Continuous games II

Polynomial Games. Reduction to an SDP Problem.

Tomáš Kroupa

Faculty of Electrical Engineering
Artificial Intelligence Center

CZECH TECHNICAL UNIVERSITY IN PRAGUE

Agenda

1. Polynomial games. Properties of solutions. Minimax theorem.
2. An algorithm to solve polynomial games based on SDP.

This presentation is based on

P. Parrilo.
Polynomial games and sum of squares optimization.
45th IEEE Conference on Decision and Control, 2855–2860, 2006.

Polynomial games and their solutions

Polynomial games

A strategic game $(N, (S_i)_{i \in N}, (u_i)_{i \in N})$ is polynomial if

- Each strategy set $S_i \subseteq \mathbb{R}^{d_i}$ is **compact**, where $d_i \in \mathbb{N}$
- Each utility function $u_i \colon S \to \mathbb{R}$ is **polynomial**, $S := S_1 \times \cdots \times S_n$

Fundamental facts

- Polynomial games are separable, so there is an **atomic equilibrium** $(\mu_1, \ldots, \mu_n)$ such that $|\operatorname{spt} \mu_i| \leq m_i + 1$ for all $i \in N$
- Each μ_i is a convex combination of at most $m_i + 1$ pure strategies

Two-person zero-sum polynomial games

Our setting

- $N = \{1, 2\}$
- $S_1 = S_2 = [-1, 1]$
- $u\colon [-1, 1]^2 \to \mathbb{R}$ is a bivariate polynomial, put $u_1 := u$
- $u_1 + u_2 = 0$

The payoff of the 1st player is a polynomial of the form

$$u(x, y) = \sum_{i=0}^{k} \sum_{j=0}^{m} c_{ij} x^i y^j, \qquad x, y \in [-1, 1].$$

Two-person zero-sum polynomial games – minimax theorem

The game is determined by the function

$$u(x,y) = \sum_{i=0}^{k}\sum_{j=0}^{m} c_{ij}x^i y^j, \qquad x,y \in [-1,1].$$

Dresher, Karlin, Shapley (1950)

There exists an equilibrium (μ^*, σ^*) satisfying the saddle-point condition, where $|\operatorname{spt}\mu^*| \leq k+1$ and $|\operatorname{spt}\sigma^*| \leq m+1$.

Analytical solution of a two-player zero-sum polynomial game

$$u(x, y) = (x - y)^2$$

Atomic equilibrium (μ^*, σ^*) exists.

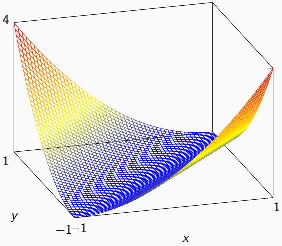

- The support of σ^* may contain only those $y \in [-1, 1]$ minimizing the function $U(\mu^*, y)$, which is **strictly convex**. Its unique minimizer is 0, hence $\sigma^* = \delta_0$
- The support of μ^* may contain only those $x \in [-1, 1]$ maximizing the function $U(x, \sigma^*) = u(x, 0) = x^2$, hence $\operatorname{spt} \mu^* \subseteq \{-1, 1\}$
- The value of game is $V = \max_{x \in [-1,1]} U(x, \sigma^*) = 1$
- The unique optimal strategy is $\mu^* = \frac{1}{2}(\delta_{-1} + \delta_1)$

Algorithm for polynomial games

Semidefinite programming

$\mathcal{S}^n$	real symmetric matrices $n \times n$
$\mathbf{X} \in \mathcal{S}^n$	matrix of variables
$\mathbf{X} \succeq 0$	matrix $\mathbf{X}$ is positive semidefinite
$\mathbf{C} \in \mathbb{R}^{n \times n}$	matrix of coefficients
a_{ijk}, b_k	real coefficients

SD program

$$\text{Maximize} \qquad \sum_{i=1}^{n} \sum_{j=1}^{n} c_{ij} x_{ij} \qquad \text{subject to}$$

$$\sum_{i=1}^{n} \sum_{j=1}^{n} a_{ijk} x_{ij} = b_k \qquad k = 1, \ldots, m$$

$$\mathbf{X} \succeq 0$$

SDPs and their solutions

- Linear programming is a special case
- Convex optimization problem
- Efficiently solvable by interior point methods
- Many solvers are available nowadays
- Algorithms find a solution up to an error ε in the time polynomial in the problem description and $\log \varepsilon^{-1}$

Formulation of an optimization problem

LP formulation for matrix games (variables v and σ)

$$\text{Minimize } v \quad \text{subject to} \quad \begin{cases} U(x,\sigma) \le v & \forall x \in S_1 \\ \sigma \in \Delta(S_2) \end{cases}$$

By analogy with matrix games we try to compute the optimal strategy σ of Player 2 in a polynomial game as follows:

Formulation in our setting (variables v and σ)

$$\text{Minimize } v \quad \text{subject to} \quad \begin{cases} U(x,\sigma) \le v & \forall x \in [-1,1] \\ \sigma \in \Delta([-1,1]) \end{cases}$$

This is an infinite-dimensional optimization problem.

Reduction to a finite-dimensional problem (1)

For any $x \in [-1, 1]$ and any $\sigma \in \Delta([-1, 1])$:

$$U(x, \sigma) = \int_{-1}^{1} u(x, y)\, d\sigma(y) = \sum_{i=0}^{k} \sum_{j=0}^{m} c_{ij} x^i \overbrace{\left(\int_{-1}^{1} y^j\, d\sigma(y) \right)}^{\sigma_j :=}$$

$$= \sum_{i=0}^{k} \sum_{j=0}^{m} c_{ij} \sigma_j x^i$$

This is a univariate polynomial!

Finite-dimensional problem with variables v and $(\sigma_0, \ldots, \sigma_m)$

Minimize v subject to
$$\begin{cases} v - \sum_{i=0}^{k} \sum_{j=0}^{m} c_{ij} \sigma_j x^i \geq 0 & \forall x \in [-1, 1] \\ (\sigma_0, \ldots, \sigma_m) \in \mathbb{R}^{m+1} \text{ is a tuple of moments} \end{cases}$$

Reduction to a finite-dimensional problem (2)

We need to find a computationally convenient description of these sets:

1. Univariate polynomials nonnegative in $[-1, 1]$
2. Moment space in $\mathbb{R}^{m+1}$ of probability measures over $[-1, 1]$

Method

- The problem above leads to a single **semidefinite program**
- Its solution provides the value and the moments of optimal strategies for the 2nd player
- The **dual program** computes the moments of optimal strategies for the 1st player

Hankel matrix

$\mathcal{H}$ linear operator $\mathbb{R}^{2\ell-1} \to \mathcal{S}^\ell$

$\mathcal{H}^*$ adjoint linear operator $\mathcal{S}^\ell \to \mathbb{R}^{2\ell-1}$

$$\mathcal{H}\colon \begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_{2\ell-1} \end{pmatrix} \mapsto \begin{pmatrix} a_1 & a_2 & \dots & a_\ell \\ a_2 & a_3 & \dots & a_{\ell+1} \\ \vdots & \vdots & \ddots & \ddots \\ a_\ell & a_{\ell+1} & \dots & a_{2\ell-1} \end{pmatrix}$$

$$\mathcal{H}^*\colon \begin{pmatrix} a_{11} & a_{12} & \dots & a_{1\ell} \\ a_{12} & a_{22} & \dots & a_{2\ell} \\ \vdots & \vdots & \ddots & \ddots \\ a_{1\ell} & a_{2\ell} & \dots & a_{\ell\ell} \end{pmatrix} \mapsto \begin{pmatrix} a_{11} \\ 2a_{12} \\ a_{22} + 2a_{13} \\ \vdots \\ a_{\ell\ell} \end{pmatrix}$$

When is a polynomial $p \in \mathbb{R}[x]$ nonnegative?

Proposition

The following are equivalent:

- $p(x) \geq 0$ for all $x \in \mathbb{R}$
- p is a sum of squares: $p = \sum_{i=1}^{\ell} q_i^2$ for some $q_i \in \mathbb{R}[x]$ and $\ell \in \mathbb{N}$
- There is a matrix $\mathbf{Q} \succeq 0$ of order $d+1$ satisfying $\mathcal{H}^*(\mathbf{Q}) = \mathbf{c}$, where $\deg p = 2d$ and $\mathbf{c} = (c_0, c_1, \ldots, c_{2d})^T$ are coefficients of p

The last condition is a **semidefinite constraint**.

When is a polynomial $p \in \mathbb{R}[x]$ nonnegative in $[-1, 1]$?

Proposition

The following are equivalent:

- $p(x) \geq 0$ for all $x \in [-1, 1]$
- There are sums of squares $z, w \in \mathbb{R}[x]$ satisfying

$$p(x) = z(x) + (1 - x^2)w(x), \qquad x \in [-1, 1]$$

- There are $\mathbf{Z} \in \mathcal{S}^{d+1}$, $\mathbf{W} \in \mathcal{S}^{d}$, where $\mathbf{Z} \succeq 0$ and $\mathbf{W} \succeq 0$, such that

$$\mathcal{H}^*(\mathbf{Z} + \mathbf{L}_1 \mathbf{W} \mathbf{L}_1^T - \mathbf{L}_2 \mathbf{W} \mathbf{L}_2^T) = \mathbf{c}$$

  for appropriately chosen matrices $\mathbf{L}_i$ of 0s and 1s

The last condition is a **semidefinite constraint**.

Semidefinite representation of the moment space in $\mathbb{R}^3$

The moment space $\mathcal{M}_3$ in $\mathbb{R}^3$ is the set of moments $\boldsymbol{\sigma} = (\sigma_0, \sigma_1, \sigma_2)$ with respect to all probability measures σ over $[-1, 1]$. We know that

$$\begin{aligned} \mathcal{M}_3 &= \text{conv}\{(1, x, x^2) \in \mathbb{R}^3 \mid x \in [-1, 1]\} \\ &= \{\boldsymbol{\sigma} \in \mathbb{R}^3 \mid \sigma_0 = 1,\ \ \sigma_2 - \sigma_1^2 \geq 0,\ \ 1 - \sigma_2 \geq 0\} \end{aligned}$$

Equivalent semidefinite constraints

$$\sigma_0 = 1 \qquad \mathcal{H}(\boldsymbol{\sigma}) = \begin{pmatrix} \sigma_0 & \sigma_1 \\ \sigma_1 & \sigma_2 \end{pmatrix} \succeq 0 \qquad 1 - \sigma_2 \geq 0$$

Semidefinite representation of the moment space in $\mathbb{R}^{m+1}$

The moment space $\mathcal{M}_{m+1}$ in $\mathbb{R}^{m+1}$ is the set of moments $\boldsymbol{\sigma} = (\sigma_0, \sigma_1, \ldots, \sigma_m)$ and we know that

$$\mathcal{M}_{m+1} = \boxed{\mathrm{conv}\{(1, x, x^2, \ldots, x^m) \in \mathbb{R}^{m+1} \mid x \in [-1, 1]\}}$$

Equivalent semidefinite constraints

$$\sigma_0 = 1 \qquad \mathcal{H}(\boldsymbol{\sigma}) \succeq 0 \qquad \mathbf{L}_1^T \mathcal{H}(\boldsymbol{\sigma}) \mathbf{L}_1 - \mathbf{L}_2^T \mathcal{H}(\boldsymbol{\sigma}) \mathbf{L}_2 \succeq 0$$

SDP formulation

- Utility function $u(x, y) = \sum_{i=0}^{k} \sum_{j=0}^{m} c_{ij} x^i y^j$, where $x, y \in [-1, 1]$
- $v \in \mathbb{R}$
- $\boldsymbol{\sigma} = (\sigma_0, \ldots, \sigma_m) \in \mathbb{R}^{m+1}$
- $\mathbf{C} \in \mathbb{R}^{(k+1)\times(m+1)}$ contains the coefficients of polynomial u

Semidefinite program for the 2nd player

Minimize v subject to

$$\begin{aligned} \mathbf{Z}, \mathbf{W} &\succeq 0 \\ \mathcal{H}^*(\mathbf{Z} + \mathbf{L}_1 \mathbf{W} \mathbf{L}_1^T - \mathbf{L}_2 \mathbf{W} \mathbf{L}_2^T) &= v\mathbf{e}_1 - \mathbf{C}\boldsymbol{\sigma} \\ \sigma_0 &= 1 \\ \mathcal{H}(\boldsymbol{\sigma}) &\succeq 0 \\ \mathbf{L}_1^T \mathcal{H}(\boldsymbol{\sigma}) \mathbf{L}_1 - \mathbf{L}_2^T \mathcal{H}(\boldsymbol{\sigma}) \mathbf{L}_2 &\succeq 0 \end{aligned}$$

Example

- Utility function $u(x, y) = 5xy - 2x^2 - 2xy^2 - y$, where $x, y \in [-1, 1]$
- $v \in \mathbb{R}$
- $\boldsymbol{\sigma} = (1, \sigma_1, \sigma_2) \in \mathbb{R}^3$
- $\mathbf{Z} \in \mathcal{S}^2$, $\mathbf{W} \in \mathcal{S}^1$, where $\mathbf{Z} \succeq 0$ and $\mathbf{W} \succeq 0$

How to arrive at an SD formulation

1. $\begin{pmatrix} 1 & \sigma_1 \\ \sigma_1 & \sigma_2 \end{pmatrix} \succeq 0 \qquad 1 - \sigma_2 \geq 0$
2. Decomposition of $v - U(x, \boldsymbol{\sigma})$ by using two s.o.s. polynomials:

$$v - 5\sigma_1 x + 2x^2 + 2\sigma_2 x + \sigma_1 = \begin{pmatrix} 1 & x \end{pmatrix} \mathbf{Z} \begin{pmatrix} 1 \\ x \end{pmatrix} + (1 - x^2)\mathbf{W}$$

3. Equating corresponding coefficients yields the above SD constraints

Example (ctnd.)

The solution of SDP for the 2nd player is

$$v = -0.48 \qquad \boldsymbol{\sigma} = (1, 0.56, 1) \qquad \mathbf{Z} = \begin{pmatrix} 0.08 & -0.4 \\ -0.4 & 2 \end{pmatrix} \qquad \mathbf{W} = 0$$

How to recover the optimal mixed strategies from their moments

1. **Dual SDP** provides the optimal moments $\boldsymbol{\mu} = (1, 0.2, 0.04)$.
2. The support of optimal mixed strategy for the 1st player is localized by checking the best response condition $U(x, \boldsymbol{\sigma}) = v$, which is a quadratic equation $\Rightarrow$ $x = 0.2$
3. The support of optimal mixed strategy for the 2nd player is localized by checking the best response condition $U(\boldsymbol{\mu}, y) = v$, which is a quadratic equation $\Rightarrow$ $y = \pm 1$
4. The only unknown probability p of -1 is a unique solution of the moment constraint $\sigma_1 = (-1) \cdot p + (1 - p) \Rightarrow$ $p = 0.22$

Summary

- The presented algorithm is a natural generalization of linear programming for matrix games
- Polynomial games over more general strategy spaces can be solved by a hierarchy of SDPs (Laraki and Lasserre, 2012)
- `GloptiPoly 3` package - moments, optimization and SDP
- It is an open problem to compute efficiently Nash equilibria of a general-sum two-person polynomial game

References

G. Blekherman, P. A. Parrilo, and R. R. Thomas, editors.
Semidefinite optimization and convex algebraic geometry.
SIAM, 2012.

R. Laraki and J. B. Lasserre.
Semidefinite programming for min–max problems and games.
Mathematical Programming, 131(1-2):305–332, 2012.

P. Parrilo.
Polynomial games and sum of squares optimization.
45th IEEE Conference on Decision and Control, 2855–2860, 2006.