Artificial Intelligence
Answers to CIS Ph.D. Comprehensive Exam., Spring 1997, Closed-Book

1. How may machine learning be viewed as a task involving heuristic search? Identify the states, operators/moves, heuristics, and goal states, preferably in the context of a specific machine learning algorithm.

   ANSWER: Consider the example of the conjunct-dropping approach:

   - Each state is a set of rules of the form

     $\forall x$.Property1(x)=something & Property2(x)=something & ... $\Rightarrow$ Concept(x)

   - The application of a conjunct-dropping operator results in modifying a rule in the rule-set, eliminating one element from the left-hand-side of a rule if it does not violate the constraints suggested by the positive and negative examples. Duplicates of rules are also deleted from the rule-set.
   - Among various possible components of the left-hand-sides of various rules, one possible heuristic is to omit that component which occurs most often in negative examples for the concept.
   - A goal state is a state with the smallest possible number of left-hand-side components in the rules.

2. Consider the sentences

   > The apprentice painted the hull of the ship with a toothbrush.
   > The apprentice painted the ship with green paint after dinner.
   > The apprentice painted the ship with the toothbrush with soft bristles.

   Describe a simple context-free grammar which will yield a correct parse of these sentences. What are all the parses obtained for these sentences? From among these, how can a language recognition system find the semantically meaningful correct parses?

   ANSWER:
   Terminals: the, apprentice, painted, hull, of, ship, with, a, toothbrush, green, paint, after, dinner, soft, bristles.
   Nonterminals: S (sentence/start-symbol), N (noun), V (verb), A (adjective), R (article), P (preposition), N' (noun phrase), V' (verb phrase).
   Grammar rules:

   $$
   \begin{aligned}
   N &\rightarrow \mathit{apprentice} \mid \mathit{hull} \mid \mathit{ship} \mid \mathit{toothbrush} \mid \mathit{paint} \mid \mathit{dinner} \mid \mathit{bristles} \\
   A &\rightarrow \mathit{green} \mid \mathit{soft} \\
   V &\rightarrow \mathit{painted} \\
   R &\rightarrow \mathit{a} \mid \mathit{the} \\
   P &\rightarrow \mathit{with} \mid \mathit{after}
   \end{aligned}
   $$

$$
\begin{aligned}
S &\rightarrow N'V' \\
N' &\rightarrow N \mid AN \mid RN' \mid N'PN' \\
V' &\rightarrow V \mid VN'
\end{aligned}
$$

Possible parses for the sentences using the above grammar are easy to construct. In the last sentence, for instance, two possible parses are suggested by the following parentheses:

> The apprentice painted N'(N'(N'(the ship) P(with) N'(the toothbrush)) P(with) N'(soft bristles)).
> The apprentice painted N'(N'(the ship) P(with) N'(the toothbrush with soft bristles)).

Among these, a language recognition system may choose the more meaningful second parse by methods such as maintaining an association network that contains information such as that "bristles" and "toothbrush" are more strongly associated than "toothbrush" and "ship".

3. (a) Precisely what is meant when it is claimed that resolution provides a proof procedure that is complete?
   ANSWER: Whenever a set of clauses is unsatisfiable, there exists a sequence of resolution steps (from that set of clauses) that leads to the empty clause.

   (b) Why is resolution theorem-proving harder (or easier) when the equality predicate is included?
   ANSWER: Clauses that express the properties of the equality predicated include reflexivity, symmetry, transitivity and substitution of equals by equals for each predicate symbol and function. When these clauses are introduced, the number of possible resolution steps that can be performed increases drastically, making it harder to find the "right" sequence of resolution steps that lead to the empty clause.

   (c) What is the "factor" of a clause in first order logic?
   ANSWER: When two elements of a clause are unifiable, apply the most general unifier to the clause (and eliminate duplicates).

4. Speakers of one language often import grammatical features from the language they first learnt (mother tongue) into a language they learn later. For instance, a speaker of Spanish may speak English words arranged in a manner that would be more reasonable in Spanish than in English. Suggest the design of a language understanding system tailor-made to understand such sequences of English words, e.g., if the system is informed that the speaker's original language is Spanish.
   ANSWER: One possible solution is as follows. Attempt to perform bottom-up parsing for the sentence, using a grammar for English. During the parsing process, identify the combinations of words for which parsing doesn't succeed. For each such combination, attempt to apply a rule from

the Spanish grammar using the same primitive categories (noun, verb, adjective, etc.). If this step is successful, replace the Spanish non-terminal thus obtained by the corresponding English non-terminal, and continue the English bottom-up parsing process.

5. Discuss the main problems to be addressed in three-dimensional object recognition, assuming that algorithms are available to perform two-dimensional exact shape matching invariant to rotation, translation and scaling. Structure your answer, separating issues related to representation, matching, inference, etc.
   ANSWER:
   Representation: 3D representations (such as vertices and edges of polyhedra) need to be stored and accessed, not an easy task.
   Matching: Faces of 3D objects are distorted (e.g. by perspective projections), and their images do not match exactly with the ideal representation of such faces. For instance, a square may appear as a parallelogram.
   Inference: Only some of the faces of an object will be visible in general. Each polygon in the image may match with several possible faces of several different objects. From the set of successfully matching faces, one must infer which object may be present.

6. Answer one of the following:

   (a) Suggest an adaptive algorithm that adds hidden nodes to a feedforward network as needed for a specific problem.
   ANSWER: Here's one outline; more details are expected:

   Start with a small 1-hidden layer network and train it;
   WHILE (system error is unsatisfactory) and (network size not too big), DO:
   Add another hidden node, and train the network until system error converges.

   (b) Suggest an adaptive algorithm that constructs and modifies a probabilistic inference system as new example data become available for a problem.
   ANSWER: Here's one outline; more details are expected:

   Use initially available data to estimate various probabilities and construct a probabilistic inference network;
   WHILE (new data keeps becoming available), DO:

   Update the probabilities associated with connections in the network.