

Game Theory

Lecturer: Ji Liu

Thanks for Jerry Zhu's slides

[based on slides from Andrew Moore <http://www.cs.cmu.edu/~awm/tutorials>]

Overview

- Matrix normal form
- Chance games
- Games with hidden information
- Non-zero sum games

Pure strategy

- A **pure strategy** for a player is the mapping between all possible states the player can see, to the move the player would make.

- Player A has 4 pure strategies:

A's strategy I: $(1 \rightarrow L, 4 \rightarrow L)$

A's strategy II: $(1 \rightarrow L, 4 \rightarrow R)$

A's strategy III: $(1 \rightarrow R, 4 \rightarrow L)$

A's strategy IV: $(1 \rightarrow R, 4 \rightarrow R)$


- Player B has 3 pure strategies:

B's strategy I: $(2 \rightarrow L, 3 \rightarrow R)$

B's strategy II: $(2 \rightarrow M, 3 \rightarrow R)$

B's strategy III: $(2 \rightarrow R, 3 \rightarrow R)$

- How many pure strategies if each player can see N states, and has b moves at each state?



Matrix Normal Form of games

A's strategy I: (1→L, 4→L)

A's strategy II: (1→L, 4→R)


A's strategy III: (1→R, 4→L)

A's strategy IV: (1→R, 4→R)

B's strategy I: (2→L, 3→R)

B's strategy II: (2→M, 3→R)

B's strategy III: (2→R, 3→R)




- The matrix normal form is the game value matrix indexed by each player's strategies.

	B-I	B-II	B-III
A-I	7	3	-1
A-II	7	3	4
A-III	5	5	5
A-IV	5	5	5


The matrix encodes every outcome of the game! The rules etc. are no longer needed.

Matrix normal form example



- How many pure strategies does A have?
- How many does B have?
- What is the matrix form of this game?

Matrix normal form example



	B-I	B-II	B-III	B-IV
A-I	-1	-1	2	2
A-II	4	4	2	2
A-III	5	2	5	2
A-IV	5	2	5	2

- How many pure strategies does A have? 4

A-I (1→L, 4→L) A-II (1→L,4→R) A-III (1→R,4→L) A-IV (1→R, 4→R)


- How many does B have? 4

B-I (2→L, 3→L) B-II (2→L,3→R) B-III (2→R,3→L) B-IV (2→R, 3→R)

- What is the matrix form of this game?

Minimax in Matrix Normal Form


- Player A: for each strategy, consider all B's counter strategies (a row in the matrix), find the **minimum value** in that row. Pick the row with the maximum minimum value.
- Here $\text{maximin}=5$



	B-I	B-II	B-III
A-I	7	3	-1
A-II	7	3	4
A-III	5	5	5
A-IV	5	5	5

Minimax in Matrix Normal Form

- Player B: find the **maximum value** in each column. Pick the column with the minimum maximum value.
- Here minimax = 5



Fundamental game theory result (proved by von Neumann):

In a 2-player, zero-sum game of perfect information, Minimax==Maximin. And there always exists an optimal pure strategy for each player.

	B-I	B-II	B-III
A-I	7	3	-1
A-II	7	3	4
A-III	5	5	5
A-IV	5	5	5

Minimax in Matrix Normal Form

Interestingly, A can tell B in advance what strategy A will use (the maximin), and this information will not help B!

Similarly B can tell A what strategy B will use.


In fact A knows what B's strategy will be.

And B knows A's too.

And A knows that B knows

...


The game is at an equilibrium



	B-I	B-II	B-III
A-I	7	3	-1
A-II	7	3	4
A-III	5	5	5
A-IV	5	5	5

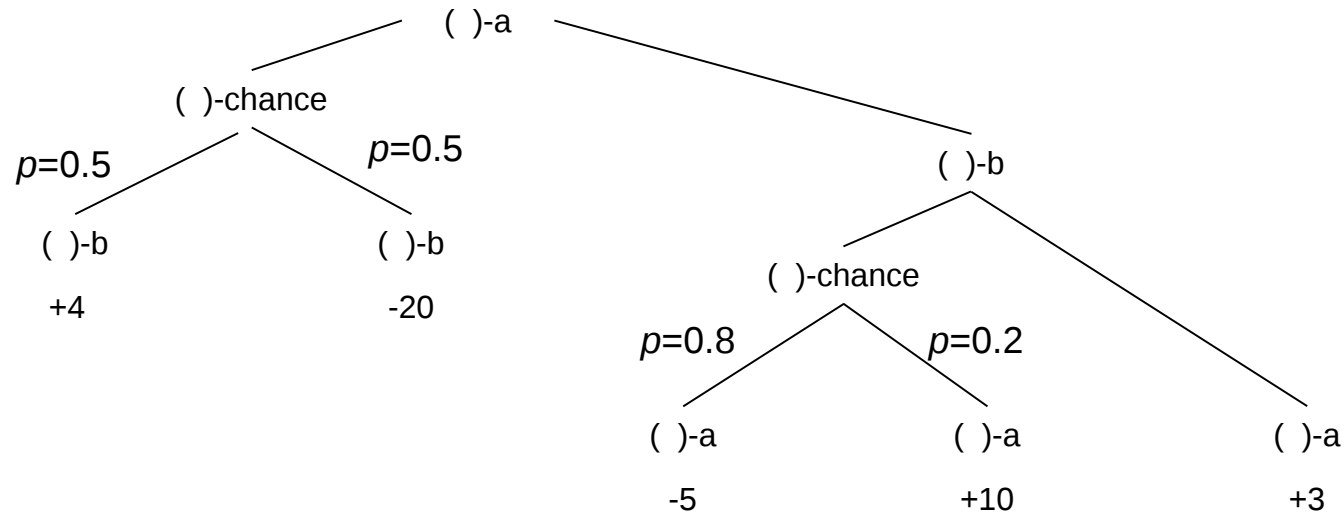
Matrix Normal Form for **NONdeterministic** games

- Recall the chance nodes (coin flip, die roll etc.): neither player moves, but a random move is made according to the known probability



- The game theoretic value is the expected value if both players are optimal
- What's the matrix form of this game?

Matrix Normal Form for NONdeterministic games



- A-I: L, A-II: R, B-I: L, B-II: R
- The i,j^{th} entry is the **expected** value with strategies A-i,B-j
- von Neumann's result still holds
- Minimax == Maximin

	B-I	B-II
A-I	-8	-8
A-II	-2	3

Non-zero sum games

Non-zero sum games

- One player's gain is not the other's loss
- Matrix normal form: simply lists all players' gain

	B-I	B-II
A-I	-5, -5	-10, 0
A-II	0, -10	-1, -1

Convention: A's gain first, B's next

Note B now wants to maximize the blue numbers.

- Previous zero-sum games trivially represented as

	O-I	O-II
E-I	2, -2	-3, 3
E-II	-3, 3	4, -4

Prisoner's dilemma

	B-testify	B-refuse
A-testify	-5, -5	0, -10
A-refuse	-10, 0	-1, -1

Strict domination

- A's strategy i dominates A's strategy j , if for every B's strategy, A is better off doing i than j .

	B-testify	B-refuse
A-testify	-5, -5	0, -10
A-refuse	-10, 0	-1, -1

If B-testify: A-testify (-5) is better than A-refuse (-10)

If B-refuse: A-testify (0) is better than A-refuse (-1)

A: Testify is always better than refuse.

A-testify **strictly dominates** (all outcomes strictly better than) A-refuse.

Strict domination

- Fundamental assumption of game theory: **get rid of strictly dominated strategies – they won't happen.**
- In some cases like prisoner's dilemma, we can use strict domination to predict the outcome, if both players are **rational**.

	B-testify	B-refuse
A-testify	-5, -5	0, -10
A-refuse	-10, 0	-1, -1

Strict domination

- Fundamental assumption of game theory: **get rid of strictly dominated strategies – they won't happen.**
- In some cases like prisoner's dilemma, we can use strict domination to predict the outcome, if both players are **rational**.

	B-testify	B-refuse
A-testify	-5, -5	0, -10
A-refuse	10, 0	1, 1



	B-testify	B-refuse
A-testify	-5, -5	0, -10

Strict domination

- Fundamental assumption of game theory: **get rid of strictly dominated strategies – they won't happen.**
- In some cases like prisoner's dilemma, we can use strict domination to predict the outcome, if both players are **rational**.

	B-testify	B-refuse
A-testify	-5, -5	0, -10
A-refuse	10, 0	1, 1



	B-testify	B-refuse
A-testify	-5, -5	0, -10



	B-testify
A-testify	-5, -5

Another strict domination example

- Iterated elimination of strictly dominated strategies

		Player B			
		I	II	III	IV
Player A	I	3 , 1	4 , 1	5 , 9	2 , 6
	II	5 , 3	5 , 8	9 , 7	9 , 3
	III	2 , 3	8 , 4	6 , 2	6 , 3
	IV	3 , 8	3 , 1	2 , 3	4 , 5

Strict domination?

- Strict domination doesn't always happen...

	I	II	III
I	0 , 4	4 , 0	5 , 3
II	4 , 0	0 , 4	5 , 3
III	3 , 5	3 , 5	6 , 6

- What do you think the players will do?

Nash equilibria

- (player 1's strategy s_1^* , player 2's strategy s_2^* , ... player n's strategy s_n^*) is a **Nash equilibrium**, iff

$$s_i^* = \operatorname{argmin}_s V(s_1^*, \dots, s_{i-1}^*, s, s_{i+1}^*, \dots, s_n^*)$$

- This says: if everybody else plays at the Nash equilibrium, player i will hurt itself unless it also plays at the Nash equilibrium.

N.E. is a local maximum in unilateral moves.

	I	II	III
I	0 , 4	4 , 0	5 , 3
II	4 , 0	0 , 4	5 , 3
III	3 , 5	3 , 5	6 , 6

Nash equilibria examples

	B-testify	B-refuse
A-testify	-5, -5	0, -10
A-refuse	-10, 0	-1, -1

1. Is there always a Nash equilibrium?
2. Can there be more than one Nash equilibrium?

Player B

I II III IV

Player A

I	3, 1	4, 1	5, 9	2, 6
II	5, 3	5, 8	9, 7	9, 3
III	2, 3	8, 4	6, 2	6, 3
IV	3, 8	3, 1	2, 3	4, 5

Example: no N.E. with pure strategies

- two-finger Morra

	O-I	O-II
E-I	2, -2	-3, 3
E-II	-3, 3	4, -4

- No pure strategy Nash equilibrium, but...

Two-player zero-sum deterministic game with **hidden** information

- Hidden information: something you don't know but your opponent knows, e.g. hidden cards, or simultaneous moves
- Example: two-finger Morra
 - Each player (O and E) displays 1 or 2 fingers
 - If sum f is odd, O collects $\$f$ from E
 - If sum f is even, E collects $\$f$ from O
 - Strategies?
 - Matrix form?

Two-player zero-sum deterministic game with hidden information

- Hidden information: something you don't know but your opponent knows, e.g. hidden cards, or simultaneous moves
- Example: two-finger Morra
 - Each player (O and E) displays 1 or 2 fingers
 - If sum f is odd, O collects $\$f$ from E
 - If sum f is even, E collects $\$f$ from O
 - Strategies?
 - Matrix form?
 - **Maximin = -3, minimax = 2**
 - The two are not the same!
 - What should O and E do?

	O-I	O-II
E-I	2,-2	-3,3
E-II	-3,3	4,-4

Game theoretic value when there is hidden information

- It turns out O can win a little over 8 cents on average in each game, if O does the right thing.
- Again O can tell E what O will do, and E can do nothing about it!
- The trick is to use a **mixed strategy** instead of a pure strategy.
 - A mixed strategy is defined by a probability distribution (p_1, p_2, \dots, p_n) . $n = \#$ of pure strategies the player has
 - At the start of each game, the player picks number i according to p_i , and uses the i^{th} pure strategy for this round of the game
- von Neumann: every two-player zero-sum game (even with hidden information) has an optimal (mixed) strategy.

Boring math: Two-finger Morra

- E's mixed strategy: $(p:I, (1-p):II)$
- O's mixed strategy: $(q:I, (1-q):II)$
- What is p, q ?
- step 1: let's fix p for E, and O knows that.
 - What if O always play O-I ($q=1$)? $v_1=p*2+(1-p)*(-3)$
 - What if O always play O-II ($q=0$)? $v_0=p*(-3)+(1-p)*4$
 - And if O uses some other q ? $q*v_1+(1-q)*v_0$
 - O is going to pick q to minimize $q*v_1+(1-q)*v_0$
 - Since this is a linear combination, such q must be 0 or 1, not something in between!
 - The value for E is $\min(p*2+(1-p)*(-3), p*(-3)+(1-p)*4)$
- step 2: E choose the p that maximizes the value above.

	O-I	O-II
E-I	2	-3
E-II	-3	4

More boring math

- step 1: let's fix p for E.
 - The value for E is $\min(p^*2+(1-p)^*(-3), p^*(-3)+(1-p)^*4)$, in case O is really nasty
- step 2: E choose the p^* that maximizes the value above.
$$p^* = \operatorname{argmax}_p \min(p^*2+(1-p)^*(-3), p^*(-3)+(1-p)^*4)$$
- Solve it with (proof by “it's obvious”)
$$p^*2+(1-p)^*(-3) = p^*(-3)+(1-p)^*4$$
- E's optimal $p^* = 7/12$, value = $-1/12$ (expect to lose \$!)
That's the best E can do!)
- Similar analysis on O shows $q^* = 7/12$, value = $1/12$

This is a zero-sum,
but unfair game.

Recipe for computing A's optimal mixed strategy for a n*m game

- n*m game = A has n pure strategies and B has m. $v_{ij} = (i,j)^{\text{th}}$ entry in the matrix form.
- Say A uses mixed strategy (p_1, p_2, \dots, p_n) .

A's expected gain if B uses pure strategy 1: $g_1 = p_1 v_{11} + p_2 v_{21} + \dots + p_n v_{n1}$

A's expected gain if B uses pure strategy 2: $g_2 = p_1 v_{12} + p_2 v_{22} + \dots + p_n v_{n2}$

...

A's expected gain if B uses pure strategy m: $g_m = p_1 v_{1m} + p_2 v_{2m} + \dots + p_n v_{nm}$

- Choose (p_1, p_2, \dots, p_n) to maximize

$$\min(g_1, g_2, \dots, g_m)$$

$$\text{Subject to: } p_1 + p_2 + \dots + p_n = 1$$

$$0 \leq p_i \leq 1 \text{ for all } i$$

Fundamental theorems

- In a n-player pure strategy game, if iterated elimination of strictly dominated strategies leaves all but one cell $(s_1^*, s_2^*, \dots, s_n^*)$, then it is the unique NE of the game
- Any NE will survive iterated elimination of strictly dominated strategies
- [Nash 1950]: If n is finite, and each player has finite strategies, then there exists at least one NE (possibly involving mixed strategies)

