C1 Artificial Intelligence (25 points)

1. **Bayesian inference [15 points]**

   Consider the following Bayesian network:

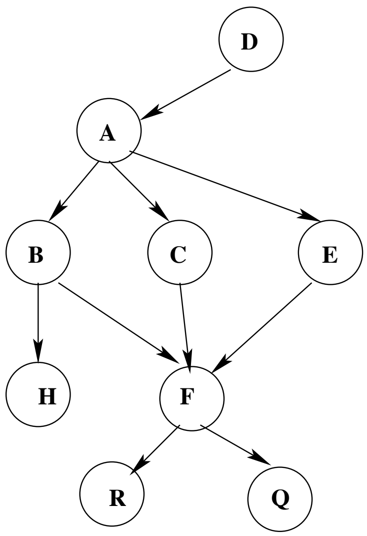

D	P(D)
t	.6
f	.4

D	P(A)
t	.3
f	.4

A	P(B)
t	.2
f	.7

A	P(C)
t	.1
f	.7

A	P(E)
t	.5
f	.7

B	C	E	P(F)
t	t	t	.3
t	t	f	.4
t	f	t	.2
t	f	f	.8
f	t	t	.5
f	t	f	.3
f	f	t	.5
f	f	f	.4

F	P(R)
t	.4
f	.8

F	P(Q)
t	.2
f	.3

B	P(H)
t	.6
f	.4

   What is $P(C = t, F = t \mid A = t, B = t)$?

2. **Decision tree learning [10 points]**

   Suppose you are given eight flowers sampled from two different species of iris and asked to train a decision tree to classify them by petal width and length. The labeled data looks like this:

species	petal width	petal length
virginica	thin	short
virginica	fat	short
virginica	thin	short
virginica	fat	long
versicolor	thin	long
versicolor	fat	short
versicolor	fat	short
versicolor	fat	short

   (a) (7 points) Show the information gain formulas for the two features. (You do not have to calculate information gain; just write out the formulas.)

   (b) (3 points) What feature would be chosen as the root of the decision tree?

C2 Artificial Intelligence (25 points)

1. **Constraint satisfaction [12 points]**

   Two trains, the A-line and the B-line, use the rail network shown below. Each train departs

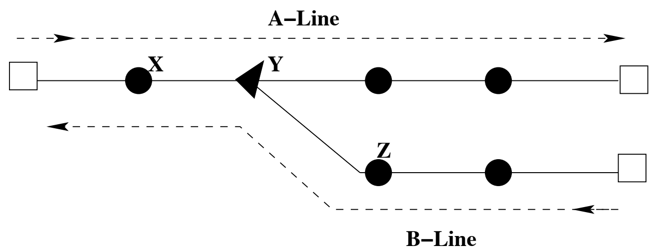

   on the hour, between 1pm and 7pm inclusive. The trains run at the same speeds. There are 8 segments to the track, where a segment is the track between any two dark circles or triangles or unfilled squares; each segment is the same length and it takes 1 hour to cover a segment. The trains cannot pass each other in the part of the track that they share and will collide if they are not properly scheduled. The **only** points on the shared part of the track where trains can pass or touch without colliding are the terminals (represented as squares) and the intersection labelled Y.

   - For example, if the A-line leaves at 4pm and the B-line at 2pm, then at 5pm the A-line will have reached point X and the B-line will have reached the intersection Y. They will thus collide about a half hour later.
   - For example, if the A-line leaves at 2pm and the B-line leaves at 1pm, then at 3pm the A-line will be at point X and the B-line will be at point Z. Thus at 4pm they will both reach the intersection Y where they can safely pass one another.

   (a) (7 points) Consider the scheduling of the trains departure times as a constraint satisfaction problem with the trains as the variables and the departure times as the domains of the variables. State the constraints between the variables as a formal arithmetic inequality (not just a statement that the trains should not collide).

   (b) (5 points) What are the domains of each variable after applying arc consistency?

2. **Search [13 points]**

   You are in a start-up company working on a *shortest time* algorithm to bundle with a global positioning system. Thus you need to determine the path from point A to point B that can be traveled in the shortest amount of time (assuming, of course, that no one breaks the speed limit while driving). One issue is dealing with unknowns such as whether the driver will get stuck at traffic lights, or behind a school bus unloading children, or behind a garbage truck. You therefore want to extend heuristic search to deal with "chance nodes" which represent intermediate points along a route that could add to the time.

   (a) (3 points) What is an admissible heuristic for this problem *without* chance nodes?

   (b) (3 points) Suppose there are two different paths, one of which is more optimal using the heuristic in part (a). What are some chance conditions that might make the other route better in practice?

   (c) (7 points) How could you incorporate chance nodes into your search? Provide an admissible heuristic when chance nodes are involved. Is this heuristic likely to provide the best routes? Why/Why not?

C3 Artificial Intelligence (25 points)

1. **Adversarial search [15 points]**

   Two different programs playing the same game at the same stage produce the game trees shown below. The square nodes are moves for which the evaluation function is known (note that MAX sometimes has only one move, sometimes three).

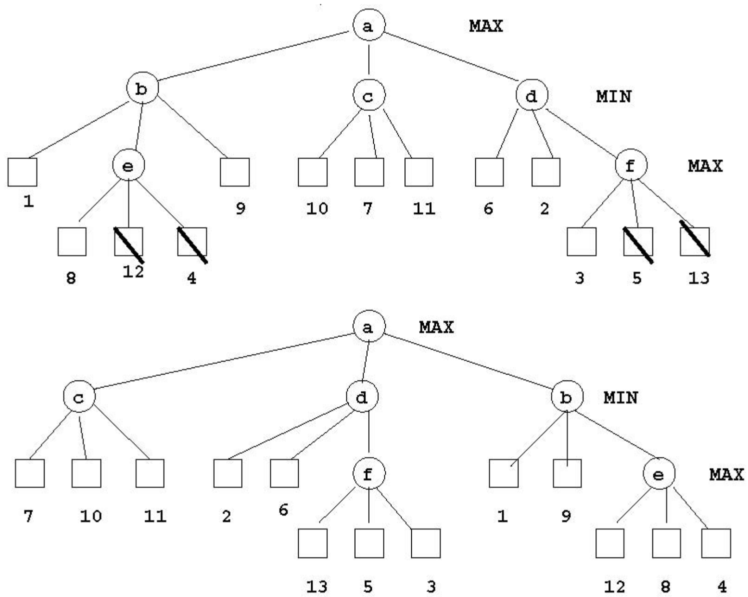

   Notice that the trees differ only in the order in which moves are generated; in other words, the two programs use different move generators.

   (a) (3 points) What score is MAX guaranteed?

   (b) (5 points) The upper tree shows the result of alpha-beta pruning: crossed nodes are not in fact evaluated. Copy the second tree and similarly mark the unevaluated nodes. If there are more or fewer, say why.

   (c) (7 points) Rewrite the tree for a move generator that produces nodes in the worst possible order for alpha-beta pruning. How many nodes are pruned in this case?

2. **Learning [10 points]** ADABOOST is a well-known method for creating an ensemble of learned decision tree models. Suppose that you want to create an ensemble of three decision trees and you have 1000 instances in your training set.

   (a) (2 points) Explain how the first model is created.

   (b) (3 points) Suppose that the first model has an error rate of .25. How is the second model created? You must be precise and give the formulas that are used to treat the instances in the training set differently.

   (c) (2 points) Suppose that the second model has an error rate of .30. How is the third model created? Again, be precise.

   (d) (3 points) Suppose that the third model has an error rate of .10. Given a new instance, how is this instance's class determined? Be precise, giving the formulas that would be used.

C4 Artificial Intelligence (25 points)

1. **Planning [16 points]**

   Consider the following world where there are 3 literals p, q, r and two actions $A1, A2$.

   Initial state: q

   Action $A1$**:**
   Precondition: (none)
   Effect: $p \wedge \neg q$

   Action $A2$**:**
   Precondition: q
   Effect: r

   Goal State: $p \wedge r$

   Here is the planning graph up to level 1 (without marking the mutexes):

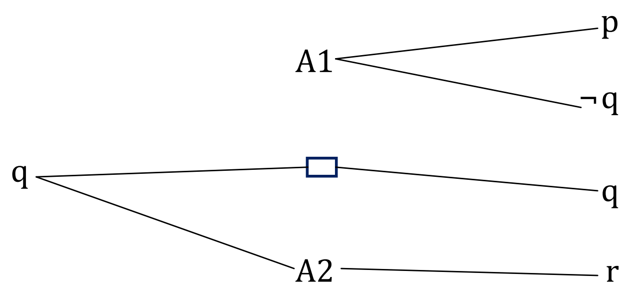

   (a) (6 points) Copy the plangraph (leaving room for part b below) and mark all of the mutexes on it.

   (b) (10 points) Extend the graph to the second level and mark mutexes. Show a valid subgraph of the 2-level planning graph that corresponds to a solution.

2. **Search [9 points]**

   The traveling salesman problem is to visit every city in a set of cities exactly once, returning to the start, covering the shortest distance on the ground. We'd like to set up the search for the solution route as an A* search over cities. That is, we are going to do a best-first search with an evaluation function:

$$F(n) = g(n) + h(n)$$

   where $g(n)$ gives the cost of the path so far, which will just be the total distance the salesman has driven.

   For each of the choices for $h(n)$ below, state whether it is an admissible heuristic and justify your answer. Assume that at city n there are M cities left to visit.

   (a) (3 points) $h(n) = M\times$ the maximum distance between any two unvisited cities.

   (b) (3 points) $h(n) = M\times$ the average distance between all pairs of unvisited cities.

   (c) (3 points) $h(n) = M\times$ the minimum distance between any two unvisited cities.