Computational Game Theory (BE4M36MAS)

Normal Form Games

Branislav Bošanský

Artificial Intelligence Center,
Department of Computer Science,
Faculty of Electrical Engineering,
Czech Technical University in Prague

branislav.bosansky@agents.fel.cvut.cz

October 5, 2021

Previously ... on computational game theory.

1. What is game theory?
2. Classification of games
3. Applications, Motivation

Today, you should learn ...

1. Baseline formal representation for games.
2. What does it mean "solution of the game"?
3. Basic solution concepts.

Games in Game Theory

What do we need to specify if we want to talk about (almost any) game:

- **Who?** – Which **agents (players)** are participating in the game?
- **What?** – What are the **actions** the agents can choose to play? What is the **outcome** of the game if agents choose their actions? What do the players **know** during the game?

Formal Representation of Games

There are many possible formal representations for games (we will see later). **Normal-form** (or matrix) representation is the most basic one.

Definition (Normal Form Game (NFG))

We call a triplet $\mathcal{G} = (\mathcal{N}, \mathcal{A}, u)$ *a normal-form game*, where

- $\mathcal{N}$ is a finite set of players, we use $n = |\mathcal{N}|$,
- $\mathcal{A}_i$ is a finite set of actions (pure strategies; hence, we also use $\mathcal{S}_i$ in some definitions) for player i,
- u_i is a utility function of player i that assigns the reward for joint action $a \in \mathcal{A}$, $a = (a_1, a_2, \dots, a_n)$ to player i.

We assume that players are **rational** and they only maximize their expected utility value. (there are parts of game theory that deal with imperfectly rational players)

Normal-Form Game Examples

Rock Paper Scissors

	R	**P**	**S**
R	$(0, 0)$	$(-1, 1)$	$(1, -1)$
P	$(1, -1)$	$(0, 0)$	$(-1, 1)$
S	$(-1, 1)$	$(1, -1)$	$(0, 0)$

Prisoners' Dilemma

	C	**D**
C	$(-1, -1)$	$(0, -5)$
D	$(-5, 0)$	$(-4, -4)$

Normal-Form Game Examples

Matching Pennies

	H	**T**
H	$(1, -1)$	$(-1, 1)$
T	$(-1, 1)$	$(1, -1)$

Battle of Sexes

	M	**F**
M	$(1, 2)$	$(0, 0)$
F	$(0, 0)$	$(2, 1)$

Classes of Games

Depending on players, actions, and outcomes, there are many different classes of games:

- Depending on the number of players, we can focus on 2, 3, or n-player games.
- Games can be **one-shot** or **dynamic** (sequential) with **finite** (or **infinite**) **horizon**.
- Games can be with **perfect** or **imperfect information**.
- Games can be **zero-sum** or **general-sum**.
- Games can be discrete or continuous (any of the set of players, actions, set of states can be infinitely large).

Strategies in Games

Choices players make in a game are called **strategies** (they do not necessarily to correspond to only a single action).

- We denote s_i to be a strategy of player $i \in \mathcal{N}$. $\mathcal{S}_i$ is a set of all strategies of player i.
- A set of strategies of all players is called a **strategy profile**

$$s = \langle s_1, s_2, \ldots, s_n \rangle$$

- Often, we need to refer to strategies all other player except player $i \in \mathcal{N}$:

$$s_{-i} = \langle s_1, s_2, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n \rangle$$

.

What Strategies Should an Agent Play?

What is a desirable outcome of a game?

- An outcome s, such that there is no other outcome s' where one player would be better of and all other players have at least the utility as in s – **pareto optimal outcome**
- An outcome that maximizes the sum of all players – **social welfare optimization**

	L	**C**	**R**
T	$(1,0)$	$(-1,1)$	$(1,-1)$
M	$(2,2)$	$(0,0)$	$(3,1)$
B	$(-1,1)$	$(1,-1)$	$(0,3)$

What Strategies Should an Agent Play?

Some strategies can be better than others.

	C	**D**
A	$(2, 1)$	$(3, 4)$
B	$(-1, 0)$	$(1, 1)$

Which strategy would you recommend to be played?

Strategy **A** yields a better outcome for player 1 than strategy **B** regardless of the action of player 2.

We say that strategy **A dominates** strategy **B**
(or that strategy **B is dominated** by strategy **A**)

Dominance

Definition (Strong Dominance)

Let $\mathcal{G} = (\mathcal{N}, \mathcal{A}, u)$ be a normal-form game. We say that s_i *strongly dominates* s_i' if $\forall s_{-i} \in \mathcal{S}_{-i}, u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$.

Definition (Weak Dominance)

Let $\mathcal{G} = (\mathcal{N}, \mathcal{A}, u)$ be a normal-form game. We say that s_i *weakly dominates* s_i' if $\forall s_{-i} \in \mathcal{S}_{-i}, u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$ and $\exists s_{-i} \in \mathcal{S}_{-i}$ such that $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$.

Definition (Very Weak Dominance)

Let $\mathcal{G} = (\mathcal{N}, \mathcal{A}, u)$ be a normal-form game. We say that s_i *very weakly dominates* s_i' if $\forall s_{-i} \in \mathcal{S}_{-i}, u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$.

Removal Dominated Strategies

Question

Would a rational agent play a (strongly/weakly/very weakly) dominated strategy?

Rational agent would never choose a strongly dominated strategy, hence we can remove those strategies from the game.

Iterative Removal of Dominated Strategies – a simple algorithm that iteratively removes (strongly) dominated strategies from a game.

Removal Dominated Strategies

Iterative Removal of Dominated Strategies

	L	**C**	**R**
T	$(1, 0)$	$(-1, 1)$	$(1, -1)$
M	$(2, 2)$	$(0, 0)$	$(3, 1)$
B	$(-1, 1)$	$(1, -1)$	$(0, 3)$

- **T** is dominated by **M**
- **C** is then dominated by **R** (and not before)
- **B** is then dominated by **M** (and not before)
- **R** is then dominated by **L** (and not before)

What Strategies Should an Agent Play? – Deviations

While players are rational, they may not choose to play the *best outcome* (in pareto or social-welfare sense). Given a strategy of the opponents s_{-i}, if there is a better strategy for player i, he is going to deviate:

	C	**D**
A	$(5,5)$	$(0,6)$
B	$(6,0)$	$(1,1)$

$$(A,C) \rightarrow (B,C) \rightarrow (B,D)$$

Definition (Best Response)

Let $\mathcal{G} = (\mathcal{N}, \mathcal{A}, u)$ be a normal-form game and let $BR_i(s_{-i}) \subseteq \mathcal{S}_i$ such that $s_i^* \in BR_i(s_{-i})$ iff $\forall s_i \in \mathcal{S}_i, u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$.

Nash Equilibrium

John Forbes Nash Jr. (1928 - 2015)
Nash (1950) "Equilibrium Points in N-person Games". Proceedings of the National Academy of Sciences of the United States of America.

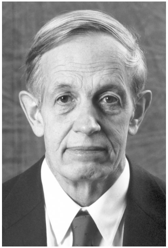

Definition (Nash Equilibrium)

Let $\mathcal{G} = (\mathcal{N}, \mathcal{A}, u)$ be a normal-form game. Strategy profile $s = \langle s_1, \ldots, s_n \rangle$ is a Nash equilibrium iff $\forall i \in \mathcal{N}, s_i \in BR_i(s_{-i})$.

Nash Equilibrium

Question

How do we look for a Nash equilibrium?

So far, we considered only actions being played in a game. Hence, if all players choose a strategy, exactly one outcome is selected.

It is sufficient to check whether there is some agent that wants to deviate or not.

If not, this outcome is a Nash equilibrium.

Nash Equilibrium

What are Nash equilibria in these games?

	C	**D**
A	$(5, 5)$	$(0, 6)$
B	$(6, 0)$	$(1, 1)$

	L	**C**	**R**
T	$(1, 0)$	$(-1, 1)$	$(1, -1)$
M	$(2, 2)$	$(0, 0)$	$(3, 1)$
B	$(-1, 1)$	$(1, -1)$	$(0, 3)$

Rock Paper Scissors

	R	**P**	**S**
R	$(0,0)$	$(-1,1)$	$(1,-1)$
P	$(1,-1)$	$(0,0)$	$(-1,1)$
S	$(-1,1)$	$(1,-1)$	$(0,0)$

What is the best strategy to play in Rock-Paper-Scissors?

Every time we are about to play we randomly select an action we are going to use.

The concept of pure strategies is not sufficient.

Mixed Strategies

Definition (Mixed Strategies)

Let $\mathcal{G} = (\mathcal{N}, \mathcal{A}, u)$ be a normal-form game. Then the set of *mixed strategies* $\mathcal{S}_i$ for player i is the set of all probability distributions over $\mathcal{A}_i$; $\mathcal{S}_i = \Delta(\mathcal{A}_i)$.

Player selects a pure strategy according to the probability distribution.

We extend the utility function to correspond to *expected utility*:

$$u_i(s) = \sum_{a \in A} u_i(a) \prod_{j \in \mathcal{N}} s_j(a_j)$$

We can extend existing concepts (dominance, best response, ...) to mixed strategies.

Existence of Nash equilibria?

	M	**F**
M	$(1, 2)$	$(0, 0)$
F	$(0, 0)$	$(2, 1)$

	R	**P**	**S**
R	$(0, 0)$	$(-1, 1)$	$(1, -1)$
P	$(1, -1)$	$(0, 0)$	$(-1, 1)$
S	$(-1, 1)$	$(1, -1)$	$(0, 0)$

Theorem (Nash)

Every game with a finite number of players and action profiles has at least one Nash equilibrium in mixed strategies.

Nash Equilibrium

Characteristics of a Nash equilibrium (NE)

- NE is a descriptive solution concept – it *describes* which strategy profile is stable, it does not prescribe which strategies the players should be playing!
- NE is generally not unique and there may exist many NE. If one agent plays a strategy from a NE strategy profile, there are generally no guarantees on an (expected) outcome.
- NE is optimal in a sense of unilateral deviations. Strong NE is a variant that is optimal in a sense of group deviations.

Regret

The concept of regret is useful when the utility of other players is unknown.

	L	**R**
U	$(100, a)$	$(1 - \varepsilon, b)$
D	$(2, c)$	$(1, d)$

Definition (Regret)

A player i's *regret* for playing an action a_i if the other agents adopt action profile a_{-i} is defined as

$$\left[\max_{a'_i \in \mathcal{A}_i} u_i(a'_i, a_{-i})\right] - u_i(a_i, a_{-i})$$

Regret

Definition (MaxRegret)

A player is *maximum regret* for playing an action a_i is defined as

$$\max_{a_{-i} \in \mathcal{A}_{-i}} \left(\left[\max_{a'_i \in \mathcal{A}_i} u_i(a'_i, a_{-i}) \right] - u_i(a_i, a_{-i}) \right)$$

Definition (MinimaxRegret)

Minimax regret actions for player i are defined as

$$\arg\min_{a_i \in \mathcal{A}_i} \max_{a_{-i} \in \mathcal{A}_{-i}} \left(\left[\max_{a'_i \in \mathcal{A}_i} u_i(a'_i, a_{-i}) \right] - u_i(a_i, a_{-i}) \right)$$

Correlated Equilibrium

Consider again the following game:

	L	**R**
U	$(2,1)$	$(0,0)$
D	$(0,0)$	$(1,2)$

Wouldn't it be better to coordinate 50:50 between the outcomes (U,L) and (D,R)? Can we achieve this coordination? We can use *a correlation device*—a coin, a streetlight, commonly observed signal—and use this signal to avoid unwanted outcomes.

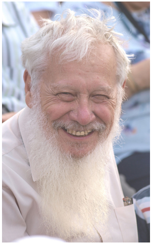

Robert Aumann

Correlated Equilibrium

Definition (Correlated Equilibrium (simplified))

Let $\mathcal{G} = (\mathcal{N}, \mathcal{A}, u)$ be a normal-form game and let σ be a probability distribution over joint pure strategy profiles $\sigma \in \Delta(\mathcal{A})$. We say that σ is a correlated equilibrium if for every player i, every signal $a_i \in \mathcal{A}_i$ and every possible action $a_i' \in \mathcal{A}_i$ it holds

$$\sum_{a_{-i} \in \mathcal{A}_{-i}} \sigma(a_i, a_{-i}) u_i(a_i, a_{-i}) \geq \sum_{a_{-i} \in \mathcal{A}_{-i}} \sigma(a_i, a_{-i}) u_i(a_i', a_{-i})$$

Corollary

For every Nash equilibrium there exists a corresponding Correlated Equilibrium.

Stackelberg Equilibrium

Finally, consider a situation where an agent is a central public authority (police, government, etc.) that needs to design and publish a policy that will be observed and reacted to by other agents.

- *the leader* – publicly commits to a strategy
- *the follower(s)* – play a Nash equilibrium with respect to the commitment of the leader

Stackelberg equilibrium is a strategy profile that satisfies the above conditions and maximizes the expected utility value of the leader:

$$\arg\max_{s \in \mathcal{S}; \forall i \in \mathcal{N} \setminus \{1\} s_i \in BR_i(s_{-i})} u_1(s)$$

Stackelberg Equilibrium

Consider the following game:

	$\mathbf{L}$	$\mathbf{R}$
$\mathbf{U}$	$(4,2)$	$(6,1)$
$\mathbf{D}$	$(3,1)$	$(5,2)$

$(\mathbf{U}, \mathbf{L})$ is a Nash equilibrium.

What happens when the row player commits to play strategy $\mathbf{D}$ with probability 1? Can the row player get even more?

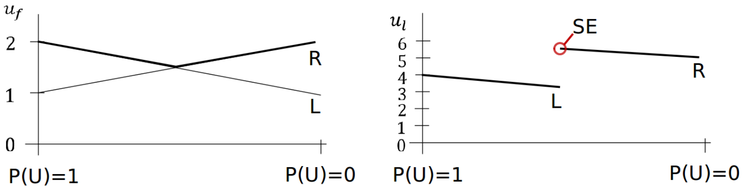

There may be Multiple Nash Equilibria

The followers need to break ties in case there are multiple NE:

- arbitrary but fixed tie breaking rule
- *Strong SE* – the followers select such NE that maximizes the outcome of the leader (when the tie-braking is not specified we mean SSE),
- *Weak SE* – the followers select such NE that minimizes the outcome of the leader.

Exact Weak Stackelberg equilibrium does not have to exist.

Different Stackelberg Equilibria

Exact Weak Stackelberg equilibrium does not have to exist.

$1 \setminus 2$	a	b	c	d	e
T	$(2,4)$	$(6,4)$	$(9,0)$	$(1,2)$	$(7,4)$
B	$(8,4)$	$(0,4)$	$(3,6)$	$(1,5)$	$(0,0)$