CMSC 671 (Introduction to AI)

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 start with your last name(s) and have your full name(s) in the file, at/near the top.

Reminder: Assignments **must be turned in on time. If Blackboard says it's late, it is late.**

PART I. 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:

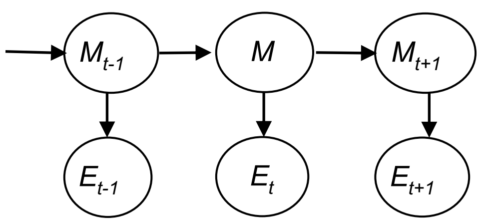

M_{t-1}	$P(M_t \mid M_{t-1})$
T	0.75
F	0.12

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

1. 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?

PART II. MULTI-AGENT SYSTEMS (10 POINTS)

Assignment: For each of the two normal-form two-player games below, identify the Nash equilibria (if there are any). Explain why these strategy sets are the Nash equilibria of the game, or why no Nash equilibria exist if this is the case.

2. Chicken: Two drivers are headed for a one-lane pass. If they both swerve out of the way, they "tie" (nobody scores). If one swerves and the other drives straight on, the one who swerved loses a point and the other gains a point. If neither swerves, they both lose big.

	Straight	**Swerve**
Straight	-10, -10	+1, -1
Swerve	-1, +1	0, 0

3. Rock-Paper-Scissors: each of two players chooses one of Rock, Paper, or Scissors, then the players reveal their choices to each other simultaneously. Rock beats scissors, scissors beats paper, and paper beats rock, as shown.

	R	**P**	**S**
R	0, 0	-1,+1	+1, -1
P	+1, -1	0, 0	-1, +1
S	-1, +1	+1, -1	0, 0

Part III. Learning in the wild (15 pts.)

Assignment: Consider the problem faced by a small child learning to solve a shape sorter puzzle. *(~300 words)*

- Explain how this problem fits into the general *learning model*.
- Describe the percepts (sensor inputs) and actions of the child, and the types of learning the child must do.
- Describe the subfunctions the child is trying to learn in terms of inputs, outputs, and available training data.

Part IV. Decision tree learning (55 pts.)

It's Halloween, and you're trying to figure out whether to give out candy, and if so, whether to give out the candy you're thinking of or choose something else to give out (so there are three possible outcomes). For each type of candy, you know whether you already have some, how far you would have to go to obtain more, and how much candy of that kind was left over in previous years, out of 100 starting pieces. You start by polling your friends about what their decision would be, resulting in the chart shown.

Assignment: Answer the questions using this data:

Candy Decisions

Type	Do I already have some? (*H*)	Where can I buy it? (*B*)	Do I or my roommates like it? (*R*)	Left over (*L*)	Class
Candy Corn	No	Very far	No	80	No candy
Chewing Gum	Yes	Far	Yes	40	Something else
Good'n'Plenty	No	Far	No	90	No candy
Gummi Bears	Yes	Near	Yes	0	Give this out
Hershey Kisses	Yes	Far	No	70	Give this out
Jelly Beans	No	Far	Yes	80	Something else
Jolly Ranchers	Yes	Very far	Yes	30	No candy
Kit-Kats	Yes	Very far	No	60	No candy
Lollipops	Yes	Near	No	20	Give this out
M&Ms	Yes	Near	Yes	50	Something else
Nerds	Yes	Far	No	40	Give this out
Pocky	No	Very far	Yes	10	Something else
Reese's Pieces	No	Near	Yes	70	Something else

4. **Information Gain (20 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 *B*? *(10 pts.)*

   b. What would it be for a split at the root on the attribute *L*? (Use a threshold of 50 for *L* (i.e., assume a binary split: $L \leq 50$). *(10 pts.)*

5. **Gain Ratios (15 pts.)**

   a. Again at the root node, what is the *gain ratio* associated with the attribute *B*?

   b. What is the gain ratio for the *L* attribute at the root (using the same threshold as in 4b)?

6. **Decision Tree (20 pts.)**

   a. Suppose you build a decision tree that splits on the *H* attribute at the root node. How many child nodes are there are at the *first level* of the decision tree? *(2 pts.)*

   b. After *H*, which branches require a further split? *(2 pts.)*

   c. Draw the smallest (fewest nodes) decision tree you can you construct for this dataset. The tree should show which attribute you split on for each branch, and show the decisions (class predictions) at the leaves. *(8 pts.)*

   d. What method(s) did you use to find that tree? Show all calculations. *(8 pts.)*