Game Theory

Louis Oliphant
oliphant@cs.wisc.edu
CS540-2

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→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)
- Player B has 3 pure strategies:
  - B's strategy I: (2→L, 3→R)
  - B's strategy II: (2→M, 3→R)
  - B's strategy III: (2→R, 3→R)
- How many pure strategies if each player can see *N* states, and has *b* moves at each state?

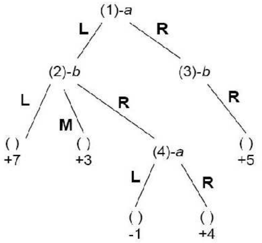

Matrix Normal Form of games

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)

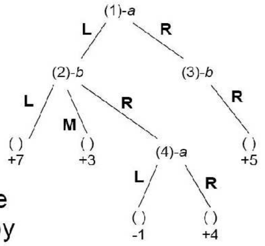

- The matrix normal form is the game value matrix indexed by each player's strategies.

	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

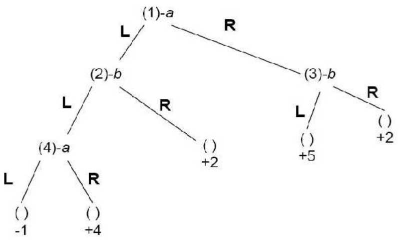

- How many pure strategies does A have?
- How many does B have?
- What is the matrix form of this game?

Matrix normal form example

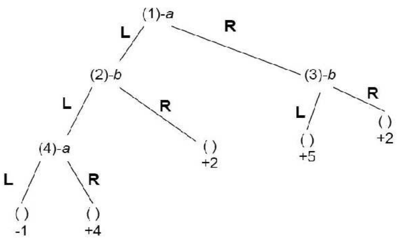

	B-I	B-II	B-III	B-IV
A-I	-1	-1	4	4
A-II	4	4	2	2
A-III	5	2	5	2
A-IV	5	2	5	2

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

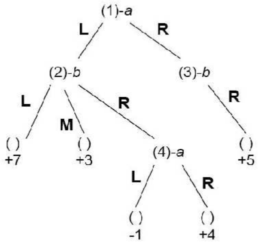

Fundamental game theory result (proved by von Neumann):

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

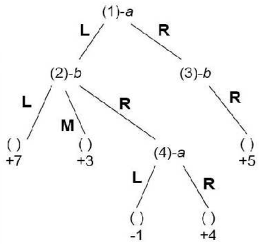

	B-I	B-II	B-III
A-I	7	3	-1
A-II	7	3	4
A-III	5	5	5
A-IV	5	5	5

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	-3
E-II	-3	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, … 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.

How To Calculate the Optimal Mixed Strategy (for 2x2 Matrix) Provided Opponent plays optimally

- Using the following probabilities:
    - 1 Finger for Even – p
    - 2 Fingers for Even– (1-p)
    - 1 Finger for Odd – q
    - 2 Fingers for Odd – (1-q)
- Now assume that Odd will set q=1 and Even knows this. Calculate the Expected Outcome
- Now assume that Odd will set q=0 and Even knows this. Calculate the Expected Outcome
- Find where the two expected outcomes are equal
    - This is the point where it doesn't matter what q is set to, the expected outcome will be the same either way

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

How To Calculate the Optimal Mixed Strategy (for 2x2 Matrix) Provided Opponent plays optimally

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

- Now assume that Odd will set q=1 and Even knows this. Calculate the Expected Outcome

  Expected Outcome=2p-3(1-p)
  Expected Outcome=5p-3
- Now assume that Odd will set q=0 and Even knows this. Calculate the Expected Outcome

  Expected Outcome=-3p+4(1-p)
  Expected Outcome=-7p+4
- Find where the two expected outcomes are equal

  5p-3=-7p+4
  12p=7
  p=7/12

How To Calculate the Optimal Mixed Strategy (for 2x2 Matrix) Provided Opponent plays optimally

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

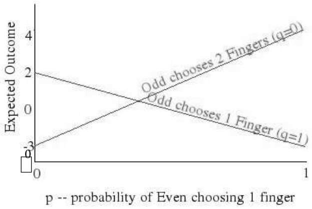

What if q is something between q=0 and q=0? say q=1/2.
Expected Outcome=p(1/2*2+1/2*-3)+(1-p)(1/2*-3+1/2*4)
Expected Outcome=1/2-p

How To Calculate the Optimal Mixed Strategy (for 2x2 Matrix) Provided Opponent plays optimally

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

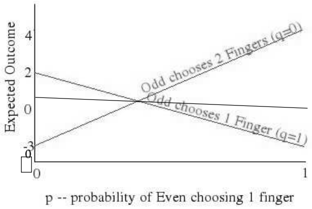

What if q is something between q=0 and q=0? say q=1/2.
Expected Outcome=p(1/2*2+1/2*-3)+(1-p)(1/2*-3+1/2*4)
Expected Outcome=1/2-p

How To Calculate the Optimal Mixed Strategy (for 2x2 Matrix) Provided Opponent plays optimally

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

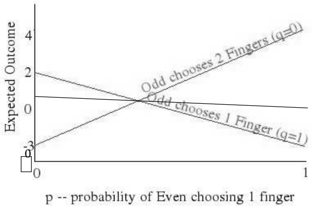

What if q is something else? say q=3/4.
Expected Outcome=p(3/4*2+1/4*-3)+(1-p)(3/4*-3+1/4*4)
Expected Outcome=2p-5/4

How To Calculate the Optimal Mixed Strategy (for 2x2 Matrix) Provided Opponent plays optimally

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

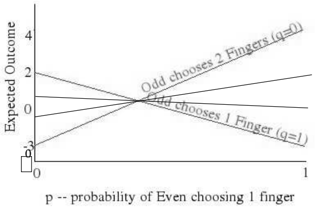

What if q is something else? say q=3/4.
Expected Outcome=p(3/4*2+1/4*-3)+(1-p)(3/4*-3+1/4*4)
Expected Outcome=2p-5/4

How To Calculate the Optimal Mixed Strategy (for 2x2 Matrix) Provided Opponent plays optimally

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

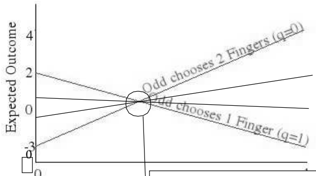

At This Point it Doesn't matter what q is.
This is the best I can do, if q is playing optimally.

What if q is something else? say q=3/4.
Expected Outcome=p(3/4*2+1/4*-3)+(1-p)(3/4*-3+1/4*4)
Expected Outcome=2p-5/4

How To Calculate the Optimal Mixed Strategy (for 2x2 Matrix) Provided Opponent plays optimally

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

- So what is the Expected Outcome when p=7/12
  - Just plug 7/12 back into one of the Expected Outcome formulas:

  2*7/12-3*5/12 = **-1/12**
- So if Even follows this Mixed policy he can **expect to loose 1/12 of a dollar every game**
- That is the bad news. The good news is Odd can't do anything about it. No matter what the strategy (from q=0 to q=1) the expected outcome will be the same – PROVIDED Even doesn't know Odd's strategy

---

What if you know the others strategy (and its not optimal)

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

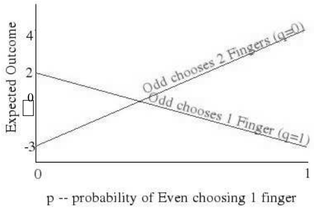

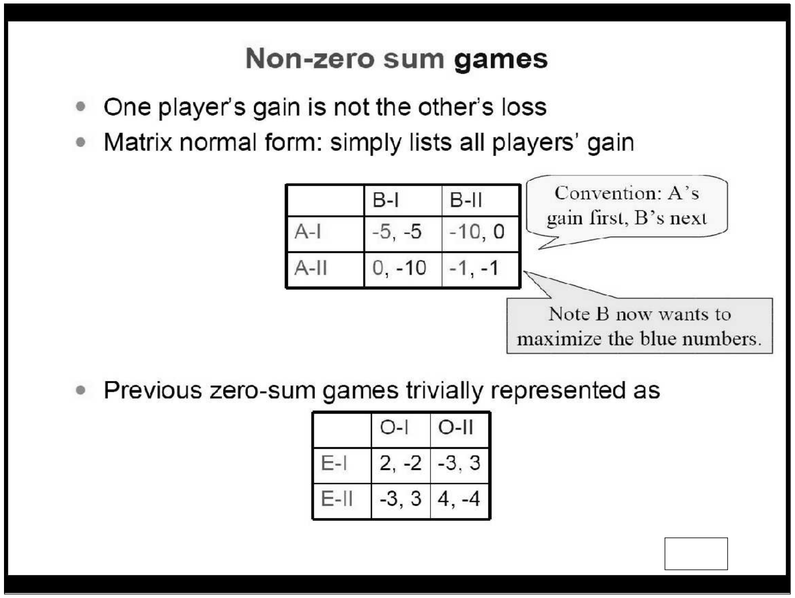

Non-zero sum games

- 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 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

Player A

	I	II	III	IV
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

$$s_i^* = \underset{s_i}{\text{argmax}}\ v(s_1^* \ldots s_{(i-1)}^*, s_i, s_{(i+1)}^* \ldots s_n^*), \text{ for all } i$$

- 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	$\boxed{6 , 6}$

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^*, … 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)