

CS 331: Artificial Intelligence Game Theory I

1

Prisoner's Dilemma

You and your partner have both been caught red handed near the scene of a burglary. Both of you have been brought to the police station, where you are interrogated **separately** by the police.



2

Prisoner's Dilemma

The police present your options:

1. You can **testify** against your partner
2. You can **refuse** to testify against your partner (and keep your mouth shut)

3

Prisoner's Dilemma

Here are the consequences of your actions:

- If you **testify** against your partner and your partner **refuses**, you are released and your partner will serve 10 years in jail
- If you **refuse** and your partner **testifies** against you, you will serve 10 years in jail and your partner is released
- If both of you **testify** against each other, both of you will serve 5 years in jail
- If both of you **refuse**, both of you will only serve 1 year in jail

4

Prisoner's Dilemma

- Your partner is offered the same deal
- Remember that you can't communicate with your partner and you don't know what he/she will do
- Will you **testify** or **refuse**?

5

Game Theory

- Welcome to the world of Game Theory!
- Game Theory defined as "the study of rational decision-making in situations of conflict and/or cooperation"
- Adversarial search is part of Game Theory
- We will now look at a much broader group of games

6

Types of games we will deal with today

- Two players
- Discrete, finite action space
- Simultaneous moves (or without knowledge of the other player's move)
- Imperfect information
- Zero sum games and non-zero sum games

7

Uses of Game Theory

- Agent design: determine the best strategy against a rational player and the expected return for each player
- Mechanism design: Define the rules of the game to influence the behavior of the agents

Real world applications: negotiations, bandwidth sharing, auctions, bankruptcy proceedings, pricing decisions

8

Back to Prisoner's Dilemma

Normal-form (or matrix-form) representation

Players: Alice, Bob

Actions: testify, refuse

	Bob: testify	Bob: refuse
Alice: testify	A = -5, B = -5	A = 0, B = -10
Alice: refuse	A = -10, B = 0	A = -1, B = -1

Payoffs for each player (non-zero sum game in this example)

9

Formal definition of Normal Form

The normal-form representation of an n-player game specifies:

- The players' strategy spaces S_1, \dots, S_n
- Their payoff functions u_1, \dots, u_n where $u_i: S_1 \times S_2 \times \dots \times S_n \rightarrow R$ i.e. a function that maps from the combination of strategies of all the players and returns the payoff for player i

10

Strategies

- Each player must adopt and execute a strategy
- Strategy = policy i.e. mapping from state to action
- Prisoner's Dilemma is a one move game:
 - Strategy is a single action
 - There is only a single state
- A pure strategy is a deterministic policy

11

Other Normal Form Games

The game of chicken: two cars drive at each other on a narrow road. The first one to swerve loses.

	B: Stay	B: Swerve
A: Stay	A = -100, B = -100	A = 1, B = -1
A: Swerve	A = -1, B = 1	A = 0, B = 0

12

Other Normal Form Games

Penalty kick in Soccer: Shooter vs. Goalie. The shooter shoots the ball either to the left or to the right. The goalie dives either left or right. If it's the same side as the ball was shot, the goalie makes the save. Otherwise, the shooter scores.

	Goalie: Left	Goalie: Right
Shooter: Left	S = -1, G = 1	S = 1, G = -1
Shooter: Right	S = 1, G = -1	S = -1, G = 1

13

Prisoner's Dilemma Strategy

	Bob: testify	Bob: refuse
Alice: testify	A = -5, B = -5	A = 0, B = -10
Alice: refuse	A = -10, B = 0	A = -1, B = -1

- What is the right pure strategy for Alice or Bob?
- (Assume both want to maximize their own expected utility)

14

Prisoner's Dilemma Strategy

	Bob: testify	Bob: refuse
Alice: testify	A = -5, B = -5	A = 0, B = -10
Alice: refuse	A = -10, B = 0	A = -1, B = -1

Alice thinks:

- If Bob testifies, I get 5 years if I testify and 10 years if I don't
- If Bob doesn't testify, I get 0 years if I testify and 1 year if I don't
- "Alright I'll testify"

15

Prisoner's Dilemma Strategy

	Bob: testify	Bob: refuse
Alice: testify	A = -5, B = -5	A = 0, B = -10
Alice: refuse	A = -10, B = 0	A = -1, B = -1

Testify is a **dominant strategy** for the game (notice how the payoffs for Alice are always bigger if she testifies than if she refuses)

16

Dominant Strategies

Suppose a player has two strategies S and S'. We say S **dominates** S' if choosing S always yields at least as good an outcome as choosing S'.

- S **strictly dominates** S' if choosing S always gives a better outcome than choosing S' (no matter what the other player does)
- S **weakly dominates** S' if there is one set of opponent's actions for which S is superior, and all other sets of opponent's actions give S and S' the same payoff.

17

Example of Dominant Strategies

	Bob: testify	Bob: refuse
Alice: testify	A = -5, B = -5	A = 0, B = -10
Alice: refuse	A = -10, B = 0	A = -1, B = -1

"testify" strongly dominates "refuse"

	Bob: testify	Bob: refuse
Alice: testify	A = -5, B = -5	A = 0, B = -10
Alice: refuse	A = -10, B = 0	A = 0, B = -1

"testify" weakly dominates "refuse"

Note →

18

Dominated Strategies (The opposite)

S is **dominated** by S' if choosing S never gives a better outcome than choosing S', no matter what the other players do

- S is **strictly dominated** by S' if choosing S always gives a worse outcome than choosing S', no matter what the other player does
- S is **weakly dominated** by S' if there is at least one set of opponent's actions for which S gives a worse outcome than S', and all other sets of opponent's actions give S and S' the same payoff.

19

Dominance

- It is irrational not to play a strictly dominant strategy (if it exists)
- It is irrational to play a strictly dominated strategy
- Since Game Theory assumes players are rational, they will not play strictly dominated strategies

20

Iterated Elimination of Strictly Dominated Strategies

	Bob: testify	Bob: refuse
Alice: testify	A = -5, B = -5	A = 0, B = -10
Alice: refuse	A = -10, B = 0	A = -1, B = -1



Simplifies to:

	Bob: testify	Bob: refuse
Alice: testify	A = -5, B = -5	A = 0, B = -10

21

Iterated Elimination of Strictly Dominated Strategies

	Bob: testify	Bob: refuse
Alice: testify	A = -5, B = -5	A = 0, B = -10

But in this simplified game, "refuse" is also a strictly dominated strategy for Bob

22

Iterated Elimination of Strictly Dominated Strategies

	Bob: testify	Bob: refuse
Alice: testify	A = -5, B = -5	A = 0, B = -10



Simplifies to:

	Bob: testify
Alice: testify	A = -5, B = -5

This is the game-theoretic solution to Prisoner's Dilemma (note that it's worse off than if both players refuse)

23

Dominant Strategy Equilibrium

	Bob: testify	Bob: refuse
Alice: testify	A = -5, B = -5	A = 0, B = -10
Alice: refuse	A = -10, B = 0	A = -1, B = -1

- (testify, testify) is a dominant strategy equilibrium
- It's an equilibrium because no player can benefit by switching strategies given that the other player sticks with the same strategy
- An equilibrium is a local optimum in the space of policies

24

Pareto Optimal

- An outcome is **Pareto optimal** if there is no other outcome that all players would prefer
- An outcome is **Pareto dominated** by another outcome if all players would prefer the other outcome
- If Alice and Bob both testify, this outcome is Pareto dominated by the outcome if they both refuse.
- This is why it's called Prisoner's Dilemma

25

Iterated Prisoner's Dilemma

- Possible to arrive at the Pareto optimal solution
- Strategies for repeated game:
 - **Perpetual punishment**: *refuse* unless opponent has ever played *testify*
 - **Tit-for-tat**: start with *refuse*; then play the opponents previous move
- This situation arose in trench warfare in WWI (see [The Evolution of Cooperation](#) by Robert Axelrod for more)

26

What If No Strategies Are Strictly Dominated?

		B		
		S1	S2	S3
A	S1	A = 0, B = 4	A = 4, B = 0	A = 5, B = 3
	S2	A = 4, B = 0	A = 0, B = 4	A = 5, B = 3
	S3	A = 3, B = 5	A = 3, B = 5	A = 6, B = 6

How do we find these equilibrium points in the game?

27

Nash Equilibrium

- A dominant strategy equilibrium is a special case of a **Nash Equilibrium**
- **Nash Equilibrium**: A strategy profile in which no player wants to deviate from his or her strategy.
- **Strategy profile**: An assignment of a strategy to each player e.g. (testify, testify) in Prisoner's Dilemma
- Any Nash Equilibrium will survive iterated elimination of strictly dominated strategies

28

Nash Equilibrium in Prisoner's Dilemma

	Bob: testify	Bob: refuse
Alice: testify	A = -5, B = -5	A = 0, B = -10
Alice: refuse	A = -10, B = 0	A = -1, B = -1

If (testify, testify) is a Nash Equilibrium, then:

- Alice doesn't want to change her strategy of "testify" given that Bob chooses "testify"
- Bob doesn't want to change his strategy of "testify" given that Alice chooses "testify"

29

How to Spot a Nash Equilibrium

		B		
		S1	S2	S3
A	S1	A = 0, B = 4	A = 4, B = 0	A = 5, B = 3
	S2	A = 4, B = 0	A = 0, B = 4	A = 5, B = 3
	S3	A = 3, B = 5	A = 3, B = 5	A = 6, B = 6

30

How to Spot a Nash Equilibrium

		B		
		S1	S2	S3
A	S1	A = 0, B = 4	A = 4, B = 0	A = 5, B = 3
	S2	A = 4, B = 0	A = 0, B = 4	A = 5, B = 3
	S3	A = 3, B = 5	A = 3, B = 5	A = 6, B = 6

Go through each square and see:

- If player A gets a higher payoff if she changes her strategy
- If player B gets a higher payoff if he changes his strategy
- If the answer is no to both of the above, you have a Nash Equilibrium

31

How to Spot a Nash Equilibrium

		B		
		S1	S2	S3
A	S1	A = 0, B = 4	A = 4, B = 0	A = 5, B = 3
	S2	A = 4, B = 0	A = 0, B = 4	A = 5, B = 3
	S3	A = 3, B = 5	A = 3, B = 5	A = 6, B = 6

↑ A won't change her Strategy of S3
Payoff of 6 > 5 (S2) and 6 > 5 (S1)

← B won't change his Strategy of S3
Payoff of 6 > 5 (S2) and 6 > 5 (S1)

32

Formal Definition of A Nash Equilibrium (n-player)

Notation:

S_i = Set of strategies for player i

$s_i \in S_i$ means strategy s_i is a member of strategy set S_i

$u_i(s_1, s_2, \dots, s_n)$ = payoff for player i if all the players in the game play their respective strategies s_1, s_2, \dots, s_n .

$s_1^* \in S_1, s_2^* \in S_2, \dots, s_n^* \in S_n$ are a Nash equilibrium iff:

$$\forall i \ s_i^* = \arg \max_{s_i} u_i(s_1^*, \dots, s_{i-1}^*, s_i, s_{i+1}^*, \dots, s_n^*)$$

33

Formal Definition of a Nash Equilibrium

		B		
		S1	S2	S3
A	S1	A = 0, B = 4	A = 4, B = 0	A = 5, B = 3
	S2	A = 4, B = 0	A = 0, B = 4	A = 5, B = 3
	S3	A = 3, B = 5	A = 3, B = 5	A = 6, B = 6

Using the notation $u_i(A's \text{ strategy}, B's \text{ strategy})$:

$$u_A(S3, S3) = \max[u_A(S1, S3), u_A(S2, S3), u_A(S3, S3)]$$

$$u_B(S3, S3) = \max[u_B(S3, S1), u_B(S3, S2), u_B(S3, S3)]$$

34

Neat fact

- If your game has a single Nash Equilibrium, you can announce to your opponent that you will play your Nash Equilibrium strategy
- If your opponent is rational, he will have no choice but to play his part of the Nash Equilibrium strategy
- Why?

35

Can you have more than one Nash Equilibrium?

- ACME, a video game hardware manufacturer, has to decide whether its next game machine will use Blu-ray or DVDs
- Best, a video game software producer, needs to decide whether to produce its next game on Blu-ray or DVD
- Profits for both will be positive if they agree and negative if they disagree

36

Can you have more than one Nash Equilibrium?

	Best: bluray	Best: dvd
ACME: bluray	A = 9, B = 9	A = -3, B = -1
ACME: dvd	A = -4, B = -1	A = 5, B = 5

37

Can you have more than one Nash Equilibrium?

	Best: bluray	Best: dvd
ACME: bluray	A = 9, B = 9	A = -3, B = -1
ACME: dvd	A = -4, B = -1	A = 5, B = 5

There are two Nash Equilibria in this game. In general, you can have multiple Nash Equilibria. This creates a big problem. Can you see what that problem is?

38

Dealing with Multiple Nash Equilibria

- Could choose the Pareto-optimal Nash Equilibrium e.g. (bluray, bluray) but
 - What if there are multiple Pareto-optimal Nash Equilibria?
 - Or it's too computationally expensive to find all the Nash Equilibria?
 - Or there are an infinite number of Nash Equilibria?
- Could communicate before the game
 - But what if you can't compute all the Nash Equilibria beforehand?
- Take your best guess

This is a big unresolved issue

Can we have no Nash Equilibria?

Two Fingered Morra

Two players, O (for Odd) and E (for Even) simultaneously display one or two fingers. Let the total number of fingers be f .

- If f is odd, O collects f dollars from E.
- If f is even, E collects f dollars from O.

E is the max player

	O: one	O: two
E: one	E = 2, O = -2	E = -3, O = 3
E: two	E = -3, O = 3	E = 4, O = -4

40

Two Fingered Morra

	O: one	O: two
E: one	E = 2, O = -2	E = -3, O = 3
E: two	E = -3, O = 3	E = 4, O = -4

- No pure strategy Nash Equilibrium
- If total # of fingers is even, O will want to switch
- If total # of fingers is odd, E will want to switch
- Also, this is a zero-sum game (payoffs in a cell sum to zero)

41

The Big Theorem

- [Nash 1950] In the n -player normal-form game $G = \{S_1, \dots, S_n; u_1, \dots, u_n\}$, if n is finite and S_i is finite for every i then there exists at least one Nash Equilibrium, possibly involving mixed strategies

42

Mixed Strategies

- Recall that a pure strategy is a deterministic policy i.e. you pick a strategy and play it all the time
- A **mixed strategy** is a randomized policy i.e. you select your strategy based on a probability distribution
- E.g. Select strategy S1 with probability p and strategy S2 with probability $(1-p)$
- Is there a mixed strategy Nash Equilibrium in 2 Fingered Morra?

43

Formal Definition of a Mixed Strategy

In the normal-form game $G = \{S_1, \dots, S_n; u_1, \dots, u_n\}$, suppose $S_i = \{s_{i1}, \dots, s_{iK}\}$. Then a mixed strategy for a player i is a probability distribution $p_i = (p_{i1}, \dots, p_{iK})$,

where $0 \leq p_{ik} \leq 1$ for $k = 1, \dots, K$
and $p_{i1} + \dots + p_{iK} = 1$.

44

Mixed Strategy Nash Equilibrium

- The pair of mixed strategies (M_A, M_B) are a Nash Equilibrium iff
- Player A does not want to deviate from M_A (because M_A is Player A's best response to M_B and)
- Player B does not want to deviate from M_B (because M_B is Player B's best response to M_A)

45

Finding optimal mixed strategy for two-player zero-sum games

- Note: applies to zero-sum games (or, more generally, constant sum games)
- Von Neumann's maximin technique

46

Expected Payoff to E if O Uses a Mixed Strategy

	O: one	O: two
E: one	E = 2, O = -2	E = -3, O = 3
E: two	E = -3, O = 3	E = 4, O = -4

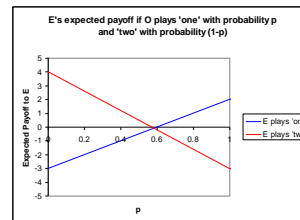
Suppose O chooses to display one finger with probability p and two fingers with probability $(1-p)$

If E chooses the pure strategy of one finger, E's expected profit is $2p - 3(1-p) = 2p - 3 + 3p = 5p - 3$

If E chooses the pure strategy of two fingers, E's expected profit is $-3p + 4(1-p) = -3p + 4 - 4p = -7p + 4$

47

Expected Payoff to E if O Uses a Mixed Strategy



$$5p - 3 = -7p + 4$$

$$\Rightarrow 12p = 7$$

$$\Rightarrow p = 7/12$$

When $p < 7/12$, E plays 'two'

When $p > 7/12$, E plays 'one'

O gets to pick p to minimize E's expected payoff. O picks the lowest point of the higher of the two lines. This happens at the intersection of the two lines.

E's expected payoff at $p = 7/12$ is $5(7/12) - 3 = -1/12$

O's mixed strategy is $(7/12$ for 'one', $5/12$ for 'two')

48

Expected Payoff to O if E Uses a Mixed Strategy

	O: one	O: two
E: one	E = 2, O = -2	E = -3, O = 3
E: two	E = -3, O = 3	E = 4, O = -4


Suppose E chooses to display one finger with probability q and two fingers with probability $(1-q)$

If O chooses the pure strategy of one finger, O's expected payoff is $-2q + 3(1-q) = -2q + 3 - 3q = -5q + 3$

If O chooses the pure strategy of two fingers, O's expected payoff is $3q - 4(1-q) = 3q - 4 + 4q = 7q - 4$

49

Expected Payoff to O if E Uses a Mixed Strategy



$$-5q + 3 = 7q - 4$$

$$\Rightarrow 7 = 12q$$

$$\Rightarrow q = 7/12$$

When $q < 7/12$, O plays 'one'

When $q > 7/12$, O plays 'two'

E gets to pick p to minimize O's expected payoff. E picks the lowest point of the higher of the two lines. This happens at the intersection of the two lines.

O's expected payoff at $q=7/12$ is $-5(7/12)+3 = -35/12 + 36/12 = 1/12$.

E's mixed strategy is $(7/12$ for 'one', $5/12$ for 'two')

50

Mixed Strategy

- E's expected payoff is $-1/12$, O's is $1/12$
- It is better to be O than to be E
- The final mixed strategy is for both players to play "one" with probability $7/12$ and "two" with probability $5/12$
 - It's a coincidence that both players have the same mixed strategy here; in general they could be different
- This is a maximin equilibrium (which is also a Nash equilibrium)

51

Theoretical Results

- Every two-player zero-sum game has a maximin equilibrium when you allow mixed strategies
- Every Nash equilibrium in a two-player zero-sum game is a maximin equilibrium for both players

52

Recipe for Computing Optimal Mixed Strategy 2x2 Constant-Sum Games

	B: S1	B: S2
A: S1	A = m_{11}	A = m_{21}
A: S2	A = m_{12}	A = m_{22}

- Let Player B use strategy S1 with probability p
- Compute Player A's expected payoff if A uses pure strategy S1: $m_{11}p + m_{21}(1-p)$
- Compute Player A's expected payoff if A uses pure strategy S2: $m_{12}p + m_{22}(1-p)$
- Find the p between 0 and 1 that minimizes $\max(m_{11}p + m_{21}(1-p), m_{12}p + m_{22}(1-p))$
- The optimum will be at $p=0$, $p=1$ or at the point where the two lines intersect
- Repeat by letting Player A use Strategy S1 with probability q but looking at B's payoffs now

Practice

	B: S1	B: S2
A: S1	A = -2, B = 2	A = 3, B = -3
A: S2	A = 1, B = -1	A = -2, B = 2

- Calculate B's Nash equilibrium strategy.
- Calculate A's expected payoff.
- Calculate A's Nash equilibrium strategy.
- Calculate B's expected payoff.

54

Recipe for Computing Optimal Mixed Strategy NxM Zero-Sum Games

- NxM game = Player A has N pure strategies, Player B has M pure strategies
- Let Player B use:
Strategy S1 with probability p_1
Strategy S2 with probability p_2
⋮
Strategy S_N with probability p_N
- Compute Player A's expected payoff if A uses:
Pure strategy S1: $e_1 = m_{11}p_1 + m_{21}p_2 + \dots + m_{N1}p_N$
Pure strategy S2: $e_2 = m_{12}p_1 + m_{22}p_2 + \dots + m_{N2}p_N$
⋮
Pure strategy SM: $e_M = m_{1M}p_1 + m_{2M}p_2 + \dots + m_{NM}p_N$
- Find p_1, p_2, \dots, p_N to minimize $\max(e_1, e_2, \dots, e_M)$ subject to $\sum p_i = 1$ and $0 \leq p_i \leq 1$ for all i
- Use a method called Linear Programming (polynomial time in number of actions)
- Repeat by letting Player A use a mixed strategy and looking at Player B's payoffs

Conclusions on Game Theory

- Game theory is mathematically elegant, but there can be problems when applying it to real world problems:
 - Assumes opponents will play the equilibrium strategy
 - What to do with multiple Nash equilibria?
 - Computing Nash equilibria for complex games is nasty (perhaps even intractable)
 - Players have non-stationary policies
- Game theory used mainly to analyze environments at equilibrium rather than to control agents within an environment
- Also good for designing environments (mechanism design)

56

What you should know

- How to find pure strategy Nash Equilibria in a game
- Problems with having multiple Nash Equilibria
- How to compute mixed strategy Nash Equilibria in two-player constant sum games

57