you are taking the depth exam,

answer any *four (4) of* the six breadth (basic) questions *B1 through B6*

AND

answer any *four (4) of* the twelve depth (advanced) questions *D1 through D12*.

Breadth (basic) questions

B1

Many different AI algorithms can be described in terms of **state-space search**. For two (2) of the following three (3) problem domains, indicate what are the states, the operators, the search strategies, and the goal descriptions.

a) Theorem proving with Prolog.

b) Two-person games.

c) Chapman's TWEAK.

B2

Consider the sentence:
"The boy saw the statue in the park with his telescope after dark."

How many different **parse trees** does this sentence have with the following phrase structure grammar?

S → NP VP
NP → DET NOUN PPS
NP → NOUN PPS
VP → VERB NP PPS
PPS → ε
PPS → PP PPS
PP → PREP NP
NOUN → boy | statue | park | telescope | dark
VERB → saw
PREP → in | with | after
DET → the

You don't have to draw all the parse trees. Just explain how you got your number.

B3

Predicate logic is often used in AI to represent the basic actions that a robot can perform.

a) Consider the action *pushbox(V)* indicating that the robot pushes the box to position *V*, where *V* is a variable. The conditions for this to occur are that the robot is at the same position as the box, the robot is not on top of the box, and its hands are empty. The resulting state of the world has the box and the robot at position *V* and nothing else has changed. Formalize this action in first-order predicate logic (FOPC).

b) How is the **frame problem** addressed when using FOPC to represent actions?

B4

Recent critics of "classical" AI **planning systems** contend that the planning process in these systems is too slow, arguing that planning should be "reactive" in a way it is not in the classical systems.

a) Define the classical approach to planning and illustrate it in terms of the behavior of a particular system. (STRIPS would be a good example to use, but a number of other systems could be used to make your points.) Are the critics correct about the efficiency of this approach?

b) How do the creators of SOAR address this efficiency issue?

c) Typically, in reactive planning the planner monitors the world, chooses only the next operator to apply, applies it, and then repeats the process. While this is fast, there are problems with it. What are they?

B5

It is sometimes argued that for the GOFAI (good old-fashioned AI) **AI paradigm** a problem must be translated so that it becomes one of explicit symbol manipulation *before* AI techniques such as heuristic search, means-ends analysis, symbolic learning, and hierarchical planning can be applied; that the hard part of realizing intelligence is this initial translation, *not* the application of the techniques; and that GOFAI therefore begs the question.

a) Evaluate this argument in terms of some specific system that has been important in the history of AI. (GPS would be a good example to use, but a number of other systems could be used to make your points.)

b) Discuss the extent to which the argument applies to older AI systems but not to ones that have been developed recently. Give examples.

B6

Consider the following claim:

The most important property of a connectionist **learning algorithm** is that it completely learns to properly classify its training examples. Getting stuck in a local minimum is sure to be a disaster.

Criticize this claim, both positively and negatively. To what extent do the claim and your arguments apply to symbolic learning systems like Michalski's AQ?

Depth (advanced) questions

D1

Suppose you are charged with the problem of using the latest AI representations and techniques to design an **intelligent tutoring system**.

a) Specify a feasible kind and level of tutoring for the system.

b) Briefly describe the main kinds of representations and techniques from the field of **knowledge-based systems (expert systems)** that should be considered for incorporation into the design.

c) For each kind, indicate whether and why you would or would not include it in the design.

d) Speak to the problem of integrating the different kinds of representations and techniques that you choose for inclusion.

D2

Suggest possible **parallel algorithms for** the **inference** engine of a **production-based expert system** (for example, one implemented in a "parallelized" OPS-5) and for the inference engine of a **logic-based expert system** (for example, one implemented in a "parallelized" Prolog). Make estimates of and present arguments for the speed-up you would expect, and comment on whether you think such parallelizing might break through any of the problems that are preventing the once widely predicted (but now largely stalled) wide applicability of expert systems.

D3

The **utility problem** refers to an important issue in **explanation-based learning**.

a) Define this problem and discuss its significance.

b) For any three (3) of the following four (4) **machine learning paradigms**, discuss whether or not there is a corresponding issue that could be viewed as its utility problem:

i) Case-based learning.

ii) Conceptual clustering.

iii) Genetic algorithms.

iv) Neural networks.

D4

Define the problem of **constructive induction** in machine learning and explain its significance.

Briefly describe one approach for addressing this problem used in **symbolic inductive learners** and one approach used in **neural networks**. Discuss one major strength and one major weakness of each.

How does the problem of constructive induction relate to integrated approaches that learn from both existing knowledge and training examples?

D5

What are the **philosophical origins** of the **Turing test**? For what kind of AI theories or models is the Turing test accordingly inappropriate because of inconsistencies between its philosophical roots and the philosophical underpinnings of the theories or models?

D6

Discuss the kinds of possible **anaphoric relations in texts** and computational techniques possibly useful for their detection.

D7

Convert the following sentences to clauses and derive the empty clause by binary **resolution** and **paramodulation**.

1) $\forall x\ (x = x)$

2) $\neg\forall v\exists x\ (x * x + (x + x) = v)$

3) $\forall v\exists x\ (x * x = v)$

4) $\forall x\exists y\ (y+1 = x)$

5) $\forall x\forall y\ ((x + 1) * (y + 1) = ((x * y) + (x + y)) + 1)$

6) $\forall x\forall y\forall z\ (x + z = y + z \rightarrow x = y)$

D8

Explain what **non-monotonic logic** is, and then give an argument in support of each of the following:

a) Non-monotonic logic provides a better explanation for human reasoning than does standard monotonic logic.

b) Human reasoning is never non-monotonic, and non-monotonic constructs only arise from misunderstanding what the premises are.

D9

You are given the problem of developing a system to **recognize objects** that are simple enough to be adequately resolved in a 32-by-32 array (1024 8-bit grey-scale spots). There are ten such object classes that this system must recognize. The individual members of five are always pretty much the same (like a coke bottle or the type of pliers your company makes) but five can vary (like the different chairs and tables in the factory).

a) Describe two programs you would try to implement to do this, one a **neural network that learns**, the other a **pre-programmed structure of feature detectors**. Explain and justify your decisions.

b) If you were told to make a system before the specific objects to be recognized have been specified:

i) Which of these two (neural net or pre-programmed structure) would you choose, and why?

ii) What are the major problems you would encounter using each, and what appear to be the best ways to overcome them?

D10

Computer vision systems are often developed as follows: Pre-program a model for each object class to be recognized. Then match each such model to the input image (usually after a *low-level* set of processes has extracted an array of features such as connected contours and regions).

a) Estimate the time and space complexities (and justify these) of such a system:

i) Using a traditional serial computer.

ii) Using an appropriate network.

b) Describe the network you chose, explain how the models would be mapped into it, and explain how the models would be handled by it.

c) What would be the problems in extending such a system to handle larger numbers (100, 1000, 10000, . . .) of different and more complex objects?

D11

Traditional approaches to **image understanding** have constructed a *sequence* of modules for performing subtasks in low-level and intermediate-level vision. Recently, a number of researchers have advocated so-called **integrated methods** in which modules cooperatively solve multiple subtasks simultaneously.

a) Give an example consisting of at least two subtasks that can mutually benefit from information obtained during the solution of the other task(s). Briefly describe how the integrated approach would work in this example.

b) Some people have stated that the integrated approach is *necessary* in order to correctly solve particular sets of subtasks. Briefly argue for or against this position.

D12

Object recognition requires searching for a match between image features and corresponding model features to determine a consistent set of matching pairs for solving for the pose of an object in the image. Briefly describe and contrast the following two methods for solving this matching problem:

a) Generalized Hough transform.

b) Interpretation trees.

Assume a single two-dimensional polygonal object model that can occur in any position and orientation (but fixed size). Also assume you detect n corner features in the image and the model contains m corners.