Doctoral Exam in Artificial Intelligence
March 1995

Part I: (5 points each)

Briefly define each of the following terms: (2 or 3 sentences each should suffice.)

A. Non-linear planning

B. Hill climbing

C. Data dependency

D. Horn clause

E. Script application

Part II. Answer 5 of the following 6 problems (15 points each)

1. Let $\mathcal{L}$ be a first-order language containing the following predicates:

   e(O, S) — O is an element of S.
   b$(S1, S2)$ — $S1$ is a subset of $S2$.
   g(X) — X is green.

   Let s1, s2 be constant symbols denoting particular sets.

   A. Express the following facts in $\mathcal{L}$:
      i. If every element of $S1$ is an element of $S2$ then $S1$ is a subset of $S2$.
      ii. Every green object is an element of s2.
      iii. s1 is not a subset of s2.
      iv. Some element of s1 is not green.

   B. Show how (iv) can be proven from (i), (ii), and (iii) using resolution. You must show all the clauses and resolutions used. You need not show the intermediate steps of Skolemization.

2. What is the difference between forward chaining and backward chaining in automated inference? Give an example where forward chaining is more efficient and one where backward chaining is more efficient.

3. Describe the candidate elimination (Version Spaces) algorithm for concept learning.

4. What is the STRIPS representation of action? Name three important limitations on the expressive power of this representation.

5. What is the role of semantic analysis in natural language processing? Illustrate your answer by the sentence, "All birds have a beak". In processing this sentence, what would be the input to the semantic analyzer and what would be the output from it?

6. Edge detection and smoothing are two of the central operations in low-level vision. Explain the significance of each. In what sense does smoothing tend to interfere with edge detection?