Artificial Intelligence

If you are taking the **Breadth Exam**:

Answer any four (4) out of the six Basic questions B1–B6.

If you are taking the **Depth Exam**:

Answer any four (4) out of the six Basic questions B1–B6 *and*
answer any four (4) out of the 12 Advanced Depth questions D1–D12.

Breadth (Basic) Questions

B1. No one ***search*** method is best, and each method is most appropriate in different circumstances. For any *three* (3) of the following four pairs of search methods, describe general problem characteristics when the first method is more appropriate than the second.

Breadth-First Search vs. Depth-First Search
Beam Search vs. Best-First Search
Iterative Deepening vs. Breadth-First Search
A* vs. Hill-Climbing

B2. Convert the following English sentences to formulas in ***predicate logic***. Then, use the method of ***resolution*** to show that the last sentence follows from the first three.

Two *different* people are siblings if they have a common parent.
Pat is a parent of Dick and of Rich.
Dick and Rich aren't the same person.
Dick and Rich are siblings.

B3. Describe the general structure of ***production systems***. For each aspect of such a system, discuss the major design issues faced by an implementor and state a common solution for each.

B4. Explain what is meant by the ***physical symbol system hypothesis***. Give an example of research that assumes it. Also, comment on its current status within AI.

B5. A machine with failure ***X*** is referred to you for further diagnosis. Failure ***X*** comes in exactly three disjoint types: ***A, B,*** and ***C***. You decide to use the ***Dempster-Shafer*** technique for ***reasoning about uncertainty*** to interpret some tests you run.

(a) What is the *frame of discernment* (Θ) for this problem? Draw the *hypothesis space*.

(b) The first test provides evidence that the machine experienced failure type ***A***. However, this test is only 75% reliable. What is the *belief probability assignment* (bpa) for this evidence?

(c) The second test provides evidence that the machine did *not* have failure type ***B***; this test is 50% reliable. What is the *bpa* for this data?

(d) Calculate the *bpa* that results from combining these two pieces of evidence.

B6. Michalski's `AQ` and Rumelhart's `backpropagation` ***machine learning*** algorithms can both be viewed as performing search.

(a) What are *states* and *operators* in these two methods?

(b) What kind of *search control* strategy does each use?

(c) Briefly describe any *heuristics* employed by each.

Advanced Depth Questions

D1. Convert the following sentences to clauses, and then use the clauses to derive the empty clause, □, using paramodulation and binary ***resolution***.

$$\neg\, \forall\, x\, (\, (x \cdot x = x) \Rightarrow (\, (x = 0\,) \vee (x = 1)\,)\,)$$

$$\forall\, x\, (\, (x \neq 0) \Rightarrow \exists\, y\, (x \cdot y = 1))$$

$$\forall\, x\, \forall\, y\, \forall\, z\ \ (\, (x \cdot y) \cdot z\ = x \cdot (y \cdot z)\,)$$

$$\forall\, x\, \forall\, y\, (\, x \cdot y = y \cdot x\,)$$

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

$$\forall\, x\, (\, x \cdot 0 = 0\,)$$

$$\forall\, x\, (x = x)$$

D2. Logic programming, in languages such as ***Prolog***, is described as being *programming by specification*. That is, the program is really a set of logical axioms describing the answer being computed, so that any answer obtained is guaranteed to be correct by logic. To what extent is this true of Prolog, and to what extent is this false?

D3. ***Production systems*** and ***predicate logic*** are often presented as alternative knowledge formalisms. How are they different? How, if at all, do they overlap? How do the inference engines typically used in expert systems for manipulating these formalisms differ, e.g., the `OPS5` inference engine for production systems and the Prolog inference engine for predicate logic? Show in general (that is, you don't need to present detailed code) how you would make an `OPS5` system behave like a Prolog-based system and how you would make a Prolog system behave like an `OPS5`-based system.

D4. ***Expert systems*** are often criticized as being "fragile." By this is meant that expert systems can apply knowledge only within an extremely narrow domain, and that they themselves cannot recognize when a question with which they are presented falls outside that domain. One proposal often made for overcoming expert system "fragility" is that they be provided with "common sense." Evaluate this proposal in terms of what it means and in terms of how successful attempts to realize "common sense" within AI systems have been.

D5. Commonly in ***empirical learning***, one divides the full set of examples into a training set and a testing set. Frequently the training set is further subdivided into two subsets, one for standard training and the other for tuning or verification. Briefly explain why this is done. What is the major cost, if any, of doing this subdividing?

Discuss how one could effectively use a tuning set in (a) Quinlan's `ID3` and (b) Rumelhart's `backpropagation` algorithms. Be specific; use high-level pseudo code as much as possible.

D6. Assume one has an approximately-correct "domain theory" for a given task (e.g., automobile diagnosis), as well as a set of classified training examples. Using the taxonomy given by Mitchell, Keller and Kedar-Cabelli, briefly describe three general ways the domain theory could be imperfect.

One might choose to use ***machine learning*** to refine this faulty domain theory. Sketch how the genetic-algorithm approach could be used to achieve this refinement. Be as explicit as possible. (You may assume that the domain theory is expressed as a set of propositional inference rules.)

D7. In ***computer vision***, the occluding contour is an important indicator of the three-dimensional shape of an object that is projected into an image.

(a) Define which surface points generate the *occluding contour* (be as formal as possible), and say how it differs from the *silhouette* of the object.

(b) Briefly describe two (2) different properties of the occluding contour that can be used to obtain qualitative or quantitative information about the surface shape. Be specific about what information is recoverable.

(c) Instead of using the occluding contour from a single image to recover shape properties, what advantages and disadvantages would there be in using properties of the *dynamic* occluding contour as produced by a moving camera in recovering shape information?

D8. "Snakes" have been used in ***computer vision*** for solving a number of problems such as line and curve detection, stereo, and motion tracking.

(a) Define the basic snake model for a parameterized curve $\mathbf{v}(s) = (x(s),y(s))$, where x and y specify the image coordinates.

(b) When is there *not* an *aperture problem* when using snakes for motion tracking?

(c) Briefly describe how snakes could be used for model-based motion tracking of two-dimensional objects. How might this approach be an improvement over other approaches, e.g., P. Burt's hierarchical pattern tree approach? Give your answer in terms of differences in restrictions on the types of objects and/or the types of scenes that could be handled.

D9. You are hired to write a ***computer vision*** program that will recognize several dozen different objects that lie – each nicely separated – on a belt that moves through a factory. This will be incorporated into a robot that will recognize and pick up the appropriate object when needed for some job (e.g., it is a tool to be used, or a part of some finished product to be assembled). Each object can lie on the belt in only 2 or 3, or at most 13, different ways, as is the case for a chisel or crescent wrench (if an object can rest on a curved surface it always settles into only one stable position). You are told to use an expert system shell and a *frame* representation for each object. The shell will then try to match these frames with an input image.

(a) Describe in detail what each such frame would look like, illustrating with two examples of objects (e.g., chisel, crescent-wrench, 2-inch-staple, widget-you-create).

(b) Describe the total system, making sure to explain how these frames are applied, and to what.

(c) Evaluate how good this system will likely be (in terms of errors, speed, ease of programming, potentiality for expanding it to handle larger numbers of objects).

(d) Suggest and explain how you could improve this system if you didn't have to use the shell framework given you.

D10. You are asked to build a ***computer vision*** system to handle the following four problems:

A. Recognize strings of printed letters (as in a line of a book).
B. Recognize spoken speech (as in an audiotape).
C. Recognize strings of written letters (as in a personal diary).
D. Recognize scenes of objects like houses, trees, cars and bikes.

If you could only write one program, and you had to handle only 3 of these 4 problems, describe that program and explain why you chose it and why you chose the 3 problems. NOTE: the speech problem starts with one two-dimensional array that adequately resolves frequency over time; the vision problems start with one grayscale image (that is, there's just one scene, with no motion).

D11. Describe GPSG (Generalized Phrase Structure Grammar) theory and modifications required for its use in ***computational linguistics*** application programs.

D12. Discuss ways in which a ***machine translation*** system might handle semantically-based grammatical concepts requiring representation in the surface structure of a target language, but which do not occur in the source language, e.g., aspect in Russian in an English-to-Russian MT system.