

Foundations of AI

17. Strategic Games

Strategic Reasoning and Acting
 Wolfram Burgard and Bernhard Nebel

Strategic Game

- A **strategic game** G consists of
 - a finite set N (the set of **players**)
 - for each player $i \in N$ a non-empty set A_i (the set of **actions** or **strategies** available to player i), whereby $A = \prod_i A_i$
 - for each player $i \in N$ a function $u_i: A \rightarrow R$ (the **utility** or **payoff** function)
 - $G = (N, (A_i), (u_i))$
- If A is finite, then we say that the game is *finite*

Playing the Game

- Each player i makes a **decision** which action to play: a_i
- All players make their moves simultaneously leading to the **action profile** $a^* = (a_1, a_2, \dots, a_n)$
- Then each player gets the **payoff** $u_i(a^*)$
- Of course, each player tries to maximize its own payoff, but what is the right decision?
- **Note:** While we want to maximize our payoff, we are not interested in harming our opponent. It just does not matter to us what he will get!
 - If we want to model something like this, the payoff function must be changed

Notation

- For *2-player games*, we use a matrix, where the strategies of **player 1** are the **rows** and the strategies of **player 2** the **columns**
- The payoff for every action profile is specified as a pair x,y , whereby x is the value for player 1 and y is the value for player 2
- Example: For (T,R), **player 1** gets x_{12} , and **player 2** gets y_{12}

	Player 2 L action	Player 2 R action
Player1 T action	x_{11}, y_{11}	x_{12}, y_{12}
Player1 B action	x_{21}, y_{21}	x_{22}, y_{22}

Example Game: Bach and Stravinsky

- Two people want to out together to a concert of music by either Bach or Stravinsky. Their main concern is to go out together, but one prefers Bach, the other Stravinsky. Will they meet?
- This game is also called the *Battle of the Sexes*

	Bach	Stravinsky
Bach	2,1	0,0
Stravinsky	0,0	1,2

18/5

Example Game: Hawk-Dove

- Two animals fighting over some prey.
- Each can behave like a dove or a hawk
- The best outcome is if oneself behaves like a hawk and the opponent behaves like a dove
- This game is also called *chicken*.

	Dove	Hawk
Dove	3,3	1,4
Hawk	4,1	0,0

18/6

Example Game: Prisoner's Dilemma

- Two suspects in a crime are put into separate cells.
- If they both confess, each will be sentenced to 3 years in prison.
- If only one confesses, he will be freed.
- If neither confesses, they will both be convicted of a minor offense and will spend one year in prison.

	Don't confess	Confess
Don't confess	3,3	0,4
Confess	4,0	1,1

18/7

Solving a Game

- What is the right move?
- Different possible **solution concepts**
 - Elimination of strictly or weakly **dominated** strategies
 - **Maximin** strategies (for minimizing the loss in zero-sum games)
 - **Nash equilibrium**
- How difficult is it to compute a solution?
- Are there always solutions?
- Are the solutions unique?

18/8

Strictly Dominated Strategies

- **Notation:**
 - Let $a = (a_i)$ be a strategy profile
 - $a_{-i} := (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$
 - $(a_{-i}, a'_i) := (a_1, \dots, a_{i-1}, a'_i, a_{i+1}, \dots, a_n)$
- Strictly dominated strategy:
 - An strategy $a_j^* \in A_j$ is **strictly dominated** if there exists a strategy a'_j such that for all strategy profiles $a \in A$:

$$u_j(a_{-j}, a'_j) > u_j(a_{-j}, a_j^*)$$
- Of course, it is **not rational** to play **strictly dominated strategies**

Iterated Elimination of Strictly Dominated Strategies

- Since strictly dominated strategies will never be played, one can **eliminate** them from the game
- This can be done **iteratively**
- If this converges to a single strategy profile, the result is **unique**
- This can be regarded as the **result** of the game, because it is the **only rational outcome**

Iterated Elimination: Example

- Eliminate:
 - , dominated by
 - , dominated by
 - , dominated by
 - , dominated by
 - , dominated by

	b1	b2	b3	b4
a1	1,7	2,5	7,2	0,1
a2	5,2	3,3	5,2	0,1
a3	7,0	2,5	0,4	0,1
a4	0,0	0,-2	0,0	9,-1

➤ Result:

Iterated Elimination: Prisoner's Dilemma

- Player 1 reasons that “not confessing” is strictly dominated and eliminates this option
- Player 2 reasons that player 1 will not consider “not confessing”. So he will eliminate this option for himself as well
- So, they both confess

	Don't confess	Confess
Don't confess	3,3	0,4
Confess	4,0	1,1

Weakly Dominated Strategies

- Instead of strict domination, we can also go for weak domination:
 - An strategy $a_j^* \in A_j$ is **weakly dominated** if there exists a strategy a_j' such that for all strategy profiles $a \in A$:

$$u_j(a_{-j}, a_j') \geq u_j(a_{-j}, a_j^*)$$

and for at least one profile $a \in A$:

$$u_j(a_{-j}, a_j') > u_j(a_{-j}, a_j^*).$$

Results of Iterative Elimination of Weakly Dominated Strategies

- The result is not necessarily unique
- Example:
 - Eliminate
 - Eliminate:

	L	R
T	2,1	0,0
M	2,1	1,1
B	0,0	1,1

Analysis of the *Guessing 2/3 of the Average* Game

- All strategies above 67 are weakly dominated, since they will *never ever* lead to winning the prize, so they can be eliminated!
- This means, that all strategies above $\frac{2}{3} \times 67$ can be eliminated
- ... and so on
- ... until all strategies above 1 have been eliminated!
- So: The rationale strategy would be to play 1!

Existence of Dominated Strategies

- Dominating strategies are a convincing **solution concept**
- Unfortunately, often dominated strategies do not exist
- What do we do in this case?

	Dove	Hawk
Dove	3,3	1,4
Hawk	4,1	0,0

➤ **Nash equilibrium**

Nash Equilibrium

- A **Nash equilibrium** is an action profile $a^* \in A$ with the property that for all players $i \in N$:
 $u_i(a^*) = u_i(a^*_{-i}, a^*_i) \geq u_i(a^*_{-i}, a_i) \forall a_i \in A_i$
- In words, it is an action profile such that there is **no incentive** for any agent **to deviate** from it
- While it is less convincing than an action profile resulting from iterative elimination of dominated strategies, it is still a **reasonable solution concept**
- If there exists a **unique solution** from **iterated elimination of strictly dominated strategies**, then it is also a **Nash equilibrium**

18/17

Example Nash-Equilibrium: Prisoner's Dilemma

- Don't – Don't
– not a NE
- Don't – Confess (and vice versa)
– not a NE
- Confess – Confess
– NE

	Don't confess	Confess
Don't confess		0,4
confess	3,3	
Confess	4,0	1,1

18/18

Example Nash-Equilibrium: Hawk-Dove

- Dove-Dove:
– not a NE
- Hawk-Hawk
– not a NE
- Dove-Hawk
– is a NE
- Hawk-Dove
– is, of course, another NE
- So, NEs are not necessarily unique

	Dove	Hawk
Dove	3,3	1,4
Hawk	4,1	0,0

18/19

Auctions

- An **object** is to be **assigned** to a player in the set $\{1, \dots, n\}$ in exchange for a payment.
- Players i **valuation** of the object is v_i , and $v_1 > v_2 > \dots > v_n$.
- The mechanism to assign the object is a **sealed-bid auction**: the players simultaneously submit bids (non-negative real numbers)
- The object is given to the player with the lowest index among those who submit the highest bid in exchange for the payment
- The payment for a **first price** auction is the highest bid.
- What are the Nash equilibria in this case?

18/20

Formalization

- Game $G = (\{1, \dots, n\}, (A_i), (u_i))$
- A_i : bids $b_i \in \mathbb{R}^+$
- $u_i(b_{-i}, b_i) = v_i - b_i$ if i has won the auction, 0 otherwise
- Nobody would bid more than his valuation, because this could lead to negative utility, and we could easily achieve 0 by bidding 0.

18/21

Nash Equilibria for First-Price Sealed-Bid Auctions

- The Nash equilibria of this game are all profiles b with:
 - $b_i \leq b_1$ for all $i \in \{2, \dots, n\}$
 - No i would bid more than v_2 because it could lead to negative utility
 - If a b_i (with $< v_2$) is higher than b_1 , player 1 could increase its utility by bidding $v_2 + \epsilon$
 - So 1 wins in all NEs
 - $v_1 \geq b_1 \geq v_2$
 - Otherwise, player 1 either loses the bid (and could increase its utility by bidding more) or would have itself negative utility
 - $b_j = b_1$ for at least one $j \in \{2, \dots, n\}$
 - Otherwise player 1 could have gotten the object for a lower bid

18/22

Another Game: Matching Pennies

- Each of two people chooses either **Head** or **Tail**. If the choices differ, player 1 pays player 2 a euro; if they are the same, player 2 pays player 1 a euro.
- This is also a **zero-sum** or **strictly competitive** game
- No NE at all! What shall we do here?

	Head	Tail
Head	1,-1	-1,1
Tail	-1,1	1,-1

18/23

Randomizing Actions ...

- Since there does not seem to exist a rational decision, it might be best to **randomize** strategies.
- Play **Head** with probability p and **Tail** with probability $1-p$
- Switch to **expected utilities**

	Head	Tail
Head	1,-1	-1,1
Tail	-1,1	1,-1

18/24

Some Notation

- Let $G = (N, (A_i), (u_i))$ be a **strategic game**
- Then $\Delta(A_i)$ shall be the set of probability distributions over A_i - the set of **mixed strategies** $\alpha_i \in \Delta(A_i)$
- $\alpha_i(a_i)$ is the probability that a_i will be chosen in the mixed strategy α_i
- A profile $\alpha = (\alpha_i)$ of mixed strategies induces a probability distribution on A :
 $p(a) = \prod \alpha_i(a_i)$
- The expected utility is $U_i(\alpha) = \sum_{a \in A} p(a) u_i$ ^{18/25}

Example of a Mixed Strategy

- Let
 - $\alpha_1(H) = 2/3, \alpha_1(T) = 1/3$
 - $\alpha_2(H) = 1/3, \alpha_2(T) = 2/3$
- Then
 - $p(H,H) = 2/9$
 - $p(H,T) =$
 - $p(T,H) =$
 - $p(T,T) =$
 - $U_1(\alpha_1, \alpha_2) =$

	Head	Tail
Head	1,-1	-1,1
Tail	-1,1	1,-1

18/26

Mixed Extensions

- The **mixed extension** of the strategic game $(N, (A_i), (u_i))$ is the strategic game $(N, \Delta(A_i), (U_i))$.
- The **mixed strategy Nash equilibrium of a strategic game** is a Nash equilibrium of its mixed extension.
- Note that the **Nash equilibria in pure strategies** (as studied in the last part) are just a special case of mixed strategy equilibria.

18/27

Nash's Theorem

Theorem. Every finite strategic game has a mixed strategy Nash equilibrium.

- Note that it is essential that the game is **finite**
- So, there **exists** always a solution
- What is the **computational complexity**?
- **Identifying** a NE with a value larger than a particular value is **NP-hard**

18/28

The Support

- We call all pure actions a_i that are chosen with non-zero probability by α_i the **support** of the mixed strategy α_i

Lemma. Given a finite strategic game, α^* is a *mixed strategy equilibrium* if and only if for every player i every *pure strategy in the support* of α_i^* is a **best response** to α_{-i}^* .

18/29

Using the Support Lemma

- The **Support Lemma** can be used to compute all types of Nash equilibria in 2-person 2x2 action games.
 - There are 4 potential Nash equilibria in **pure strategies**
 - ❖ *Easy to check*
 - There are another 4 potential Nash equilibrium types with a **1-support** (pure) against **2-support** mixed strategies
 - ❖ Exists only if the **corresponding pure strategy profiles** are already Nash equilibria (follows from **Support Lemma**)
 - There exists one other potential Nash equilibrium type with a **2-support** against a **2-support** mixed strategies
 - ❖ Here we can use the **Support Lemma** to compute an NE (if there exists one)

18/30

A Mixed Nash Equilibrium for Matching Pennies

	Head	Tail
Head	1,-1	-1,1
Tail	-1,1	1,-1

- $U_1((1,0), (\alpha_2(H), \alpha_2(T))) = U_1((0,1), (\alpha_2(H), \alpha_2(T)))$
- $U_1((1,0), (\alpha_2(H), \alpha_2(T))) = 1\alpha_2(H) + -1\alpha_2(T)$
- $U_1((0,1), (\alpha_2(H), \alpha_2(T))) = -1\alpha_2(H) + 1\alpha_2(T)$
- $\alpha_2(H) - \alpha_2(T) = -\alpha_2(H) + \alpha_2(T)$
- $2\alpha_2(H) = 2\alpha_2(T)$
- $\alpha_2(H) = \alpha_2(T)$
- Because of $\alpha_2(H) + \alpha_2(T) = 1$:
 - $\alpha_2(H) = \alpha_2(T) = 1/2$
 - Similarly for player 1!
- ❖ $U_1(\alpha^*) = 0$

18/31

Mixed NE for BoS

	Bach	Stra-vinsky
Bach	2,1	0,0
Stra-vinsky	0,0	1,2

- $U_1((1,0), (\alpha_2(B), \alpha_2(S))) = U_1((0,1), (\alpha_2(B), \alpha_2(S)))$
- $U_1((1,0), (\alpha_2(B), \alpha_2(S))) = 2\alpha_2(B) + 0\alpha_2(S)$
- $U_1((0,1), (\alpha_2(B), \alpha_2(S))) = 0\alpha_2(B) + 1\alpha_2(S)$
- $2\alpha_2(B) = 1\alpha_2(S)$
- Because of $\alpha_2(B) + \alpha_2(S) = 1$:
 - $\alpha_2(B) = 1/3$
 - $\alpha_2(S) = 2/3$
- Similarly for player 1!
- ❖ $U_1(\alpha^*) = 2/3$

18/32

The 2/3 of Average Game

- You have n players that are allowed to choose a number between 1 and K .
- The players coming **closest to 2/3 of the average** over all numbers win. A fixed prize is **split equally** between all the winners
- What number would **you** play?
- What **mixed strategy** would you play?

18/33

A Nash Equilibrium in Pure Strategies

- All playing 1 is a NE in pure strategies
 - A deviation does not make sense
- All playing the same number different from 1 is **not a NE**
 - Choosing the number just below gives you more
- Similar, when all play different numbers, some not winning anything could get closer to 2/3 of the average and win something.
- So: **Why did you not choose 1?**
- Perhaps **you acted rationally** by assuming that the **others do not act rationally?**

18/34

Are there Proper Mixed Strategy Nash Equilibria?

- Assume there exists a mixed NE α different from the pure NE $(1,1,\dots,1)$
- Then there exists a maximal $k^* > 1$ which is played by some player with a probability > 0 .
 - Assume player i does so, i.e., k^* is in the support of α_i .
- This implies $U_i(k^*, \alpha_{-i}) > 0$, since k^* should be as good as all the other strategies of the support.
- Let a be a realization of α s.t. $u_i(a) > 0$. Then at least one other player must play k^* , because not all others could play below 2/3 of the average!
- In this situation player i could get more by playing k^*-1 .
- This means, playing k^*-1 is better than playing k^* , i.e., k^* cannot be in the support, i.e., **α cannot be a NE**

18/35

Summary

- **Strategic games** are one-shot games, where everybody plays its move simultaneously
- Each player gets a payoff based on its **payoff function** and the resulting **action profile**.
- **Iterated elimination of strictly dominated strategies** is a convincing solution concept.
- **Nash equilibrium** is another solution concept: Action profiles, where **no player has an incentive to deviate**
- It also might **not be unique** and there can be even infinitely many NEs or none at all!
- For every finite strategic game, there exists a Nash equilibrium in **mixed strategies**
- Actions in the support of mixed strategies in a NE are always best answers to the NE profile, and therefore have the same payoff \leadsto **Support Lemma**
- Computing a NE in mixed strategies is NP-hard

18/36