CMPSCI 683 Fall 2006 December 18, 2006

Final Exam

120 minutes, closed book. For each question, explain your answer clearly and concisely.

Question 1 (points): Learning and Search:

Suppose you were given a search problem that involved a very large search space where any state could be the initial starting state of the search and any state could be the goal state. Each state could be defined in terms of a relatively small number of attributes and there were also a relatively small number of operators that could be applied to a state. You want to minimize the cost of getting from a specific start state to a specific goal state; the start and goal states are different for different problems. Suppose you did not know how to construct a heuristic function for this search problem. Instead you decided to use a learning approach to construct a heuristic function.

1a (points): Of the learning techniques we have discussed in class choose one that you think is most appropriate and explain in detail how it would be applied to learning a heuristic function for this search problem.

Suppose after the learning is complete, you discovered that the learned heuristic function predicted within 5% of the true cost of the path; This was a very positive result however this learned heuristic function is not an admissible heuristic since it could be a 5% overestimate.

1b (points): How could you modify the learned heuristic so that it could be used directly in an A* search to find the optimal answer.

1c (points): Another way of using this learned heuristic function would be to use it in an Anytime A* search. Explain how this learned heuristic could be used so in Anytime A* so that the optimal answer could be found.

1d (points): Do you think using the learned heuristic in A* or Anytime A* would lead to less search? Explain your reasoning.

1e (points): Suppose in looking at the data on the performance of the learned heuristic that for certain operators the heuristic never overestimated the cost. Could you use that knowledge in the A* search to speed up the search.

Question 2 (points) Decision Networks:

In class we discussed the Mildew example which was defined in the following way.

"Two months before the harvest of a wheat field, the farmer observes the state Q of the crop, and he observes whether it has been attacked by mildew, M. If there is an attack, he will decide on a treatment with fungicides.
There are five variables:
-Q: fair (f), not too bad (n), average (a), good (g)
-M: no (no), little (l), moderate (m), severe (s)\
-H: state of Q plus rotten (r),bad (b), poor (p)
-OQ: observation of Q
OM: observation of M"

This led to the following Decision Network:

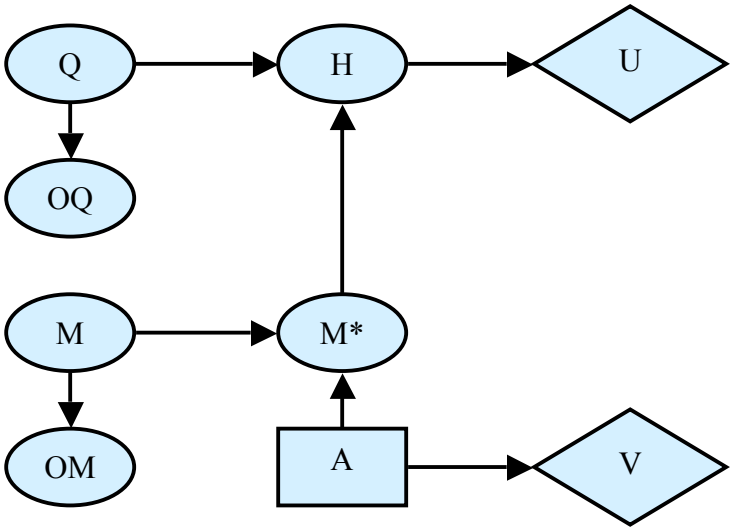

2a (points) Explain why M* is introduced into the decision network

2b (points) Explain what knowledge is required in the problem formulation to detail the Q→OQ link, (Q, M*) → H links, and A → V link

Question 3 (points): Baysean Inference

An admissions committee for a college is trying to determine the probability that an admitted candidate is really qualified; the relevant probabilities are given in the Bayes network shown here. **Calculate p(A|D).**

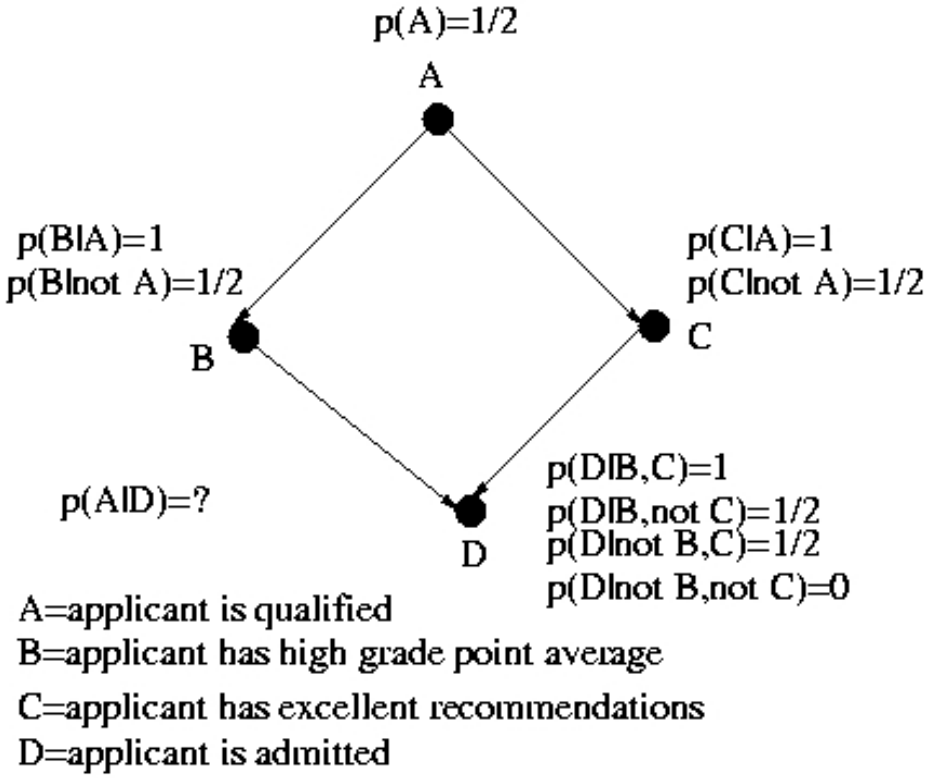

Question 4 (points) Short Questions:

4a (points): Define the Markov assumption. Why does this significantly simplify the necessary information to represent a policy for an MDP? Why, for certain applications, is the size of the MDP state space very large to maintain the Markov assumption.

4b (points): Explain why RL will generally converge faster in environments which there are intermediate rewards, than they would in environments where all the rewards come from terminal states.

4c (points): What are the differences and/or similarities between the credit assignment problem for neural networks and for that of reinforcement learning.

4d (points): How is D-separation used in Baysean Networks to speed up inferencing.

4e (points): Hidden Markov Processes has been used to take a sequence of words and attach the correct part of speech (noun, pronoun, adjective, …) with each word. Discuss the specific knowledge in the HMM that would be needed to do this process.

4f (points): How does the number of hidden nodes in a neural network affect its expressability, robustness and learning time?