


1 Game Theory Review

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

An extensive form game is one where every player is represented, as well as every opportunity that each of those players has to make a move, what moves they can make at that time, what information they have at the time, and the possible payoffs to be received.



On the other hand, a normal form game is one where games are approximated as a single shot. It represents only actions and utilities. This is the table format that we're used to using at this point.

World outcomes

	ACTIONS		
	CRAM	DO HW	PLAY GAME
EASY	98	100	85
HARD	97	90	65

UTILITIES

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

	CRAM	DO HW	PLAY GAME
	98	100	85
	97	90	65

$P(\text{EASY}) = .2$
 $P(\text{HARD}) = .8$

- What is the utility of the pure strategy: cram?

$$98 * 0.2 + 97 * 0.8 = 97.2$$

- What is the utility of the pure strategy: do HW?

$$100 * 0.2 + 90 * 0.8 = 92$$

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

$$\frac{1}{2} * 97.2 + \frac{1}{2} * 92 = 94.6$$

- (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.

Yes. If $x_1 = x_2 = 2$, then neither player has any incentive to unilaterally change their strategy. There is no option to choose a lower number, and any player who chooses to change their number to anything greater than 2 will then receive a utility of 0, which is strictly worse than their current utility of 2.

- (d) Describe the following voting strategies:

- Plurality

Each voter gets one vote for their top-ranked preference. Alternative with the most votes wins.

- Borda Count

Each voter awards $m - k$ points to their rank k . Alternative with the most votes wins.

- Single Transferable Vote

Each voter gets 1 vote per round. In each round, the alternative with the least number of plurality votes is eliminated. Alternative left standing is winner

- Pairwise Elections

Alternative x beats y in pairwise election if majority of voters prefer x to y .

- Plurality with Runoff

First Round: Top 2 plurality winners advance to second round. Second Round: Pairwise election between two winners.

- Condorcet Winner

Alternative x beats y in pairwise election if majority of voters prefer x to y . Condorcet winner x beats every other alternative y in pairwise election. Condorcet paradox is a cycle in majority preferences.

- (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

Neither! In the first case, the plurality vote is A; borda count is B; condorcet is A. In the second case, the plurality vote is A; the borda count is B; and the condorcet is B.

(f) Which voting rule is the best?

Hard to say! [Mao, Procaccia, Chen 2013] Compared ranking strategies Plurality, Borda, Condorcet. Found that Borda finds the winners most consistently even with noisy human responses, and Plurality performs the worst.

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>Elmo's Christmas Countdown</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

Elmo's Christmas Countdown. In Plurality, we only count the top candidate for each voter, so *Elmo's Christmas Countdown* has the highest plurality votes at 45 votes.

(b) Borda Count

Frankenstein. In Borda Count, we give each voter's lowest choice candidate zero votes, second lowest candidate choice one vote, and so on. If we count these votes and multiply by the number of votes for each order, we see that *Frankenstein* has the highest number of Borda Count votes at $35 \cdot 5 + 45 \cdot 4 + 10 \cdot 1 = 365$.

(c) Single Transferable Vote (STV)

Elmo's Christmas Countdown. In Single Transferable Vote, voters rank their candidates and each round the least popular candidate gets "eliminated". If a voter's top choice candidate is eliminated, their vote now goes towards their second ranked candidate until there is one winner remaining. With this system, everything besides *Arrival*, *Frankenstein*, and *Elmo's Christmas Countdown* get eliminated in the first three rounds because they have 0 top choice votes. In the fourth round, *Arrival* gets eliminated with only 10 votes and category 1 votes get sent to *Elmo's Christmas Countdown*. Then, in round five, *Frankenstein* is eliminated leaving *Elmo's Christmas Countdown* as the winner.

(d) Condorcet Winner

None. A Condorcet Winner is the alternative that beats every other alternative in a pairwise election. In these votes, there is no such candidate that beats every other candidate pairwise, so the Condorcet Winner does not exist.

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)?

Deep Blue Sea. If we count votes given this rule, Deep Blue Sea has the most votes at $35 + 45 = 80$. This is tied with Frankenstein, but we break ties alphabetically favoring Deep Blue Sea.

- (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.

Yes. A person who voted in the first category can change their vote to make Frankenstein win in this new voting scheme. This is because Frankenstein and Deep Blue Sea are currently tied at 80 votes, so a person from category 1 can change their vote to include Frankenstein and not Deep Blue Sea in the first three. This will put Frankenstein at 81 votes and Deep Blue Sea still at 80.

- (g) Which properties among the following does this new voting rule satisfy?

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

Onto.

If voting rule is majority consistent, then if an alternative gets more than 50 percent of the vote then that alternative wins. (Notably, the Borda Count voting rule is not majority consistent). If category 1 only had 9 voters instead, Elmo's Christmas Countdown would have more than 50 percent of the vote, but Deep Blue Sea would still win, so this does not satisfy majority consistency.

A voting rule is Condorcet Consistent if it picks the Condorcet winner if one exists as the final winner. We can see this is not true from a simple counterexample of $((a,b,c,d), (a,d,c,b), (b,d,c,a))$. We can verify that candidate A beats every other candidate pairwise. However, candidate C wins in our new voting scheme, as it gets 3 votes while the others all get 2.

This is not strategy proof because if one player in Category one changed their voting rules like in the previous problem, the winner could change.


A voting rule is dictatorial if there is a voter who always gets her most preferred alternative, which does not occur with this voting rule since there is equal weight to the votes from all voters.

A voting rule is constant if the same alternative is always chosen (regardless of the stated preferences). This does not occur with this voting rule because the winner depends on the voter's preferences.

A voting rule is onto if any alternative can win for some set of stated preferences, which this voting rule satisfies.

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		

Method	$\langle +r, -w, +m \rangle$	$\langle +r, +w, -m \rangle$
Prior Sampling	$0.4 * 0.8 * 0.45 = 0.144$	$0.4 * 0.2 * 0.9 = 0.072$
Likelihood Weighting	$P(+r)P(-w) = 0.4 * 0.8 = 0.32$	0

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
0	0.5	0	1	0.1
1	0.5	1	0	0.5
1	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	

X_1	$P(X_1)$
0	$P(X_0 = 0)P(X_1 = 0 X_0 = 0) + P(X_0 = 1)P(X_1 = 0 X_0 = 1) = .5 * .9 + .5 * .5 = .7$
1	$P(X_0 = 0)P(X_1 = 1 X_0 = 0) + P(X_0 = 1)P(X_1 = 1 X_0 = 1) = .5 * .1 + .5 * .5 = .3$

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 .

$$\begin{aligned}
 P(X_1 = 0|E_1 = 0) &= \frac{P(E_1 = 0|X_1 = 0)P(X_1 = 0)}{\sum_x P(E_1 = 0|X_1 = x)P(X_1 = x)} \\
 &= \frac{\beta * .7}{\beta * .7 + (1 - \beta) * .3}
 \end{aligned}$$

- (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 $P(X_1 = 0|E_1 = 0) > P(X_1 = 0)$?

$\beta \in (0.5, 1]$. Intuitively, observing $E_1 = 0$ will only increase the belief that $X_1 = 0$ if $E_1 = 0$ is more likely given $X_1 = 0$ than not.

- (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	

X_t	$\hat{P}(X_t)$
0	$2/5$
1	$3/5$

- (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 $X_t = 0, X_t = 1$.

We can first find the joint probability $\hat{P}(X_t, e_t)$ (e_t being 1):

X_t	$\hat{P}(X_t, e_t)$
0	$\hat{P}(X_t = 0) * P(E_t = 1 X_t = 0) = 2/5 * (1 - 0.6) = 0.16$
1	$\hat{P}(X_t = 1) * P(E_t = 1 X_t = 1) = 3/5 * 0.6 = .36$

We then normalize to get the posterior $\hat{P}(X_t|e_t)$:

X_t	$\hat{P}(X_t e_t)$
0	$.16 / (.16 + .36) = .307$
1	$.36 / (.16 + .36) = .693$

Finally, using our fixed order and the given random number list, we resample:

$$0.182 < 0.307 \rightarrow X_t = 0$$

$$0.703 \geq 0.307 \rightarrow X_t = 1$$

$$0.471 \geq 0.307 \rightarrow X_t = 1$$


$$0.859 \geq 0.307 \rightarrow X_t = 1$$

$$0.382 \geq 0.307 \rightarrow X_t = 1,$$

getting the samples [0, 1, 1, 1, 1].

There are 4 particles with state value 1, and 1 particle with state value 0.

- (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) $P(X_1 | E_1)$

$$\begin{aligned} P(X_1 | E_1) &= \frac{P(X_1, E_1)}{P(E_1)} = \frac{\sum_{y_1, x_0} P(X_1, E_1, x_0, y_1)}{\sum_{y_1, x_0, x_1} P(x_1, E_1, x_0, y_1)} \\ &= \frac{\sum_{y_1, x_0} P(y_1)P(X_1 | x_0, y_1)P(E_1 | X_1)P(x_0)}{\sum_{y_1, x_1, x_0} P(y_1)P(x_1 | x_0, y_1)P(E_1 | x_1)P(x_0)} \end{aligned}$$

(ii) $P(Y_1 | X_1, E_1)$ $P(Y_1 | X_1, E_1) = P(Y_1 | X_1)$ since $Y_1 \perp\!\!\!\perp E_1 | X_1$ (Markov property).

$$P(Y_1 | X_1) = \frac{P(X_1, Y_1)}{P(X_1)} = \frac{\sum_{x_0} P(X_1, Y_1, x_0)}{\sum_{x_0, y_1} P(X_1, y_1, x_0)} = \frac{\sum_{x_0} P(Y_1)P(X_1 | x_0, Y_1)P(x_0)}{\sum_{x_0, y_1} P(y_1)P(X_1 | x_0, y_1)P(x_0)}$$

(iii) $P(Y_1 | E_1)$

$$P(Y_1 | E_1) = \frac{P(E_1, Y_1)}{P(E_1)} = \frac{\sum_{x_0, x_1} P(Y_1)P(E_1 | x_1)P(x_1 | x_0, Y_1)P(x_0)}{\sum_{x_0, x_1, y_1} P(y_1)P(E_1 | x_1)P(x_1 | x_0, y_1)P(x_0)}$$

(iv) $P(Y_2 | E_1, E_2)$

$$\begin{aligned} P(Y_2 | E_1, E_2) &= \frac{P(Y_2, E_1, E_2)}{P(E_1, E_2)} \\ &= \frac{\sum_{x_0, x_1, x_2, y_1} P(Y_2)P(y_1)P(x_2 | x_1, Y_2)P(x_1 | x_0, y_1)P(E_2 | x_2)P(E_1 | x_1)P(x_0)}{\sum_{x_0, x_1, x_2, y_1, y_2} P(y_2)P(y_1)P(x_2 | x_1, y_2)P(x_1 | x_0, y_1)P(E_2 | x_2)P(E_1 | x_1)P(x_0)} \end{aligned}$$