CS 416 Artificial Intelligence

Lecture 23

Making Complex Decisions

Chapter 17

Final Exam

- Reminder
  - Final Exam is Tuesday, May 6th at 7 p.m.
  - Let me know if you have a legitimate conflict

Zero-sum games

- Payoffs in each cell sum to zero
- Morra
  - Two players (Odd and Even)
  - Action
    - Each player simultaneously displays one or two fingers
  - Evaluation
    - f = total number of fingers
      - if f == odd, Even gives f dollars go to Odd
      - if f == even, Odd gives f dollars go to Even

Optimal strategy

- von Neumann (1928) developed optimal mixed strategy for two-player, zero-sum games
  - Because what one player wins, the other loses
    - just keep track of one player's payoff in each cell (Even)
    - assume this player wishes to maximize
  - Maximin technique
    - make game a turn-taking game and analyze

Maximin

- Change the rules of Morra for analysis
  - Force Even to reveal strategy first
    - apply minimax algorithm
    - Odd has an advantage and thus the outcome of the game is Even's worst case and Even *might* do better in real game
      - The utility of this game to Even is >= $-3

	O: one	*O: two*
E: one	$E = 2, O = -2$	$E = -3, O = 3$
E: two	$E = -3, O = 3$	$E = 4, O = -4$

Maximin

- Change the rules of Morra for analysis
  - Force Odd to reveal strategy first
    - Apply minimax algorithm
      - Odd would always select *one* to minimize Odd's loss
      - Even would always select *one* to maximize Even's gain
    - This game favors Even
      - The utility of this game to Even is <= +$2

	O: one	*O: two*
E: one	$E = 2, O = -2$	$E = -3, O = 3$
E: two	$E = -3, O = 3$	$E = 4, O = -4$

Combining two games

- Even's combined utility
  - EvenFirst_Utility <= Even's_Utility <= OddFirst_Utility
    - -3 <= Even's_Utility <= 2

Considering mixed strategies

- Mixed strategy
  - select one finger with prob: p
  - select two fingers with prob: 1 – p
- If one player reveals strategy first, second player will always use a pure strategy
  - expected utility of a mixed strategy
    - $U1 = p * u_{one} + (1-p) u_{two}$
  - expected utility of a pure strategy
    - $U2 = \max(u_{one}, u_{two})$
  - U2 is always greater than U1

Modeling as a game tree

- Because the second player will always use a fixed strategy…
  - Still pretending Even goes first

	O: one	O: two
E: one	$E = 2, O = -2$	$E = -3, O = 3$
E: two	$E = -3, O = 3$	$E = 4, O = -4$

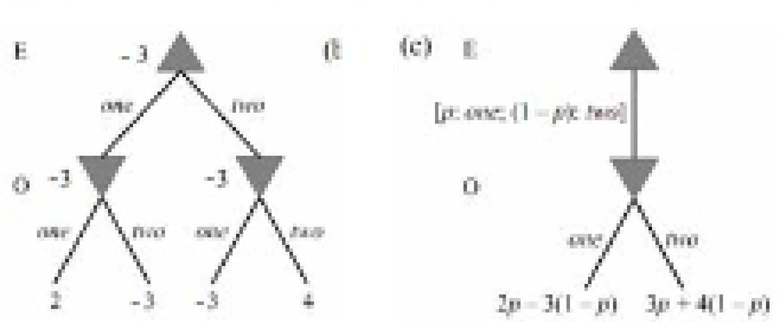

What is outcome of this game?

- Player Odd has a choice
  - Always pick the option that minimizes utility to Even
  - Represent two choices as functions of p
  - Odd picks line that is lowest (dark part on figure)
  - Even maximizes utility by choosing p to be where lines cross
    - 5p – 3 = 4 – 7p
      p = 7/12 => $E_{utility}$ = -1/12

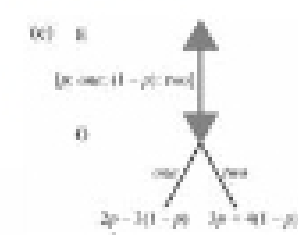

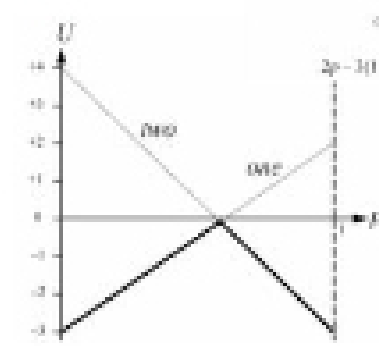

Pretend Odd must go first

- Even's outcome decided by pure strategy (dependent on q)
  - Even will always pick maximum of two choices
  - Odd will minimize the maximum of two choices
    - Odd chooses intersection point
    - 5q – 3 = 4 – 7q
      q = 7/12 => $E_{utility}$ = -1/12

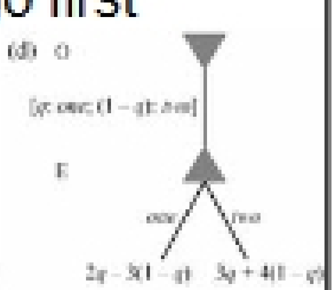

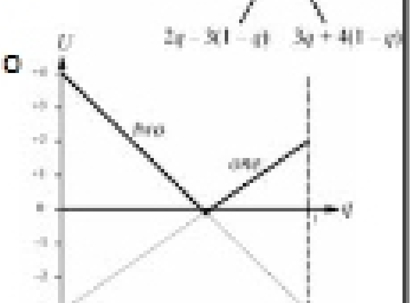

Final results

- Both players use same mixed strategy
  - p_{one} = 7/12
  - p_{two} = 5/12
  - Outcome of the game is -1/12 to Even

Generalization

- Two players with n *action choices*
  - mixed strategy is not as simple as p, 1-p
    - It is $(p_1, p_2, \ldots, p_{n-1}, 1-(p_1+p_2+\ldots+p_{n-1}))$
  - Solving for optimal p vector requires finding optimal point in (n-1)-dimensional space
    - lines become hyperplanes
    - some hyperplanes will be clearly worse for all p
    - find intersection among remaining hyperplanes
    - linear programming can solve this problem

Repeated games

- Imagine same game played multiple times
  - payoffs accumulate for each player
  - optimal strategy is a function of game history
    - must select optimal action for each possible game history
  - Strategies
    - perpetual punishment
      - cross me once and I'll take us both down forever
    - tit for tat
      - cross me once and I'll cross you the subsequent move

The design of games

- Let's invert the strategy selection process to design fair/effective games
  - Tragedy of the commons
    - individual farmers bring their livestock to the town commons to graze
    - commons is destroyed and all experience negative utility
    - all behaved rationally – refraining would not have saved the commons as someone else would eat it
      - Externalities are a way to place a value on changes in global utility
      - Power utilities pay for the utility they deprive neighboring communities (yet another Nobel prize in Econ for this – Coase)

Auctions

- English Auction
  - auctioneer incrementally raises bid price until one bidder remains
    - bidder gets the item at the highest price of another bidder plus the increment (perhaps the highest bidder would have spent more?)
    - strategy is simple… keep bidding until price is higher than utility
    - strategy of other bidders is irrelevant

Auctions

- Sealed bid auction
  - place your bid in an envelope and highest bid is selected
    - say your highest bid is v
    - say you believe the highest competing bid is b
    - bid min (v, b + ε)
    - player with highest value on good may not win the good and players must contemplate other player's values

Auctions

- Vickery Auction
  - Winner pays the price of the next highest bid
  - Dominant strategy is to bid what item is worth to you