Algorithmic Game Theory - Introduction, Complexity, Nash

Branislav Bošanský

Czech Technical University in Prague

branislav.bosansky@agents.fel.cvut.cz

February 25, 2018

About This Course

main topics of the course:

- dig deeper into game theory
- analyze the algorithmic and computational aspect of the problems in game theory
  - equilibrium computation algorithms (exact and approximate)
  - computational complexity (PLS, PPAD, FIXP, NP, $\Delta_2^P = P^{NP}$)
- extended foundations of game theory
- main theorems, their impact, generalization
- you

Grading: homework assignments (at least 2 correct out of 4) and presentation on a selected topic ($1/3$ of a research paper).

`https://cw.fel.cvut.cz/wiki/courses/xep36agt/lectures/start`

Outline of the Course

1. Introduction, Definitions
2. Nash's Theorem, Main Complexity Classes (PLS, PPAD, FIXP)
3. Computing a Nash Equilibrium (Lemke Howson, MILP)
4. Approximating Nash Equilibria
5. Computing Stackelberg Equilibria
6. Computing and Approximating Correlated Equilibria
7. Succinct Representation of Games, Correlated Equilibrium in Succinct Games
8. Repeated Games
9. Multiarmed Bandit Problems
10. Learning in Normal-Form Games, Fictitious Play
11. Regret Matching, Counterfactual Regret Minimization
12. Continual Resolving in Extensive-Form Games (DeepStack)

Standard Representation of Games

standard normal-form representation – a game is a tuple $(\mathcal{N}, \mathcal{S}, u)$

- $\mathcal{N}$ is a set of players $i \in \mathcal{N} = \{1, \ldots, n\}$, $-i$ denotes all other players except i.
- $\mathcal{S}$ is a set of actions (pure strategies) $\mathcal{S} = \times_i \mathcal{S}_i$ (we often use $|\mathcal{S}_i| = m_i$)
- u_i is a utility function $u_i : \mathcal{S} \rightarrow \mathbb{R}$ (sometimes there is a cost function $c_i : \mathcal{S} \rightarrow \mathbb{R}, u_i(s) = -c_i(s)$)

also-known-as: strategic form, matrix form

we will refer to them as NFGs

in case of only two players: *bimatrix games*

Strategies

standard normal-form representation – a game is a tuple $(\mathcal{N}, \mathcal{S}, u)$

- *pure strategies* – $s_i \in \mathcal{S}_i$ (can be infinite)
- *mixed strategies* – probability distributions over pure strategies $\Delta(\mathcal{S}_i) = \left\{ p^i \in \mathbb{R}^{|\mathcal{S}_i|} \mid \sum_{j=1}^{|\mathcal{S}_i|} p^i_j = 1 \wedge p^i_j \geq 0 \right\}$, denoted σ
- *behavioral strategies* – vector of probability distributions over actions to play in each decision step
- *convex strategies* – arbitrary convex set $X \subseteq \mathbb{R}^{|\mathcal{S}|}$
- counting strategies, strategies with states, memory strategies, turing machine strategies

Beyond Standard Representation of Games

There are other representations that capture specific types of games more compactly compared to NFGs:

- *extensive-form games* – finite sequential games (but there are also Bayesian games, multi-agent influence diagrams (MAIDs), LIMIDs, ...)
- *stochastic games* – infinite sequential games (but we also have repeated games)
- *congestion games* – abstract the network congestion games
  - We have n players, set of edges E, strategies for each player are *paths* in the network ($\mathcal{S}$), and there is a congestion function $c_e : \{0, 1, \dots, n\} \to \mathbb{Z}^+$. When all players choose their strategy path $s_i \in \mathcal{S}_i$ we have the load of edge e, $\ell(e) = |\{s_i | e \in s_i\}|$ and $u_i = \sum_{e \in s_i} c_e(\ell(e))$

Beyond Standard Representation of Games (2)

- *graphical games* – n-player games where the utility of one player typically depends only on few other players. They are represented as a graph, where agents are vertices and edge corresponds to the dependance between the two players. If the maximum degree of the graph is small ($d \ll n$), this representation offers exponentially smaller input $ns^{d+1} \ll ns^n$
- *action graph games* – even finer dependance than in graphical games based on actions
- *polymatrix games* – specific graphical games, where we consider a bimatrix game for each edge (i.e., only pairwise interactions); quadratic size in ns
- *anonymous games*, *symmetric games*, ...

Continuous/Infinite Games

games over the unit square

- X, Y are set of "pure strategies" equal to interval $[0, 1]$
- $\mathcal{M}_X, \mathcal{M}_Y$ are probability distributions over X, Y
- we can reason about them similarly (although using calculus) to discrete games
- very useful in auctions, adversarial machine learning, any time you have a naturally infinite strategy space

Example: zero-sum game, $X = [0, 1]; Y = [0, 1]$, the payoff function is

$$u(x, y) = 4xy - 2x - y + 3, \qquad \forall x \in X,\ y \in Y$$

Why do we care?

One representation does not rule them all.

Depending on the representation we can get an exponential speed-up for specific types of problems.

Even if not, algorithms that work with compact representations can be a starting point if you are looking for an approximate solution to the original problem.

Solution Concepts

we want to find optimal strategies according to different notions of optimality:

- *maxmin strategies* – $\max_{s_i \in \mathcal{S}_i} \min_{s_{-i} \in \mathcal{S}_{-i}} u_i(s_i, s_{-i})$
- *minmax strategies* – $\min_{s_{-i} \in \mathcal{S}_{-i}} \max_{s_i \in \mathcal{S}_i} u_i(s_i, s_{-i})$
- can be defined for any type of strategies

if we seek *minmax strategies* over infinite sets, maximum or minimum over function $u_i(s_i, s_{-i})$ might not exist

- $\sup_{s_i \in \mathcal{S}_i} \inf_{s_{-i} \in \mathcal{S}_{-i}} u_i(s_i, s_{-i})$

$$\max_{s_i \in \mathcal{S}_i} \min_{s_{-i} \in \mathcal{S}_{-i}} u_i(s_i, s_{-i}) \leq \min_{s_{-i} \in \mathcal{S}_{-i}} \max_{s_i \in \mathcal{S}_i} u_i(s_i, s_{-i})$$

Solution Concepts (2)

stable solution concepts

- *best response* – let σ_{-i} be a strategy of players $-i$, $\max_{s_i \in \mathcal{S}_i} u_i(s_i, \sigma_{-i})$
  - we can define pure, mixed, behavioral best response
  - it is not always true that pure best responses are sufficient
  - $BR_i(\sigma_{-i})$ is a set of all best responses
- *Nash Equilibrium* – a strategy profile σ where every player is playing the best response to the strategies of other players; $\sigma_i \in BR_i(\sigma_{-i})$
- *(Strong) Stackelberg Equilibrium* – a strategy profile σ that maximizes the expected utility of player 1 (*leader*) where all other players (*followers*) are playing Nash Equilibrium;

$$\arg\max_{\sigma; \forall i \in \mathcal{N} \setminus \{1\}, \sigma_i \in BR_i(\sigma_{-i})} u_1(\sigma)$$

Solution Concepts (3)

- *Correlated Equilibrium* – a probability distribution over pure strategy profiles $p = \Delta(\mathcal{S})$ that recommends each player i to play the best response; $\forall s_i, s_i' \in \mathcal{S}_i$:

$$\sum_{s_{-i} \in \mathcal{S}_{-i}} p(s_i, s_{-i}) u_i(s_i, s_{-i}) \geq \sum_{s_{-i} \in \mathcal{S}_{-i}} p(s_i, s_{-i}) u_i(s_i', s_{-i})$$

- *Coarse Correlated Equilibrium* – a probability distribution over pure strategy profiles $p = \Delta(\mathcal{S})$ that **in expectation** recommends each player i to play the best response; $\forall s_i \in \mathcal{S}_i$:

$$\sum_{s' \in \mathcal{S}'} p(s') u_i(s') \geq \sum_{s' \in \mathcal{S}'} p(s') u_i(s_i, s_{-i}')$$

- *Quantal Response Equilibrium* – modeling bounded rationality

$$p_j^i = \frac{\exp(u_i(s_j, \sigma_{-i}))}{\sum_{s_j' \in \mathcal{S}_i} \exp(u_i(s_j', \sigma_{-i}))}$$

Assumptions on Utilities

we can restrict to games with a specific utility function

- *zero-sum games* – meaningful for two-player games, where $u_1(s_1, s_2) = -u_2(s_1, s_2)$
- *almost zero-sum games* – games where there is an additional cost for one player $u_1(s_1, s_2) = -u_2(s_1, s_2) - c'(s_1)$
- *strategically zero-sum games* – let $A, B \in \mathbb{R}^{m_1 \times m_2}$ be the matrices of a bimatrix game. A game is SZS iff there exist $\alpha, \beta > 0$ and $D \in \mathbb{R}^{m_1 \times m_2}$ such that

$$\begin{aligned} \alpha A &= D + [\mathbf{b}^T, \mathbf{b}^T, \ldots, \mathbf{b}^T]^T \\ \beta B &= -D + [\mathbf{a}, \mathbf{a}, \ldots, \mathbf{a}] \end{aligned}$$

  for some $\mathbf{a} \in \mathbb{R}^{m_1}, \mathbf{b} \in \mathbb{R}^{m_2}$.
- *security games*, ...

Let the game begin

Complexity Classes

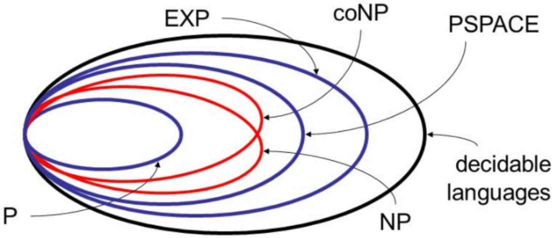

all containments believed to be proper

Main Complexity Classes in Algorithmic Game Theory

$\mathcal{P}$

- polynomial problems (linear programming of polynomial size, etc.)
- computing a Nash Equilibrium in zero-sum games (normal-form, extensive-form)
- computing a Correlated Equilibrium in general-sum games (normal-form[1])
- computing a Stackelberg Equilibrium in general-sum games (normal-form, simple security games)
- computing equilibria in many of the instances from succinctly represented games (we shall see)

---

[1] open for extensive-form games

Main Complexity Classes in Algorithmic Game Theory

$\mathcal{NP}$

- NP-hard problems, (mixed-integer linear programs, etc.)
- computing some specific Nash Equilibrium in general-sum games (normal-form, extensive-form)
- computing some specific Correlated-based Equilibrium in general-sum games (extensive-form games)
- computing a Stackelberg Equilibrium in general-sum games (extensive-form games)

Main Complexity Classes in Algorithmic Game Theory

the search problem problem that asks for *any Nash Equilibrium* is a different, potentially 'easier' problem

there is a finer description of the complexity classes

Main Complexity Classes in Algorithmic Game Theory

1. equilibria are guaranteed to exist (i.e., total problems TFNP $\subseteq$ FNP ("a function extension of a decision problem in NP")
2. we can search for them
   - pure strategy profiles
   - support enumeration

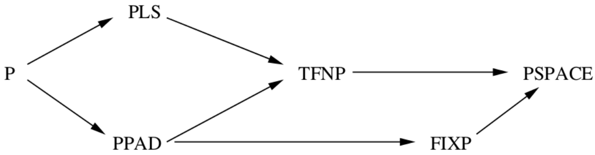

Fig. 2 – Relations between the complexity classes. [2]

[2] Figure from M. Yannakakis "Equilibria, fixed points, and complexity classes" Computer Science Review (3) 71–85, 2009

Polynomial Local Search ($\mathcal{PLS}$)

- consider an instance I of an optimization problem, $S(I)$ is a set of candidate solutions, $p_I(s)$ is a cost (or utility) associated with candidate $s \in S(I)$ that has to be minimized (or maximized, respectively)
- each candidate $s \in S(I)$ has a neighborhood $N_I(s) \subseteq S(I)$
- a candidate s is locally optimal (cost-wise) if

$$p_I(s) \leq p_I(s') \; \forall s' \in N_I(s)$$

- Sol(I) is a set of locally optimal solutions
- every step of the algorithm (generating starting solution, computing the cost, getting a better neighbor) is polynomial, but there can be an exponential number of steps

Polynomial Local Search ($\mathcal{PLS}$) (2)

- several well-known problems of this kind
- finding a local optimum in Traveling Salesman Problem, Max Cut, Max Sat, ...
  - we define a neighborhood function (e.g., *2-Opt*) and perform a greedy search
- finding a stable configuration of a neural network
- finding a pure equilibrium when it is guaranteed to exist
  - computing an optimal strategy in simple stochastic games, where pure stationary strategy is known to be optimal (the problem is in PLS, but it is open whether it is in P, or not)
  - some other variants of stochastic games (mean payoff/parity games with no chance)

From PLS to PPAD

Searching for pure Nash equilibria is not sufficient.

Pure equilibria do not have to exist.

What can we search for in mixed strategies?

Is this problem in PLS? Can we redefine the problem of finding a mixed NE as a PLS?

References I

(besides the books)

X. Chen and X. Deng.
Settling the complexity of two-player nash equilibrium.
In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 261–272, 2006.

X. Chen, X. Deng, and S.-H. Teng.
Computing Nash equilibria: Approximation and smoothed complexity.
In *Proc. 47th IEEE FOCS*, 2006.

X. Chen, S.-H. Teng, and P. Valiant.
The approximation complexity of winlose games.
In *Proc. 18th ACM SODA*, 2007.

C. Daskalakis, A. Fabrikant, and C. H. Papadimitriou.
The Game World Is Flat: The Complexity of Nash Equilibria in Succinct Games.
In *ICALP*, pages 513–524, 2006.

References II

C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou.
The Complexity of Computing a Nash Equilibrium.
In *Proceedings of the 38th annual ACM symposium on Theory of computing*, 2006.

K. Etessami and M. Yannakakis.
On the complexity of nash equilibria and other fixed points.
In *FOCS*, 2007.

Z. H. Gumus and C. A. Floudas.
Global optimization of mixed-integer bilevel programming problems.
Computational Management Science, 2:181–212, 2005.