

1 Game Theory Review

(a) What are the differences between extensive form and normal form games?

(b) We can practice computing utilities using this example that was given in lecture:

CRAM	DO HW	PLAY GAME	
98	100	85	P(EASY) = .2
97	90	65	P(HARD) = .8

- What is the utility of the pure strategy: cram?
- What is the utility of the pure strategy: do HW?
- What is the utility of the mixed strategy: $\frac{1}{2}$ cram, $\frac{1}{2}$ do HW?

(c) Consider a two player game, where each player must simultaneously choose a number from $\{2, 3, \dots, 99, 100\}$. Let x_1 represent the value chosen by player 1, and x_2 represent the value chosen by player 2. The rules of the game are such that the utility for a player 1 can be given as:

$$u(p_1) = \begin{cases} x_1 & x_1 = x_2 \\ x_2 - 2 & x_1 > x_2 \\ x_2 + 2 & x_1 < x_2 \end{cases}$$

Because the rules of the game for everyone are the same, the utility function for player 2 is symmetric to $u(p_1)$. Does there exist a pure Nash Equilibrium for this game? It may help to try to play a few rounds of this game with someone next to you.

(d) Describe the following voting strategies:

- Plurality

- Borda Count

- Single Transferable Vote

- Pairwise Elections

- Plurality with Runoff

- Condorcet Winner

(e) Which voting rule (Plurality, Borda Count, Both, Neither) is Condorcet consistent? Consider the following voting example:

3 voters	2 voters
a	b
b	c
c	a

3 voters	2 voters	2 voters
a	b	c
b	c	b
c	a	a

(f) Which voting rule is the best?

2 Voting Rules

You and your friends are deciding on a movie to watch on Netflix.

Answer the following questions given this preference profile. The first column of the table means that 10 voters put *Arrival* as their first choice, *Black Panther* as their second choice, etc.

Category 1	Category 2	Category 3
10 voters	35 voters	45 voters
<i>Arrival</i>	<i>Frankenstein</i>	<i>Category 3</i>
<i>Black Panther</i>	<i>Charlie and the Chocolate Factory</i>	<i>Frankenstein</i>
<i>Charlie and the Chocolate Factory</i>	<i>Deep Blue Sea</i>	<i>Deep Blue Sea</i>
<i>Elmo's Christmas Countdown</i>	<i>Black Panther</i>	<i>Arrival</i>
<i>Frankenstein</i>	<i>Elmo's Christmas Countdown</i>	<i>Charlie and the Chocolate Factory</i>
<i>Deep Blue Sea</i>	<i>Arrival</i>	<i>Black Panther</i>

Which movie is the winner under the following voting strategies? Assume ties are broken in alphabetical order, e.g. we would prefer *Black Panther* over *Charlie and the Chocolate Factory* if they are tied.

- (a) Plurality
- (b) Borda Count
- (c) Single Transferable Vote (STV)
- (d) Condorcet Winner


Consider a new voting rule. Under this rule, every voter's top three preferences get one point each. The movie with the most points wins. Assume ties are broken in alphabetical order, e.g. *Black Panther* over *Charlie and the Chocolate Factory*.

- (e) Which movie is the winner under this new voting rule (i.e. which movie will everyone watch)?
- (f) Can a player in the first category manipulate this new voting to make another movie the unique winner (without tie-breaking)? If yes, what is his reported preference according to the greedy algorithm? If not, briefly explain why.
- (g) Which properties among the following does this new voting rule satisfy?

Majority, Consistency, Condorcet Consistency, Strategy-proof, Dictatorial, Constant, Onto

3 Sampling Review

Consider the following Bayes Net and corresponding probability tables.



M	R	W	P(M R,W)
+m	+r	+w	0.1
-m	+r	+w	0.9
+m	+r	-w	0.45
-m	+r	-w	0.55
+m	-r	+w	0.35
-m	-r	+w	0.65
+m	-r	-w	0.9
-m	-r	-w	0.1

Consider the case where we are sampling to approximate the query $P(R | +m)$.

Fill in the following table with the probabilities of each respective sample given that we are using each of the following sampling techniques.

Method	$\langle +r, -w, +m \rangle$	$\langle +r, +w, -m \rangle$
Prior Sampling		
Likelihood Weighting		

4 HMMs and Particle Filtering Review

Consider the following hidden Markov model with a binary hidden state X . The transition probabilities and initial distribution are:

X_0	$P(X_0)$	X_t	X_{t+1}	$P(X_{t+1} X_t)$
0	0.5	0	0	0.9
1	0.5	0	1	0.1
		1	0	0.5
		1	1	0.5

- (a) After one timestep, what is the new belief distribution $P(X_1)$?

X_1	$P(X_1)$
0	
1	

Now, we incorporate sensor readings as our observations. The sensor model is parameterized by some value $\beta \in [0, 1]$:

X_t	E_t	$P(E_t X_t)$
0	0	β
0	1	$1 - \beta$
1	0	$1 - \beta$
1	1	β

- (b) At $t = 1$, we get the first sensor reading, $E_1 = 0$. Find $P(X_1 = 0|E_1 = 0)$ in terms of β .


- (c) For what range of values of β will a sensor reading $E_1 = 0$ increase our belief that $X_1 = 0$? In other words, what is the range of β for which $P(X_1 = 0|E_1 = 0) > P(X_1 = 0)$?

- (d) Now, we want to use particle filtering to predict what state value our model currently assumes. At time t , there are 2 particles in state value 0, and 3 particles in state value 1. What is the prior belief distribution $\hat{P}(X_t)$?

X_t	$\hat{P}(X_t)$
0	
1	

- (e) At some time t , we receive our first sensor reading $E_t = 1$. Given $\beta = 0.6$ and the previous table for $P(E_t|X_t)$, how many particles will be in each state value after updating our belief and resampling? When resampling, use this list of numbers as a source of randomness: $[0.182, 0.703, 0.471, 0.859, 0.382]$ and fix the order of states to be $X_t = 0, X_t = 1$.

- (f) Suppose we now have the following modified HMM structure, in which the hidden variables now have a parent variable Y_t , starting at $t = 1$:



Write expressions for answering the following queries. Make sure your expressions are solely in terms of the probability tables from the HMM, and that they are in the simplest possible form (hint: conditional independence!). You must explicitly write out any normalization constants.

- (i) $P(X_1 | E_1)$
- (ii) $P(Y_1 | X_1, E_1)$
- (iii) $P(Y_1 | E_1)$
- (iv) $P(Y_2 | E_1, E_2)$

