

Game Theory 15-451 12/05/06
 - Zero-sum games
 - General-sum games

shall we play a game?

Game Theory and Computer Science

Plan for Today


- 2-Player Zero-Sum Games (matrix games)
 - Minimax optimal strategies
 - Minimax theorem test material
and proof not test material
- General-Sum Games (bimatrix games)
 - notion of Nash Equilibrium
- Proof of existence of Nash Equilibria
 - using Brouwer's fixed-point theorem

Consider the following scenario...

- Shooter has a penalty shot. Can choose to shoot left or shoot right.
- Goalie can choose to dive left or dive right.
- If goalie guesses correctly, (s)he saves the day. If not, it's a gooooooaaaaaall!
- Vice-versa for shooter.


2-Player Zero-Sum games

- Two players R and C. Zero-sum means that what's good for one is bad for the other.
- Game defined by matrix with a row for each of R's options and a column for each of C's options. Matrix tells who wins how much.
 - an entry (x,y) means: x = payoff to row player, y = payoff to column player. "Zero sum" means that $y = -x$.



Minimax-optimal strategies

- Minimax optimal strategy is a (randomized) strategy that has the best guarantee on its expected gain, over choices of the opponent. [maximizes the minimum]
- I.e., the thing to play if your opponent knows you well.



Minimax-optimal strategies

- Minimax optimal strategy is a (randomized) strategy that has the best guarantee on its expected gain, over choices of the opponent. [maximizes the minimum]
- I.e., the thing to play if your opponent knows you well.
- In class on Linear Programming, we saw how to solve for this using LP.
 - polynomial time in size of matrix if use poly-time LP alg.

Minimax-optimal strategies

- E.g., penalty shot:

	Left	Right
Left	(0,0)	(1,-1)
Right	(1,-1)	(0,0)

Minimax optimal strategy for both players is 50/50. Gives expected gain of $\frac{1}{2}$ for shooter ($-\frac{1}{2}$ for goalie). Any other is worse.

Minimax-optimal strategies

- E.g., penalty shot with goalie who's weaker on the left.

	Left	Right
Left	$(\frac{1}{2}, -\frac{1}{2})$	(1,-1)
Right	(1,-1)	(0,0)

Minimax optimal for shooter is (2/3,1/3). Guarantees expected gain at least 2/3. Minimax optimal for goalie is also (2/3,1/3). Guarantees expected loss at most 2/3.

Minimax Theorem (von Neumann 1928)

- Every 2-player zero-sum game has a unique value V .
- Minimax optimal strategy for R guarantees R's expected gain at least V .
- Minimax optimal strategy for C guarantees C's expected loss at most V .

Counterintuitive: Means it doesn't hurt to publish your strategy if both players are optimal. (Borel had proved for symmetric 5x5 but thought was false for larger games)

Matrix games and Algorithms

- Gives a useful way of thinking about guarantees on algorithms for a given problem.
- Think of rows as different algorithms, columns as different possible inputs. E.g., sorting
- $M(i,j)$ = cost of algorithm i on input j .

- Algorithm design goal: good strategy for row player. Lower bound: good strategy for adversary.

One way to think of upper-bounds/lower-bounds: on value of this game

Matrix games and Algorithms

- Gives a useful way of thinking about guarantees on algorithms for a given problem.
- Think of rows as different algorithms, columns as different possible inputs. E.g., sorting

- $M(i,j)$ = cost of algorithm i on input j .
- Algorithm design goal: good strategy for row player. Lower bound: good strategy for adversary.

Of course matrix may be HUGE. But helpful conceptually.

Matrix games and Algs


- What is a deterministic alg with a good worst-case guarantee? Adversary
- A row that does well against all columns.
- What is a lower bound for deterministic algorithms? Alg player
- Showing that for each row i there exists a column j such that $M(i,j)$ is bad.
- How to give lower bound for randomized algs?
- Give randomized strategy for adversary that is bad for all i . Must also be bad for all distributions over i .

E.g., hashing

- Rows are different hash functions.
- Cols are different sets of n items to hash.
- $M(i,j) = \# \text{collisions incurred by alg } i \text{ on set } j$.


We saw:

- For any row, can reverse-engineer a bad column.
- Universal hashing is a randomized strategy for row player that has good behavior for every column.
 - For any set of inputs, if you randomly construct hash function in this way, you won't get many collisions in expectation.



Nice proof of minimax thm (sketch)

- Suppose for contradiction it was false.
- This means some game G has $V_C > V_R$:
 - If Column player commits first, there exists a row that gets at least V_C .
 - But if Row player has to commit first, the Column player can make him get only V_R .
- Scale matrix so payoffs to row are in $[0,1]$. Say $V_R = V_C - \delta$.



Proof sketch, contd

- Consider exponential weighting alg from Nov16 lecture as Row, against opponent who always plays best response to Row's distrib.
- In T steps,
 - Alg gets $\geq (1-\epsilon/2)OPT - \log(n)/\epsilon$ [use $\epsilon=\delta$]
 - $OPT \geq T \cdot V_C$ [Best against opponent's empirical distribution]
 - Alg $\leq T \cdot V_R$ [Each time, opponent knows your randomized strategy]
 - Gap is δT . Contradicts assumption once $\delta T > (\epsilon/2)OPT + \log(n)/\epsilon$.

General-Sum Games

- Zero-sum games are good formalism for design/analysis of algorithms.
- General-sum games are good models for systems with many participants whose behavior affects each other's interests
 - E.g., routing on the internet
 - E.g., online auctions

General-sum games

- In general-sum games, can get win-win and lose-lose situations.
- E.g., "what side of road to drive on?":

		Left	Right	
	you	(1,1)	(-1,-1)	person driving towards you
		(-1,-1)	(1,1)	

General-sum games

- In general-sum games, can get win-win and lose-lose situations.
- E.g., "which movie should we go to?":

		Borat	Happy-feet
	Borat	(8,2)	(0,0)
	Happy-feet	(0,0)	(2,8)

No longer a unique "value" to the game.

Nash Equilibrium

- A Nash Equilibrium is a stable pair of strategies (could be randomized).
- Stable means that neither player has incentive to deviate on their own.
- E.g., "what side of road to drive on":

	Left	Right
Left	(1,1)	(-1,-1)
Right	(-1,-1)	(1,1)

NE are: both left, both right, or both 50/50.

Nash Equilibrium

- A Nash Equilibrium is a stable pair of strategies (could be randomized).
- Stable means that neither player has incentive to deviate.
- E.g., "which movie to go to":

	Borat	Happy-feet
Borat	(8,2)	(0,0)
Happy-feet	(0,0)	(2,8)

NE are: both B, both HF, or (80/20,20/80)

Uses


- Economists use games and equilibria as models of interaction.
- E.g., pollution / prisoner's dilemma:
 - (imagine pollution controls cost \$4 but improve everyone's environment by \$3)

	don't pollute	pollute
don't pollute	(2,2)	(-1,3)
pollute	(3,-1)	(0,0)

Need to add extra incentives to get good overall behavior.

NE can do strange things


- Braess paradox:
 - Road network, traffic going from s to t.
 - travel time as function of fraction x of traffic on a given edge.



Fine. NE is 50/50. Travel time = 1.5

NE can do strange things

- Braess paradox:
 - Road network, traffic going from s to t.
 - travel time as function of fraction x of traffic on a given edge.



Add new superhighway. NE: everyone uses zig-zag path. Travel time = 2.

Existence of NE

- Nash (1950) proved: any general-sum game must have at least one such equilibrium.
 - Might require randomized strategies (called "mixed strategies")
- This also yields minimax thm as a corollary.
 - Pick some NE and let V = value to row player in that equilibrium.
 - Since it's a NE, neither player can do better even knowing the (randomized) strategy their opponent is playing.
 - So, they're each playing minimax optimal.

Existence of NE

- Proof will be non-constructive.
- Unlike case of zero-sum games, we **do not know any** polynomial-time algorithm for finding Nash Equilibria in $n \times n$ general-sum games. [known to be "PPAD-hard"]
- Notation:
 - Assume an $n \times n$ matrix.
 - Use (p_1, \dots, p_n) to denote mixed strategy for row player, and (q_1, \dots, q_n) to denote mixed strategy for column player.

Proof

- We'll start with Brouwer's fixed point theorem.
 - Let S be a compact convex region in \mathbb{R}^n and let $f: S \rightarrow S$ be a continuous function.
 - Then there must exist $x \in S$ such that $f(x)=x$.
 - x is called a "fixed point" of f .
- Simple case: S is the interval $[0,1]$.
- We will care about:
 - $S = \{(p,q): p,q \text{ are legal probability distributions on } 1, \dots, n\}$. I.e., $S = \text{simplex}_n \times \text{simplex}_n$

Proof (cont)

- $S = \{(p,q): p,q \text{ are mixed strategies}\}$.
- Want to define $f(p,q) = (p',q')$ such that:
 - f is continuous. This means that changing p or q a little bit shouldn't cause p' or q' to change a lot.
 - Any fixed point of f is a Nash Equilibrium.
- Then Brouwer will imply existence of NE.

Try #1

- What about $f(p,q) = (p',q')$ where p' is best response to q , and q' is best response to p ?
- Problem: not necessarily well-defined:
 - E.g., penalty shot: if $p = (0.5, 0.5)$ then q' could be anything.

	Left	Right
Left	(0,0)	(1,-1)
Right	(1,-1)	(0,0)


Try #1

- What about $f(p,q) = (p',q')$ where p' is best response to q , and q' is best response to p ?
- Problem: also not continuous:
 - E.g., if $p = (0.51, 0.49)$ then $q' = (1,0)$. If $p = (0.49, 0.51)$ then $q' = (0,1)$.

	Left	Right
Left	(0,0)	(1,-1)
Right	(1,-1)	(0,0)

Instead we will use...


- $f(p,q) = (p',q')$ such that:
 - q' maximizes [(expected gain wrt p) - $\|q-q'\|^2$]
 - p' maximizes [(expected gain wrt q) - $\|p-p'\|^2$]



Note: quadratic + linear = quadratic.

Instead we will use...

- $f(p,q) = (p',q')$ such that:
 - q' maximizes [(expected gain wrt p) - $\|q-q'\|^2$]
 - p' maximizes [(expected gain wrt q) - $\|p-p'\|^2$]



Note: quadratic + linear = quadratic.

Instead we will use...

- $f(p,q) = (p',q')$ such that:
 - q' maximizes [(expected gain wrt p) - $\|q-q'\|^2$]
 - p' maximizes [(expected gain wrt q) - $\|p-p'\|^2$]
- f is well-defined and continuous since quadratic has unique maximum and small change to p,q only moves this a little.
- Also fixed point = NE. (even if tiny incentive to move, will move little bit).
- So, that's it!