



CS 15-888


Computational Game Solving



Lecture 1:

Game representations, solution concepts, and complexity

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University



The heart of the problem

- In 1-agent settings, the agent's expected-utility-maximizing strategy is well defined
- In multiagent settings, the outcome may depend on others' strategies also
 - => the agent's best strategy may depend on what strategies the other agent(s) choose, and vice versa
- **Game theory** is specifically developed to handle this
 - The only sound approach for handling this

Terminology

- *Agent = player*
- *Action = move* = choice that agent can make at a point in the game
- *Strategy* s_i = mapping from history (to the extent that the agent i can distinguish) to actions
 - (Can include “strategic”, “operational”, and “tactical” decisions)
- *Strategy set* S_i = strategies available to the agent
- *Strategy profile* $(s_1, s_2, \dots, s_{|A|})$, i.e., one strategy for each agent
- Agent's utility is determined after each agent (including *nature* that is used to model uncertainty) has chosen its strategy, and game has been played: $u_i = u_i(s_1, s_2, \dots, s_{|A|})$


Agenthood

- Agent attempts to *maximize its expected utility*
- Utility function u_i of agent i is a mapping from outcomes to reals
 - Incorporates agent's risk attitude (allows quantitative tradeoffs)
 - E.g. outcomes over money

Lottery 1: \$0.5M w.p. 1

Lottery 2: \$1M w.p. 0.5
\$0 w.p. 0.5

Agent's *strategy* is the
choice of lottery



Risk aversion => insurance companies

- Often in game theory we just talk about expected payoff or expected value (EV)

Utility functions are scale-invariant

- Agent i chooses a strategy that maximizes expected utility


$$\max_{\text{strategy}} \sum_{\text{outcome}} p(\text{outcome} \mid \text{strategy}) u_i(\text{outcome})$$

- If $u_i'() = a u_i() + b$ for $a > 0$, then the agent will choose the same strategy under utility function u_i' as it would under u_i
 - (u_i has to be finite for each possible outcome; otherwise expected utility could be infinite for several strategies, so the strategies could not be compared)
- Inter-agent utility comparison is problematic

Classical game representations

Extensive form
(aka **tree form**)

Matrix form
(aka **normal form**
aka **strategic form**)



Player 1's strategy

Player 2's strategy

	Left, Left	Left, Right	Right, Left	Right, Right
Up	1, 2	1, 2	3, 4	3, 4
Down	5, 6	7, 8	5, 6	7, 8

Combinatorial explosion



Dominant strategy “equilibrium”

- *Best response* s_i^* (for a *given* strategy profile s_{-i} of the other players):
for all s_i' , $u_i(s_i^*, s_{-i}) \geq u_i(s_i', s_{-i})$
- *Dominant strategy* s_i^* : s_i^* is a best response for *all* s_{-i}
 - Doesn't always exist
 - Inferior strategies are called “dominated”
- *Dominant strategy equilibrium* is a strategy profile where each agent has picked its dominant strategy
 - Doesn't always exist
 - Requires no counterspeculation
 - E.g., Prisoners' Dilemma:

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	-1,-1	-3,0
	Defect	0,-3	-2,-2

Nash equilibrium

[Nash 1950]



- In many games, an agent's best response depends on others' strategies, so the agent doesn't have a dominant strategy
- A strategy profile s^* is a *Nash equilibrium* if no player has incentive to deviate from his strategy given that others don't deviate:
for every agent i , $u_i(s_i^*, s_{-i}^*) \geq u_i(s_i', s_{-i}^*)$ for all s_i'
 - Dominant strategy equilibria are Nash equilibria but not vice versa
 - Defect-defect is the only Nash eq. in Prisoner's Dilemma
- E.g., *Battle of the Sexes* has no dominant strategies

		Woman	
		Boxing	Ballet
Man	Boxing	2, 1	0, 0
	Ballet	0, 0	1, 2

Criticisms of Nash equilibrium

- Not unique in all games, e.g., Battle of the Sexes
 - Approaches for addressing this problem
 - Refinements (= strengthenings) of the equilibrium concept
 - Eliminate weakly dominated strategies first
 - Choose the Nash equilibrium with highest welfare
 - Subgame perfection ...
 - Focal points
 - Mediation
 - Communication
 - Convention
 - Learning
- Doesn't exist in some games

1, 0	0, 1
0, 1	1, 0

Rock-paper-scissors game

Sequential moves

Rock-paper-scissors game


Simultaneous moves

Most real-world games are *imperfect-information* games

- Because a player doesn't know everything about what other players and nature have done
- Imperfect-information games are the main focus of this course
- Techniques for perfect-information games (such as AlphaZero) don't apply because there are additional issues:
 - Private information => Need to understand signals and how other players will interpret signals
 - Often involves deceiving and understanding deception by others
 - Game theory is the only sound way to handle bluffing, reveal/conceal, decoys, etc.


Example real-world imperfect-information games with sequential (& simultaneous) moves

- Negotiation
- Auctions, e.g.,:
 - Multi-stage auctions (e.g., FCC ascending, combinatorial auctions)
 - Sequential auctions of multiple items
- A robot facing adversaries in uncertain, stochastic envt
- Card games, e.g., poker
- Currency attacks
- International (over-)fishing
- Political campaigns (e.g., TV spending in each region)
- Ownership games (polar regions, moons, planets)
- Allocating and timing troops/armaments to locations
- Military spending games, e.g., space vs ocean
- Airport security, air marshals, coast guard, rail
- Cybersecurity
- ...



Imperfect-information games may require *mixed strategies*

A *mixed strategy* is an agent's chosen probability distribution over pure strategies



(Bayes-)Nash equilibrium:
Each agent uses a best-response strategy and has consistent beliefs

Rock-paper-scissors game has a (mixed-strategy) Nash equilibrium where each player plays each pure strategy with probability $1/3$

Fact: In mixed-strategy equilibrium, each strategy that occurs in the mix of agent i has equal expected utility to i

Behavioral strategy

- Agent has a probability distribution over her actions at each of her information sets
- **Kuhn's theorem:** If an agent has perfect recall, for every mixed strategy there is a behavioral strategy that has an equivalent payoff (i.e., the strategies are equivalent)
 - Applies also to infinite games

Existence of ***pure-strategy*** Nash equilibria

- **Theorem.**
 - Any finite game,
 - where **each action node is alone in its information set**
 - (i.e., at every point in the game, the agent whose turn it is to move knows what moves have been played so far)
 - is dominance solvable by backward induction (at least as long as ties are ruled out)
- **Constructive proof:** Multi-player minimax search
- Lots of interesting work has been done to tackle the computational complexity
 - Checkers is solved
 - Chess and Go are not solved but best AIs are superhuman

Existence & complexity of *mixed-strategy* Nash equilibria

- Every finite player, finite strategy game has at least one Nash equilibrium if we admit mixed-strategy equilibria as well as pure [Nash 50]
 - (Proof is based on Kakutani's fix point theorem)
- May be hard to compute
 - Complexity of finding a Nash equilibrium in a normal-form game:
 - 2-player 0-sum games can be solved in polytime with LP
 - 2-player games are
 - PPAD-complete (even with 0/1 payoffs) [Chen, Deng & Teng JACM-09; Abbott, Kane & Valiant FOCS-05; Daskalakis, Goldberg & Papadimitriou STOC-06],
 - NP-complete to find an even approximately *good* Nash equilibrium [Conitzer & Sandholm GEB-08]
 - 3-player games are FIXP-complete [Etessami & Yannakakis FOCS-07]

Properties of 2-player 0-sum games

- **Swappability:** if (x,y) and (x',y') are equilibria, then so are (x',y) and (x,y')
 => no equilibrium selection problem: player is safe playing any one of her equilibrium strategies
- A player's equilibrium strategies form a bounded convex polytope
 - Any convex combination of a player's equilibrium strategies is an equilibrium strategy
- The set of Nash equilibria are exactly the set of solutions to the *minmax problem*

$$\max_x \min_y u_1(x,y)$$
- **Minmax theorem** [von Neumann 1928]: Doesn't matter which player selects a mixed strategy first:

$$\max_x \min_y u_1(x,y) = \min_y \max_x u_1(x,y).$$
 This value is called the *game value*
- More generally:

Let $X \subset \mathbb{R}^n$ and $Y \subset \mathbb{R}^m$ be compact convex sets. If $f : X \times Y \rightarrow \mathbb{R}$ is a continuous function that is concave-convex, i.e.

$f(\cdot, y) : X \rightarrow \mathbb{R}$ is concave for fixed y , and

$f(x, \cdot) : Y \rightarrow \mathbb{R}$ is convex for fixed x .

Then we have that


$$\max_{x \in X} \min_{y \in Y} f(x, y) = \min_{y \in Y} \max_{x \in X} f(x, y).$$

Example

If $f(x, y) = x^T A y$ for a finite matrix $A \in \mathbb{R}^{n \times m}$, we have:

$$\max_{x \in X} \min_{y \in Y} x^T A y = \min_{y \in Y} \max_{x \in X} x^T A y.$$

The function $f(x,y)=y^2-x^2$ is concave-convex.



Properties of 2-player 0-sum games ...

- So, opponent can't make us get in expectation less than game value even if he knows our mixed strategy
- With ML, RL, generative AI, or any other approach that yields a non-equilibrium strategy, the strategy can't guarantee we get the game value. I.e., the strategy is *exploitable*
- Amazing **safety** property in multi-step imperfect-information games:
 - By playing a non-equilibrium strategy, our opponent can cause our beliefs to be wrong, but not in a way that the opponent's expected value increases!
 - So, it is safe to play under the assumption that the opponent plays rationally. If they don't, that will only help us.
- Solvable in polynomial time in the size of the game tree using LP
 - But what if the tree has 10^{165} nodes?