CS 188: Artificial Intelligence
Spring 2006

Lecture 26: Game Theory
4/25/2006

Dan Klein – UC Berkeley

Game Theory

- Game theory: study of strategic situations, usually simultaneous actions
- A game has:
  - Players
  - Actions
  - Payoff matrix
- Example: prisoner's dilemma

Prisoner's Dilemma

		A	
		Testify	Refuse
B	Testify	-5,-5	-10,0
	Refuse	0,-10	-1,-1

Strategies

- Strategy = policy
- Pure strategy
  - Deterministic policy
  - In a one-move game, just a move
- Mixed strategy
  - Randomized policy
  - Ever good to use one?
- Strategy profile: a spec of one strategy per player
- Outcome: each strategy profile results in an (expected) number for each player

Prisoner's Dilemma

		A	
		Testify	Refuse
B	Testify	-5,-5	-10,0
	Refuse	0,-10	-1,-1

Two-Finger Morra

		O	
		One	Two
E	One	-2,2	3,-3
	Two	3,-3	-4,4

Dominance and Optimality

- Strategy Dominance:
  - A strategy s for A (strictly) dominates s' if it produces a better outcome for A, for any B strategy
- Outcome Dominance:
  - An outcome o Pareto dominates o' if all players prefer o to o'
  - An outcome is Pareto optimal if there is no outcome that all players would prefer

Prisoner's Dilemma

		A	
		Testify	Refuse
B	Testify	-5,-5	-10,0
	Refuse	0,-10	-1,-1

Two-Finger Morra

		O	
		One	Two
E	One	-2,2	3,-3
	Two	3,-3	-4,4

Equilibria

- In the prisoner's dilemma:
  - What will A do?
  - What will B do?
  - What's the dilemma?
- Both testifying is a (Nash) equilibrium
  - Neither player can benefit from a unilateral change in strategy
  - I.e., it's a local optimum (not necessarily global)
  - Nash showed that every game has such an equilibrium
  - Note: not every game has a dominant strategy equilibrium
- What do we have to change for the prisoners to refuse?
  - Change the payoffs
  - Consider repeated games
  - Limit the computational ability of the agents
  - How would we model a "code of thieves"?

		A	
		Testify	Refuse
B	Testify	-5,-5	-10,0
	Refuse	0,-10	-1,-1

Coordination Games

- No dominant strategy
  - But, two (pure) Nash equilibria
- What should agents do?
  - Can sometimes choose Pareto optimal Nash equilibrium
  - But may be ties!
  - Naturally gives rise to communication
  - Also: correlated equilibria

Technology Choice

		A	
		DVD	HD-DVD
B	DVD	5,5	-2,-1
	HD-DVD	-2,-1	8,8

Driving Direction

		A	
		Left	Right
B	Left	1,1	-1,-1
	Right	-1,-1	1,1

Mixed Strategy Games

- What's the Nash equilibrium?
  - No pure strategy equilibrium
  - Must look at mixed strategies
- Mixed strategies:
  - Distribution over actions per state
  - In a one-move game, a single distribution
  - For Morra, a single number p_{even} specifies the strategy
- How to choose the optimal mixed strategy?

Two-Finger Morra

		O	
		One	Two
E	One	-2,2	3,-3
	Two	3,-3	-4,4

(Zero-Sum) Minimax Strategies

- Idea: force one player to chose and declare a strategy first
  - Say E reveals first
  - For each E strategy, O has a minimax response
  - Utility of the root favors O (why?) and is -3 (from E's perspective)
  - If O goes first, root is 2 (for E)
  - If these two utilities matched, we would know the utility of the maximum equilibrium
- Must look at mixed strategies

Two-Finger Morra

		O	
		One	Two
E	One	-2,2	3,-3
	Two	3,-3	-4,4

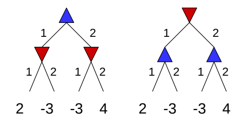

Continuous Minimax

- Imagine a minimax tree:
  - Instead of the two pure strategies, first player has infinitely many mixed ones
  - Note that second player should always respond with a pure strategy (why?)
- Here, can calculate the minimax (and maximin) values
  - Both are ½ (from O's perspective)
  - Correspond to [7/12; 1, 5/12; 2] for both players

Two-Finger Morra

		O	
		One	Two
E	One	-2,2	3,-3
	Two	3,-3	-4,4

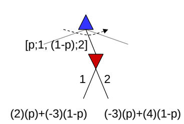

Repeated Games

- What about repeated games?
  - E.g. repeated prisoner's dilemma
  - Future responses, retaliation becomes an issue
  - Strategy can condition on past experience
- Repeated prisoner's dilemma
  - Fixed numbers of games causes repeated betrayal
  - If agents unsure of number of future games, other options
    - E.g. perpetual punishment: silent until you're betrayed, then testify thereafter
    - E.g. tit-for-tat: do what was done to you last round
  - It's enough for your opponent to believe you are incapable of remembering the number of games played (doesn't actually matter whether the limitation really exists)

Partially Observed Games

- Much harder to analyze
  - You have to work with trees of belief states
  - Problem: you don't know your opponent's belief state!
- Newer techniques can solve some partially observable games
  - Mini-poker analysis shows, e.g., that bluffing can be a rational action
  - Randomization: not just for being unpredictable, also useful for minimizing what opponent can learn from your actions

The Ultimatum Game

- Game theory can reveal interesting issues in social psychology
- E.g. the ultimatum game
  - Proposer: receives $x, offers split $k / $(x-k)
  - Accepter: either
    - Accepts: gets $k, proposer gets $(x-k)
    - Rejects: neither gets anything
- Nash equilibrium?
  - Any strategy profile where proposer offers $k and accepter will accept $k or greater
  - But that's not the interesting part…
- Issues:
  - Why do people tend to reject offers which are very unfair (e.g. $20 from $100)?
  - Irrationality?
  - Utility of $20 exceeded by utility of punishing the unfair proposer?
  - What about if x is very very large?

Mechanism Design

- One use of game theory: mechanism design
  - Designing a game which induces desired behavior in rational agents
- E.g. avoiding tragedies of the commons
  - Classic example: farmers share a common pasture
  - Each chooses how many goats to graze
  - Adding a goat gains utility for that farmer
  - Adding a goat slightly degrades the pasture
  - Inevitable that each farmer will keep adding goats until the commons is destroyed (tragedy!)
- Classic solution: charge for use of the commons
  - Prices need to be set to produce the right behavior

Auctions

- Example: auctions
  - Consider auction for one item
  - Each bidder i has value v_i and bids b_i for item
- English auction: increasing bids
  - How should bidder i bid?
  - What will the winner pay?
  - Why is this not an optimal result?
- Sealed single-bid auction, highest pays bid
  - How should bidder i bid?
  - Why is bidding your value no longer dominant?
  - Why is this auction not optimal?
- Sealed single-bid second-price auction
  - How should bidder i bid?
  - Bid v_i – why?