CMSC 671 (Introduction to AI) – Fall 2024

Homework 3: Decision making, Multi-agent systems, Decision trees.

Turnin: Blackboard.

Please submit **all parts** together as a **single PDF file** named *lastname*_hw3.pdf, with parts clearly marked and delineated. This assignment should be worked on individually.

All files must have your full name(s) in the file, at/near the top.

PART I. PRUNING (10 PTS.)

Consider the 3-ply game tree below, in which the values assigned by the evaluation function are shown at the leaves:

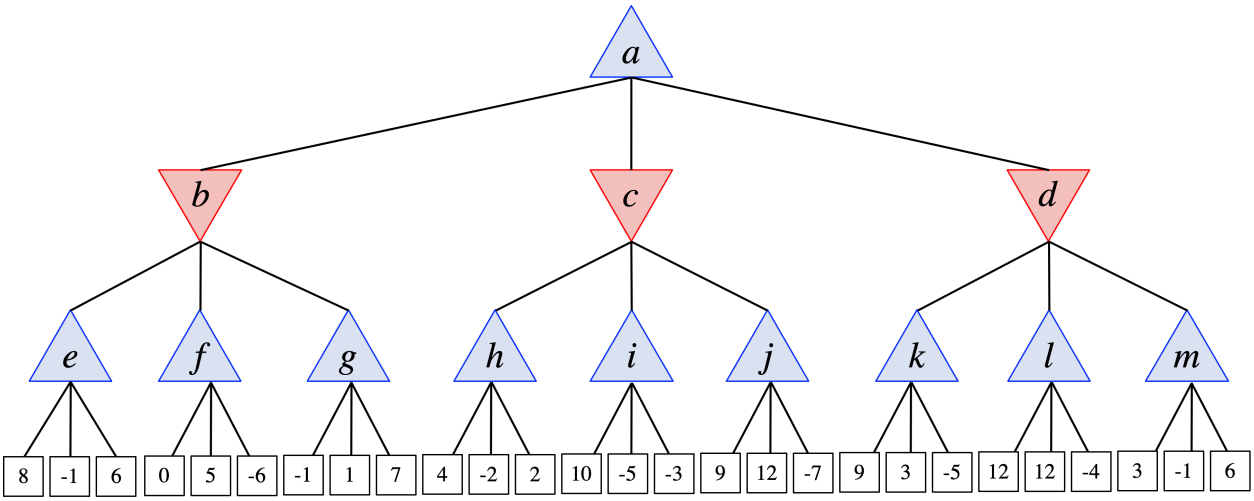

1. Suppose MAX uses alpha-beta pruning. Write down the values assigned to each node at the end of search. If a node doesn't receive a value, mark it null (∅). *(5 pts.)*

a=	*e*=	*i*=	*m*=
b=	*f*=	*j*=	
c=	*g*=	*k*=	
d=	*h*=	*l*=	

2. What nodes will be pruned by the alpha-beta algorithm? *(3 pts.)*

3. What move will MAX choose? (*b*, *c*, or *d*?) *(2 pts.)*

PART II. FILTERING (20 PTS.)

A student is working the late shift at a local movie theater. Although they don't know whether the movie playing is popular (***M***), one of their tasks is to refill the kernel hopper for the popcorn machine; they reason that if the hopper is empty (***E***), the theater sold a lot of popcorn, so it's more likely that the movie is popular. They also reason that movies tend to decrease in popularity the longer they are showing, and that 75% of movies are popular on opening night (that is, day 0). This is represented as the following model:

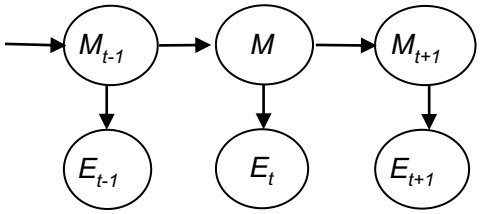

M_{t-1}	$P(M_t \mid M_{t-1})$
T	0.75
F	0.12

E_t	$P(E_t \mid M_t)$
T	0.9
F	0.1

4. Showing all work: What is the probability that a movie is popular <u>*on day 3*</u>, given that the hopper is full on day 2 and empty on day 3? *(20 pts.)*

PART II. MULTI-AGENT SYSTEMS (10 POINTS)

Assignment: For each of the two normal-form two-player games below, identify a Nash equilibrium (if there is one). Explain why it's a Nash equilibrium for that game, or why no Nash equilibria exist if this is the case.

5. Jason and Lakshmi are going to eat dinner. Jason prefers pizza, while Lakshmi would prefer burgers. However, they would both prefer to eat together rather than separately. *(5 pts.)*

Lakshmi \ Jason	*Burgers*	*Pizza*
Burgers	2, 1	0, 0
Pizza	0, 0	1, 2

6. Consider the zero-sum two-player game Evens. Each player has a die, on which they each secretly choose a number; they then simultaneously reveal their numbers. If both numbers are even ***or*** both numbers are odd, Player 1 receives a dollar from Player 2. If the numbers are split (one even and one odd), Player 2 receives a dollar from Player 1. *(5 pts.)*

Player 1 \ Player 2	*Even*	*Odd*
Even	1, -1	-1, 1
Odd	-1, 1	1, -1

PART IV. DECISION TREE LEARNING (30 PTS.)

Should you go out and play soccer with your friends, or is the weather too unpleasant? You consider what decisions you have made in the past…

Temp (*T*)	Humidity (*H*)	Wind (*W*)	Decision
Hot	Humid	Weak	No soccer
Hot	Humid	Strong	No soccer
Pleasant	Humid	Weak	No soccer
Cold	Low	Weak	Soccer
Pleasant	Low	Strong	Soccer

7. Information Gain (30 pts.)

a. At the root node for your decision tree in this domain, what is the information gain associated with a split on the attribute *T*? *(10 pts.)*

b. Build a decision tree for this dataset, using information gain to decide what attributes to split on at each level. The tree should show which attribute you split on for each branch, and show the decisions (class predictions) at the leaves. Show your tree and all calculations. *(20 pts.)*