CS 188 Spring 2011 | Introduction to Artificial Intelligence

Midterm Exam

INSTRUCTIONS

- You have 3 hours.
- The exam is closed book, closed notes except a two-page crib sheet.
- Please use non-programmable calculators only.
- Mark your answers ON THE EXAM ITSELF. If you are not sure of your answer you may wish to provide a *brief* explanation.

First name	
Last name	
SID	
Login	
Name of the person to your left	
Name of the person to your right	

For staff use only:

Total	/??

Q1. [**??** pts] Foodie Pacman

There are two kinds of food pellets, each with a different color (red and blue). Pacman is only interested in tasting the two different kinds of food: the game ends when he has eaten 1 red pellet and 1 blue pellet (though Pacman may eat more than one of each pellet). Pacman has four actions: moving up, down, left, or right, and does not have a "stay" action. There are K red pellets and K blue pellets, and the dimensions of the board are N by M.

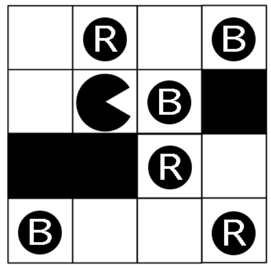

$K = 3$, $N = 4$, $M = 4$

(a) [**??** pts] Give an efficient state space formulation of this problem. Specify the domain of each variable in your state space.

(b) [**??** pts] Give a tight upper bound on the size of the state space.

(c) [**??** pts] Give a tight upper bound on the branching factor of the search problem.

(d) [**??** pts] Assuming Pacman starts the game in position (x,y), what is the initial state?

(e) [**??** pts] Define a goal test for the problem.

(f) [**??** pts] For each of the following heuristics, indicate (yes/no) whether or not it is admissible (a correct answer is worth 1 point, leaving it blank is worth 0 points, and an incorrect answer is worth -1 points).

Heuristic	Admissible?
The number of pellets remaining	
The smallest Manhattan distance to any remaining pellet	
The maximum Manhattan distance between any two remaining pellets	
The minimum Manhattan distance between any two remaining pellets of opposite colors	

Q2. [?? pts] Expectimax

Your little brother Timmy has a birthday and he was promised a toy. However, Timmy has been misbehaving lately and Dad thinks he deserves the least expensive present. Timmy, of course, wants the most expensive toy. Dad will pick the city from which to buy the toy, Timmy will pick the store and you get to pick the toy itself. You don't want to take sides so you decide to pick a toy at random. All prices (including X and Y) are assumed to be nonnegative.

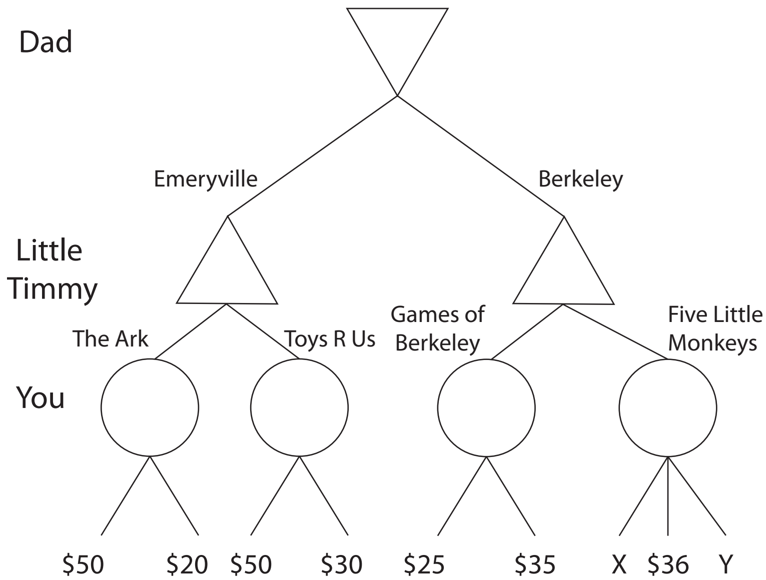

(a) [?? pts] Fill in the values of all the nodes that don't depend on X or Y.

(b) [?? pts] What values of X will make Dad pick Emeryville regardless of the price of Y?

(c) [?? pts] We know that Y is at most $30. What values of X will result in a toy from Games of Berkeley regardless of the exact price of Y?

(d) [?? pts] Normally, alpha-beta pruning is not used with expectimax. However, with some additional information, it is possible to do something similar. Which **one** of the following conditions on a problem are required to perform pruning with expectimax?

1. The children of the expectation node are leaves.
2. All values are positive.
3. The children of the expectation node have specified ranges.
4. The child to prune is last.

Q3. [**??** pts] Forced Random Policy in MDP

(a) [**??** pts] Assume you are asked to act in a given MDP $(S, A, T, R, \gamma, s_0)$. However, rather than being able to freely choose your actions, at each time step you must start by flipping a coin. If the coin lands heads, then you can freely choose your action. If the coin lands tails, however, you don't get to choose an action and instead an action will be chosen for you uniformly at random from the available actions. Can you specify a modified MDP $(S', A', T', R', \gamma', s'_0)$ for which the optimal policy maximizes the expected discounted sum of rewards under the specified restrictions on your ability to choose actions? (Hint: you may not need to change all entities in the MDP.)

$S' =$

$A' =$

$T' =$

$R' =$

$\gamma' =$

$s'_0 =$

Q4. [**??** pts] Search

(a) [**??** pts] The following implementation of graph search may be incorrect. Circle all the problems with the code.

```
function GRAPH-SEARCH(problem, fringe)
    closed ← an empty set,
    fringe ← INSERT(MAKE-NODE(INITIAL-STATE[problem]), fringe)
    loop
        if fringe is empty then
            return failure
        end if
        node ← REMOVE-FRONT(fringe)
        if GOAL-TEST(problem,STATE[node]) then
            return node
        end if
        ADD STATE[node] TO closed
        fringe ← INSERTALL(EXPAND(node,problem), fringe)
    end loop
end function
```

1. Nodes may be expanded twice.
2. The algorithm is no longer complete.
3. The algorithm could return an incorrect solution.
4. None of the above.

(b) [**??** pts] The following implementation of A* graph search may be incorrect. You may assume that the algorithm is being run with a consistent heuristic. Circle all the problems with the code.

```
function A*-SEARCH(problem, fringe)
    closed ← an empty set
    fringe ← INSERT(MAKE-NODE(INITIAL-STATE[problem]), fringe)
    loop
        if fringe is empty then
            return failure
        end if
        node ← REMOVE-FRONT(fringe)
        if STATE[node] IS NOT IN closed then
            ADD STATE[node] TO closed
            for successor IN GETSUCCESSORS(problem, STATE[node]) do
                fringe ← INSERT(MAKE-NODE(successor), fringe)
                if GOAL-TEST(problem,successor) then
                    return successor
                end if
            end for
        end if
    end loop
end function
```

1. Nodes may be expanded twice.
2. The algorithm is no longer complete.
3. The algorithm could return an incorrect solution.
4. None of the above.

Q5. [**??** pts] Probability

(a) [**??** pts] Consider the random variables A, B, and C. Circle all of the following equalities that are **always** true, if any.

1. $\mathbf{P}(A,B) = \mathbf{P}(A)\mathbf{P}(B) - \mathbf{P}(A|B)$
2. $\mathbf{P}(A,B) = \mathbf{P}(A)\mathbf{P}(B)$
3. $\mathbf{P}(A,B) = \mathbf{P}(A|B)\mathbf{P}(B) + \mathbf{P}(B|A)\mathbf{P}(A)$
4. $\mathbf{P}(A) = \sum_{b \in B} \mathbf{P}(A|B=b)\mathbf{P}(B=b)$
5. $\mathbf{P}(A,C) = \sum_{b \in B} \mathbf{P}(A|B=b)\mathbf{P}(C|B=b)\mathbf{P}(B=b)$
6. $\mathbf{P}(A,B,C) = \mathbf{P}(C|A)\mathbf{P}(B|C,A)\mathbf{P}(A)$

Now assume that A and B both can take on only the values true and false ($A \in \{\text{true}, \text{false}\}$ and $B \in \{\text{true}, \text{false}\}$). You are given the following quantities:

$$\begin{aligned} \mathbf{P}(A = \text{true}) &= \tfrac{1}{2} \\ \mathbf{P}(B = \text{true} \mid A = \text{true}) &= 1 \\ \mathbf{P}(B = \text{true}) &= \tfrac{3}{4} \end{aligned}$$

(b) [**??** pts] What is $\mathbf{P}(B = \text{true} \mid A = \text{false})$?

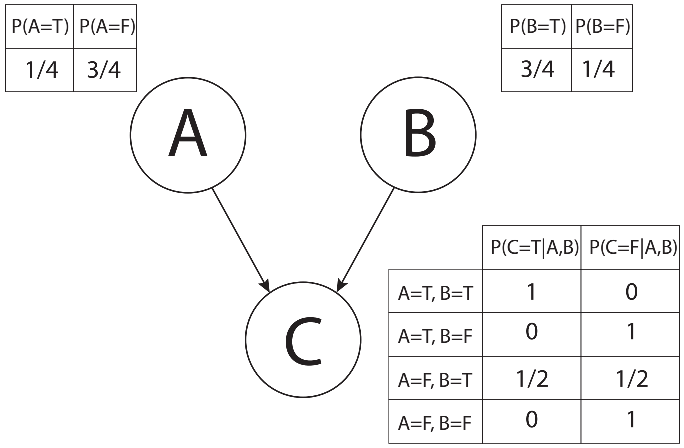

P(A=T)	P(A=F)
1/4	3/4

P(B=T)	P(B=F)
3/4	1/4

	P(C=T|A,B)	P(C=F|A,B)
A=T, B=T	1	0
A=T, B=F	0	1
A=F, B=T	1/2	1/2
A=F, B=F	0	1

(c) [**??** pts] Give the formula for the joint probability distribution induced by the above Bayes Net:

$\mathbf{P}(A, B, C) =$

Compute the values of the following probabilities:

(d) [**??** pts]

$\mathbf{P}(C = T) =$

(e) [**??** pts]

$\mathbf{P}(A = T, B = T) =$

(f) [**??** pts]

$\mathbf{P}(A = T, B = T | C = T) =$

Q6. [?? pts] Crossword Puzzles as CSPs

You are developing a program to automatically solve crossword puzzles, because you think a good income source for you might be to submit them to the New York Times (\$200 for a weekday puzzle, \$1000 for a Sunday).[1] For those unfamiliar with crossword puzzles, a crossword puzzle is a game in which one is given a grid of squares that must be filled in with intersecting words going from left to right and top to bottom. There are a given set of starting positions for words (in the grid below, the positions $1, 2, 3, 4$, and 5), where words must be placed going across (left to right) or down (top to bottom). At any position where words intersect, the letters in the intersecting words must match. Further, no two words in the puzzle can be identical. An example is the grid below, in which the down words (1, 2, and 3) are `DEN`, `ARE`, and `MAT`, while the across words (1, 4, and 5) are `DAM`, `ERA`, and `NET`.

Example Crossword Grid and Solution

[1]D	[2]A	[3]M
[4]E	R	A
[5]N	E	T

A part of your plan to make crosswords, you decide you will create a program that uses the CSP solving techniques you have learned in CS 188, since you want to make yourself obsolete at your own job from the get-go. Your first task is to choose the representation of your problem. You start with a dictionary of all the words you could put in the crossword puzzle, where the dictionary is of size K and consists of the words $\{d_1, d_2, \ldots, d_K\}$. Assume that you are given a grid with N empty squares and M different entries for words (and there are 26 letters in the English language). In the example above, $N = 9$ and $M = 6$ (three words across and three words down).

You initially decide to use words as the variables in your CSP. Let D_1 denote the first down word, D_2 the second, D_3 the third, etc., and similarly let A_k denote the kth across word. For example, in the crossword above, $A_1 = \texttt{DAM}$, $D_1 = \texttt{DEN}$, $D_2 = \texttt{ARE}$, and so on. Let $D_1[i]$ denote the letter in the ith position of the word D_1.

(a) [**??** pts] What is the size of the state space for this CSP?

(b) [**??** pts] Precisely (i.e. use mathematical notation to) describe the constraints of the CSP when we use words as variables.

[1] `http://www.nytimes.com/2009/07/19/business/media/19askthetimes.html`

After defining your CSP, you decide to go ahead and make a small crossword using the grid below. Assume that you use the words on the right as your dictionary.

Crossword Grid

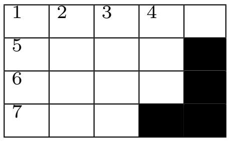

Dictionary Words

ARCS, BLAM, BEAR, BLOGS, LARD, LARP,
GAME, GAMUT, GRAMS, GPS, MDS, ORCS, WARBLER

(c) [**??** pts] Enforce all *unary* constraints by crossing out values in the table below.

D_1	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_2	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_3	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_4	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_1	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_5	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_6	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_7	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER

Here's an extra table in case you make a mistake:

D_1	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_2	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_3	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_4	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_1	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_5	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_6	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_7	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER

(d) [**??** pts] Assume that in backtracking search, we assign A_1 to be GRAMS. Enforce unary constraints, and in addition, cross out all the values eliminated by forward checking against A_1 as a result of this assignment.

D_1	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_2	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_3	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_4	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_1	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_5	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_6	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_7	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER

Here's an extra table in case you make a mistake:

D_1	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_2	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_3	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_4	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_1	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_5	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_6	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_7	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER

(e) [**??** pts] Now let's consider how much arc consistency can prune the domains for this problem, even when no assignments have been made yet. I.e., assume no variables have been assigned yet, enforce unary constraints first, and then enforce arc consistency by crossing out values in the table below.

D_1	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_2	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_3	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_4	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_1	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_5	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_6	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_7	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER

Here's an extra table in case you make a mistake:

D_1	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_2	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_3	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
D_4	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_1	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_5	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_6	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER
A_7	ARCS	BLAM	BEAR	BLOGS	LARD	LARP	GPS	MDS	GAME	GAMUT	GRAMS	ORCS	WARBLER

(f) [**??** pts] How many solutions to the crossword puzzle are there? Fill them (or the single solution if there is only one) in below.

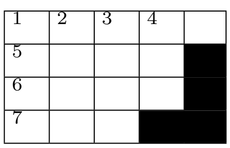

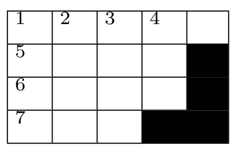

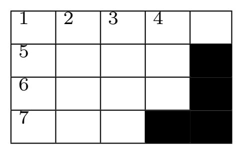

Your friend suggests using letters as variables instead of words, thinking that sabotaging you will be funny. Starting from the top-left corner and going left-to-right then top-to-bottom, let X_1 be the first letter, X_2 be the second, X_3 the third, etc. In the very first example, $X_1 = \texttt{D}$, $X_2 = \texttt{A}$, and so on.

(g) [**??** pts] What is the size of the state space for this formulation of the CSP?

(h) [**??** pts] Assume that in your implementation of backtracking search, you use the least constraining value heuristic. Assume that X_1 is the first variable you choose to instantiate. For the crossword puzzle used in parts (c)-(f), what letter(s) might your search assign to X_1?

Q7. [**??** pts] Short Answer

Each true/false question is worth 1 point. Leaving a question blank is worth 0 points. **Answering incorrectly is worth -1 point.**

(a) Assume we are running A^* graph search with a consistent heuristic h. Assume the optimal cost path to reach a goal has a cost c^*. Then we have that

- **(i)** [*true* or *false*] All nodes n reachable from the start state satisfying $g(n) < c^*$ will be expanded during the search.
- **(ii)** [*true* or *false*] All nodes n reachable from the start state satisfying $f(n) = g(n) + h(n) < c^*$ will be expanded during the search.
- **(iii)** [*true* or *false*] All nodes n reachable from the start state satisfying $h(n) < c^*$ will be expanded during the search.

(b) Running A^* *graph search* with an inconsistent heuristic can lead to suboptimal solutions. Consider the following modification to A^* graph search: replace the closed list with a cost-sensitive closed list, which stores the f-cost of the node along with the state ($f(n) = g(n) + h(n)$).

Whenever the search considers expanding a node, it first verifies whether the node's state is in the cost-senstive closed list and only expands it if either (a) the node's state is not in the cost-sensitive closed list, or (b) the node's state is in the cost-sensitive closed list with a higher f-cost than the f-cost for the node currently considered for expansion.

If a node is expanded because it meets criterion (a), its state and f-cost get added to the cost-sensitive closed list; if it gets expanded because it meets criterion (b), the cost associated with the node's state gets replaced by the current node's f-cost. Which of the following statements are true about the proposed search procedure?

- **(i)** [*true* or *false*] The described search procedure finds an optimal solution if h is admissible.
- **(ii)** [*true* or *false*] The described search procedure finds an optimal solution if h is consistent.
- **(iii)** [*true* or *false*] Assuming h is admissible (but possibly inconsistent), the described search procedure will expand no more nodes than A^* tree search.
- **(iv)** [*true* or *false*] Assuming h is consistent, the described search procedure will expand no more nodes than A^* graph search.

(c) Let H_1 and H_2 both be admissible heuristics.

- **(i)** [*true* or *false*] $\max(H_1, H_2)$ is necessarily admissible
- **(ii)** [*true* or *false*] $\min(H_1, H_2)$ is necessarily admissible
- **(iii)** [*true* or *false*] $(H_1 + H_2)/2$ is necessarily admissible
- **(iv)** [*true* or *false*] $\max(H_1, H_2)$ is necessarily consistent

(d) Let H_1 be an admissible heuristic, and let H_2 be an inadmissible heuristic.

- **(i)** [*true* or *false*] $\max(H_1, H_2)$ is necessarily admissible
- **(ii)** [*true* or *false*] $\min(H_1, H_2)$ is necessarily admissible
- **(iii)** [*true* or *false*] $(H_1 + H_2)/2$ is necessarily admissible
- **(iv)** [*true* or *false*] $\max(H_1, H_2)$ is necessarily consistent

(e) For Markov Decisions Processes (MDPs), we have that:

- **(i)** [*true* or *false*] A small discount (close to 0) encourages shortsighted, greedy behavior.
- **(ii)** [*true* or *false*] A large, negative living reward ($\ll 0$) encourages shortsighted, greedy behavior.
- **(iii)** [*true* or *false*] A negative living reward can always expressed using a discount < 1.
- **(iv)** [*true* or *false*] A discount < 1 can always be expressed as a negative living reward.

(f) You are given a game-tree for which you are the maximizer, and in the nodes in which you don't get to make a decision an action is chosen uniformly at random amongst the available options. Your objective is to maximize the probability you win \$10 or more (rather than the usual objective to maximize your expected value). Then:

- **(i)** [*true* or *false*] Running expectimax will result in finding the optimal strategy to maximize the probability of winning \$10 or more.
- **(ii)** [*true* or *false*] Running minimax, where chance nodes are considered minimizers, will result in finding the optimal strategy to maximize the probability of winning \$10 or more.
- **(iii)** [*true* or *false*] Running expectimax in a modified game tree where every pay-off of \$10 or more is given a value of 1, and every pay-off lower than \$10 is given a value of 0 will result in finding the optimal strategy to maximize the probability of winning \$10 or more.
- **(iv)** [*true* or *false*] Running minimax in a modified game tree where every pay-off of \$10 or more is given a value of 1, and every pay-off lower than \$10 is given a value of 0 will result in finding the optimal strategy to maximize the probability of winning \$10 or more.

(g) Assume we run $\alpha - \beta$ pruning expanding successors from left to right on a game with tree as shown in Figure **??** (a). Then we have that:

- **(i)** [*true* or *false*] For some choice of pay-off values, no pruning will be achieved (shown in Figure **??** (a)).
- **(ii)** [*true* or *false*] For some choice of pay-off values, the pruning shown in Figure **??** (b) will be achieved.
- **(iii)** [*true* or *false*] For some choice of pay-off values, the pruning shown in Figure **??** (c) will be achieved.
- **(iv)** [*true* or *false*] For some choice of pay-off values, the pruning shown in Figure **??** (d) will be achieved.
- **(v)** [*true* or *false*] For some choice of pay-off values, the pruning shown in Figure **??** (e) will be achieved.
- **(vi)** [*true* or *false*] For some choice of pay-off values, the pruning shown in Figure **??** (f) will be achieved.

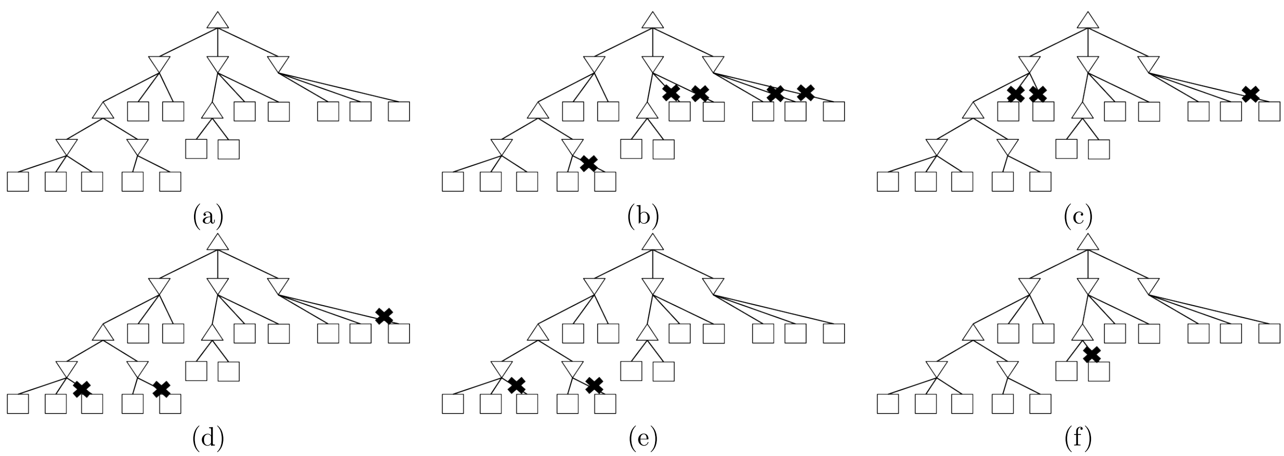

Figure 1: Game trees.

(h) Assume a probability distribution P over binary random variables A, B, C is given. Assume also a probability distribution Q is given which is defined by the Bayes net shown in Figure **??**, with conditional probability tables such that $Q(A) = P(A)$, $Q(B|A) = P(B|A)$, and $Q(C|A) = P(C|A)$. Then we have that:

- **(i)** [*true* or *false*] $\forall a, \quad Q(A = a) = P(A = a)$
- **(ii)** [*true* or *false*] $\forall b, \quad Q(B = b) = P(B = b)$
- **(iii)** [*true* or *false*] $\forall c, \quad Q(C = c) = P(C = c)$
- **(iv)** [*true* or *false*] $\forall a, b, c \quad Q(A = a, B = b, C = c) = P(A = a, B = b, C = c)$

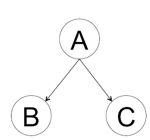

Figure 2: Bayes net