Computation of ϵ-equilibria in Separable Games

Noah Stein

January 31, 2006

Games

- Set I of interacting agents ($I = \{1, 2\}$ throughout)
- Set C_i of strategies for each player $i \in I$.
- Utility function $u_i : C_1 \times C_2 \to \mathbb{R}$.
  - Each player wants as much utility as possible.
  - Utilities capture all strategic interactions.

Equilibrium

- A **Nash Equilibrium** is a choice of strategy for each player, so that if only one player deviates, he cannot expect to improve his utility.
- An **ϵ-equilibrium** is weaker – no player can improve his payoff by more than ϵ.

Rock, Paper, Scissors

(u_1, u_2)	Rock	Paper	Scissors
Rock	$(0, 0)$	$(-1, +1)$	$(+1, -1)$
Paper	$(+1, -1)$	$(0, 0)$	$(-1, +1)$
Scissors	$(-1, +1)$	$(+1, -1)$	$(0, 0)$

- No equilibrium!
- Enlarge the set of strategies (not with dynamite).
- Allow players to choose a **mixed strategy**, i.e. a probability distribution over C_i.
- Define utility on these larger strategy spaces as expected utility.

Zero-sum Finite Games

- Both C_1 and C_2 are finite.
- Utilities satisfy $u_1 = -u_2$.
- Strictly competitive games (Rock, Paper, Scissors).
- Set of mixed strategies for each player is a simplex.
- Prove existence of a Nash Equilibrium via LP duality, and compute it efficiently with interior point methods (von Neumann 1928).

General Finite Games

- Both C_1 and C_2 are finite.
- Utilities are arbitrary.
- Allows for both competition and cooperation.
- Prove existence of an equilibrium via a non-constructive fixed-point argument (Nash 1951).
- Simple algorithms exist, e.g. the Lemke-Howson algorithm (simplex method with a different pivoting rule).
- PPAD-completeness proven in a Dec. 4, 2005 paper.

Continuous Games

- Both C_1 and C_2 are compact metric spaces.
- Utilities are continuous.
- An equilibrium always exists (Glicksberg 1952).
- But the probability measures involved can be arbitrarily complicated!
- No hope of computing equilibria in general.

Zero-sum Polynomial Games

- $C_1 = C_2 = [-1, 1]$.
- Utilities satisfy $u_1 = -u_2 =$ a polynomial.
- Space of mixed strategies is infinite-dimensional, but has a finite-dimensional representation (more on next slide).
- Can be cast as an SDP, and computed efficiently with interior point methods (Parrilo 200x).

Separable games

- A continuous game is **separable** if it has payoffs:

$$u_i(s_1, s_2) = \sum_{k=1}^{r} a_i^k f_1^k(s_1) f_2^k(s_2)$$

  where $a_i^k \in \mathbb{R}$ and $f_j^k : C_j \to \mathbb{R}$ is continuous (superscripts are not exponents).

- The separable structure allows for a finite-dimensional representation of the mixed strategy space.

- Can assume WLOG that each player randomizes among at most $r + 1$ strategies.

Computing ϵ-equilibria for two-player separable games

- Assume $C_i = [-1, 1]$ and the utilities are Lipschitz.
- Discretize the game: Choose a set $\tilde{C}_i$ of $m \propto \frac{1}{\epsilon}$ equally spaced pure strategies for each player, and sample the utilities to get $\tilde{u}_i : \tilde{C}_1 \times \tilde{C}_2 \to \mathbb{R}$.
- Compute an equilibrium of this finite game.
- This yields an ϵ-equilibrium of the separable game.

Will this work?

- In general computing an equilibrium of a finite game is not easy.
- But in this case the finite game has the same separable structure as the original game:

$$\tilde{u}_i(s_1, s_2) = \sum_{k=1}^{r} a_i^k \tilde{f}_1^k(s_1) \tilde{f}_2^k(s_2)$$

- In particular the finite game has an equilibrium in which each player mixes among at most $r+1$ strategies, independent of the choice of $m \propto \frac{1}{\epsilon}$.

Computing an equilibrium of the finite game

- Given a guess at the support of each player's mixed strategy (which pure strategies he plays with positive probability) there is a simple LP to find equilibria with that support.
-

$$\#\ \text{supports} = \binom{m}{1} + \binom{m}{2} + \ldots + \binom{m}{r+1} \leq \underbrace{\binom{m+r}{1+r}}_{\text{polynomial in } m}$$

Complexity of the algorithm

- The number of LPs and the time to solve each are both polynomial in $\frac{1}{\epsilon}$.
- Algorithm is polynomial in $\frac{1}{\epsilon}$ and exponential in r.
- Dependence on r is no worse than for finite games.
- A recently published ϵ-equilibrium algorithm for finite games is exponential in $\frac{1}{\epsilon^2}$ (LMM 2003).

Conclusion

- Future directions:
  - Apply SDP-related methods to non-zero-sum polynomial games.
  - Consider separable games with additional structure, e.g. graphical separable games.
  - Algorithms for discontinuous games.
- Thanks to:
  - Asuman Ozdaglar
  - Pablo Parrilo
  - LIDS Student Conference Committee