ARTIFICIAL INTELLIGENCE:

1. SEARCH. Describe the A* algorithm. Then describe one way in which the heuristic function, h, could be improved online (*during* search), by the algorithm. Be specific and give details.

2. BAYESIAN BELIEF NETWORKS.
(a) In a certain state, 25% of all cars emit excessive amounts of pollutants. If the probability is 0.99 that a car emitting excessive amounts of pollutants will fail the state's vehicular emission test, and the probability is 0.17 that a car not emitting excessive amounts of pollutants will nevertheless fail the test, what is the probability that a car which fails the test actually emits excessive amounts of pollutants?
(b) Using the multiply-connected belief network in Figure 1, manually (on paper) perform three runs of stochastic simulation to estimate $P(\neg E|B)$. Provide *all* details of your first run, as well as *why* you took every step. Just show the basic values for each variable for the other two runs. What is your estimate of $P(\neg E|B)$ after the three runs? How did you get it? Briefly state the principle behind stochastic simulation, i.e., why does it work?

MACHINE LEARNING:

3. DECISION TREES AND NEURAL NETWORKS.

(a) Design and draw the *simplest* decision tree you can make that represents the majority funciton on five Boolean inputs X_1, X_2, X_3, X_4, and X_5. The majority function, $majority(X_1, X_2, X_3, X_4, X_5)$, outputs 1 if at least three inputs are 1; otherwise it outputs 0. NOTE: You don't need to do any entropy calculations to answer this question because all attributes are equivalent.
(b) Design and draw the *simplest* neural network that represents the majority function on five Boolean inputs. Be precise and complete about defining every aspect of your network, including the activation function.
(c) Compare your answers to parts (a) and (b). Which representation (decision tree or neural network) yields a simpler hypothesis? Assume simplicity is measured in terms of the number of nodes. What is it about the one representation that makes it simpler than the other representation for the majority function (i.e., *why* does it result in a simpler hypothesis?

4. REINFORCEMENT LEARNING. Consider playing Tic-Tac-Toe against an opponent who plays randomly. In particular, assume the opponent chooses with uniform probability any open space, unless there is a forced move (in which case it makes the obvious correct move).
(a) Formulate the problem of learning an optimal Tic-Tac-Toe strategy in this case as a Q-learning task. What are the states, transitions, and rewards in this non-deterministic Markov decision process?
(b) Will your program succeed if the opponent plays optimally rather than randomly? Why?