Game Theory

Thomas Lukasiewicz

16 December 2003

Overview

- Normal form games
- Nash equilibria
- Extensive form games

Literature

- S. J. Russel and P. Norvig. Artificial Intelligence: A Modern Approach, 2nd edition, Chapter 17. Prentice Hall, 2002. These slides are partially based on the slides available from http://aima.eecs.berkeley.edu/.

- D. Koller and A. Pfeffer. Representations and solutions for game-theoretic problems. Artificial Intelligence, 94: 167–215, 1997.

- Craig Boutilier. Introductory game theory. Available from http://www.cs.toronto.edu/∼cebly/.

Motivation

So far: single-agent decision making under uncertainty.

But what if uncertainty is due to other agents and the decisions they make?

And what if decisions of other agents are in turn influenced by our decisions?

...dependence of decisions is in some sense circular!

In general, a solution is a strategy profile (determines one action for each player) that is in some form of equilibrium.

Example: two-finger Morra

Two players O and E simultaneously display one or two fingers.
Let f denote the total number of displayed fingers.

If f is odd, then O collects f dollars from E.

If f is even, then E collects f dollars from O.

What should the players do? Best strategy against a rational player?

Which is the expected return for each player?

Normal form game

Agents (or players): $i \in \{1, \ldots, N\}$

Set of actions for each agent: $A_1, A_2, \ldots, A_N$

Utility function for each agent:

$u_1, u_2, \ldots, u_N \colon A_1 \times A_2 \times \cdots \times A_N \rightarrow \mathbf{R}$

$u_i(a_1, \ldots, a_N)$ is utility of joint action $(a_1, \ldots, a_N)$ to agent i

A pure strategy for agent i is an action from A_i

A pure strategy profile is an assignment of a pure strategy to each agent i

Outcome under $\mathbf{a} = (a_1, \ldots, a_N)$: $(u_1(\mathbf{a}), u_2(\mathbf{a}), \ldots, u_N(\mathbf{a}))$

Example: two-finger Morra

Two players O and E simultaneously display one or two fingers.
Let f denote the total number of displayed fingers.

If f is odd, then O collects f dollars from E.

If f is even, then E collects f dollars from O.

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

Example: prisoner's dilemma

Alice and Bob are caught near the scene of a burglary and are interrogated separately by the police.

If they both confess, they will each serve 5 years in prison for burglary.

If one confesses, she/he'll go free, while the partner will serve 10 years.

If they both don't confess, they will only serve 1 year each.

	Alice: testify	Alice: refuse
Bob: testify	$A = -5$; $B = -5$	$A = -10$; $B = 0$
Bob: refuse	$A = 0$; $B = -10$	$A = -1$; $B = -1$

Example: game hardware/software producers

Video game hardware producer "Acme" has to decide whether next game machine will use DVDs or CDs.

Video game software producer "Best" has to decide whether to produce its next game on DVD or CD.

	Acme: dvd	Acme: cd
Best: dvd	$A = 9$; $B = 9$	$A = -4$; $B = -1$
Best: cd	$A = -3$; $B = -1$	$A = 5$; $B = 5$

Example: matching pennies

There are two players A and B. They simultaneously reveal one penny each in their palms, either face up or face down.

Player A wins if both pennies are face up or both are face down.

Player B wins if one penny is face up, and the other one face down.

	A: heads	A: tails
B: heads	$A = 1; B = -1$	$A = -1; B = 1$
B: tails	$A = -1; B = 1$	$A = 1; B = -1$

Towards "rational" strategies

testify is a dominant strategy for Alice:

- If Bob testifies, she gets 5 years if she also does and 10 if she does not
- If Bob refuses, she gets 0 years if she testifies and 1 if she does not
- In both cases, it is better for Alice to testify
- Similarly, testify is also a dominant strategy for Bob

A strategy s for player i strongly dominates a strategy s' if the outcome under s is better for i than the outcome under s', for every choice of strategies by the other players. It weakly dominates s' if the outcome is better for at least one strategy profile and no worse for all the others.

A dominant strategy dominates all other strategies.

...never play a strongly dominated strategy
...always play a dominant strategy if one exists

(testify,testify) is a dominant strategy equilibrium

A strategy profile is a Nash equilibrium if no player can benefit by switching strategies, given the other players do not change strategy.

Every dominant strategy equilibrium is also a Nash equilibrium.

An outcome is Pareto optimal if there is no other outcome that all players would prefer. An outcome is Pareto dominated by another outcome iff all players prefer the other outcome.

(testify,testify) is a dominant strategy and Nash equilibrium.
But its outcome $A = -5; B = -5$ is Pareto dominated by the outcome of $A = -1; B = -1$ of (refuse,refuse).

Best responses and equilibria

Notation: $\mathbf{a}_{-i}$ refers to strategy profile $\mathbf{a}$ with action for agent i removed

Action $a_i \in A_i$ is a best response to $\mathbf{a}_{-i}$ iff

$$u_i(a_i \circ \mathbf{a}_{-i}) \geq u_i(b \circ \mathbf{a}_{-i}) \text{ for all } b \in A_i.$$

Let $BR_i(\mathbf{a}_{-i})$ denote the set of all best responses to $\mathbf{a}_{-i}$.

A pure strategy Nash equilibrium is any pure strategy profile $\mathbf{a} = (a_1, a_2, \ldots, a_N)$ such that $a_i \in BR_i(\mathbf{a}_{-i})$ for all agents i.

Video game software/hardware producers: two pure strategy Nash equilibria: (dvd, dvd) and (cd, cd):

...choose Pareto-optimal equilibrium, if exactly one exists

...communicate to choose among Pareto-optimal equilibria, if several exist (coordination game)

Two-finger Morra / matching pennies: no pure strategy Nash equilibrium

Mixed strategy

A mixed strategy σ_i for agent i is a probability distribution over A_i
(randomized choice, uncorrelated with choices of other agents)

A mixed strategy profile $\sigma = (\sigma_1, \ldots, \sigma_N)$:
mixed strategy σ_i for each agent i
pure strategy profile as special case

Utility: $u_i(\sigma) = E[u_i(\mathbf{a})|\sigma]$
randomization of each agent is independent
define $BR_i(\sigma_{-i})$ in standard way

A Nash equilibrium is any mixed strategy profile $\sigma = (\sigma_1, \ldots, \sigma_N)$
such that $\sigma_i \in BR_i(\sigma_{-i})$ for all agents i.

Examples

Two-finger Morra: (E: one [7/12], two [5/12]; O: one [7/12], two [5/12]) is a mixed strategy profile, which is a Nash equilibrium.

Matching pennies: (A: heads [0.5], tails [0.5]; B: heads [0.5], tails [0.5]) is a mixed strategy profile, which is the only Nash equilibrium:

If A plays “heads [0.5], tails [0.5]”, then B can play any mixed strategy, since they all have the same expected utility. Thus, there is no need for B to deviate from “heads [0.5], tails [0.5]”.

If A plays “heads [p], tails [$1-p$]”, $p>0.5$, then B should play “heads [0], tails [1]”, but then A should no longer play “heads [p], tails [$1-p$]”.

Properties of Nash equilibria

Some games have only mixed strategy Nash equilibria:
Two-finger Morra / matching pennies

Theorem (Nash): Any finite normal form game
has a (mixed strategy) Nash equilibrium.

Two-person zero-sum games

Two-person zero-sum games embody “pure competition”:

two players (1,2); $u_1(\mathbf{a}) = -u_2(\mathbf{a})$ for all joint actions $\mathbf{a}$

usually write $u(\mathbf{a})$ assuming 1 maximizes and 2 minimizes

more generally, “constant sum” games

matching pennies, chess, soccer, etc.

An agent shows the maximin behavior iff it selects a strategy whose guaranteed (= worst-case) payoff is best:

choosing best $\mathbf{a}$ assuming other does worst $\mathbf{a}$ for you

player 1 plays σ_1 where $\min_{\sigma_2} u(\sigma_1, \sigma_2)$ is maximal

player 2 plays σ_2 where $\max_{\sigma_1} u(\sigma_1, \sigma_2)$ is minimal

Theorem (von Neumann): In a two-person zero-sum game, a strategy profile is a Nash equilibrium iff both players show maximin behavior in it.

Hence, if (σ', σ'') is a Nash equilibrium, then

$$u(\sigma', \sigma'') = \max_{\sigma_1} \min_{\sigma_2} u(\sigma_1, \sigma_2) = \min_{\sigma_2} \max_{\sigma_1} u(\sigma_1, \sigma_2).$$

So, while there may be multiple equilibria for two-person zero-sum games, they all have the same value for each player.

Theorem (von Neumann): The normal form of a two-person zero-sum game defines a linear program whose solutions are the Nash equilibria (maximin strategies) of the game.

Matching pennies: $\max_{\sigma_1} \min_{\sigma_2} u(\sigma_1, \sigma_2) = 0.5$: If A does more than 0.5 heads, then B picks tails and A loses more than 0.5 the time.

Two-finger Morra: See following slide.

(c)

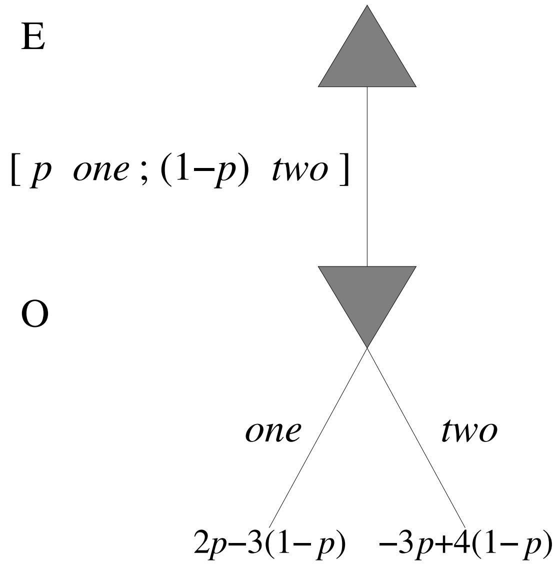

E
[p one ; (1−p) two]
O
one
two
2p−3(1−p)
−3p+4(1−p)

(d)

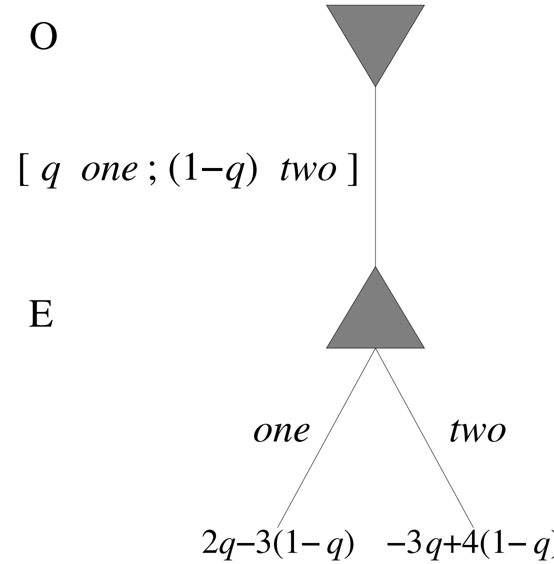

O
[q one ; (1−q) two]
E
one
two
2q−3(1−q)
−3q+4(1−q)

(e)

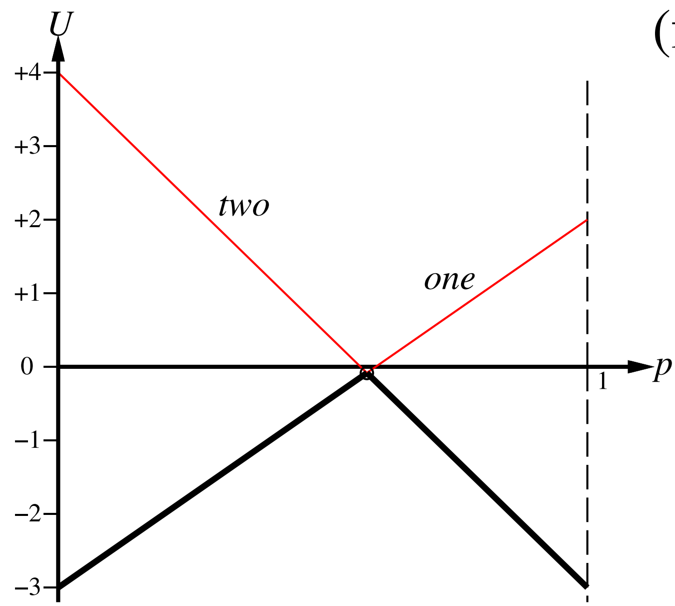

U
+4
+3
+2
+1
0
−1
−2
−3
two
one
1
p

(f)

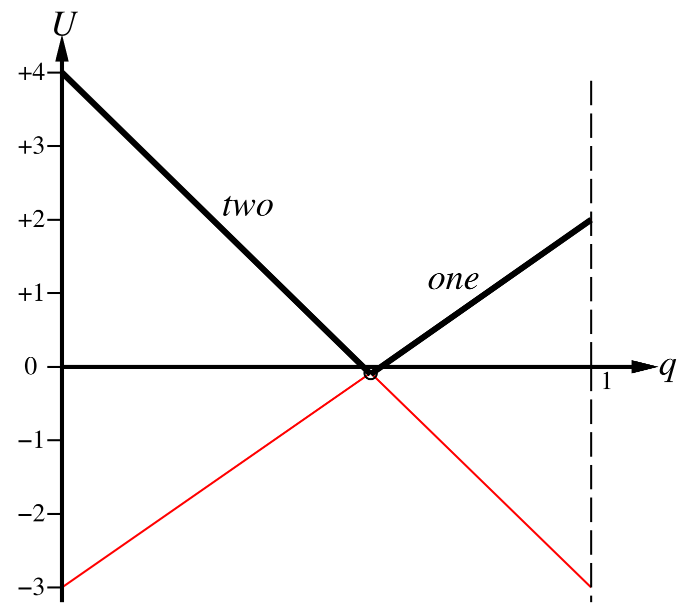

U
+4
+3
+2
+1
0
−1
−2
−3
two
one
1
q

Extensive form games

Multiple-move games allow more than a single move. For example, in repeated games, players face the same choice repeatedly, but each time with knowledge of the history of all players' previous choices.

Normal form games are very general, but they hide temporal structure.

Games with a truly sequential nature (e.g., chess, bargaining, etc.) are often more naturally represented using a game tree.

Example: sharing game

Player 1 offers a split of two goods,
and player 2 accepts or rejects.

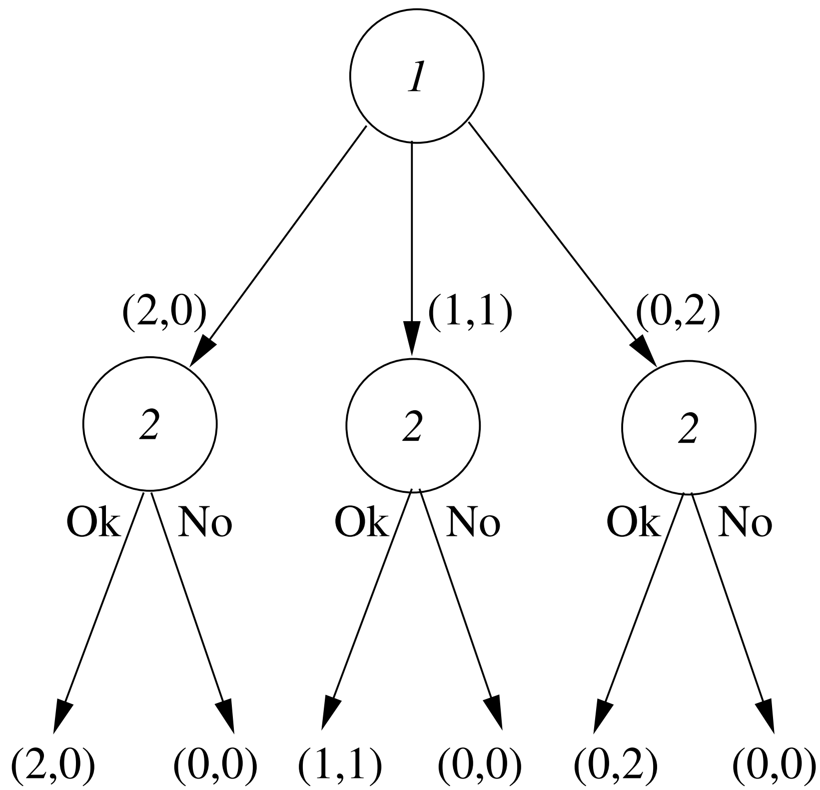

Extensive form game

Finite directed tree whose nodes denote game states.

Internal nodes are either decision nodes for some player i
or chance moves.

Outgoing edges of a decision node represent possible actions
at that node and have distinct labels called choices.

A play is a path from the root to some leaf.

A move is a choice taken on that path.

Utility function u_i for each player i:

u_i associates with each leaf p a real number

$u_i(p)$ is the utility at leaf p to player i

Set of decision nodes is partitioned into information sets. Each information set u belongs to exactly one player k. Intuitively, k cannot differentiate between the different nodes in u.

In the sequel, we assume perfect information:

every information set is a singleton

similar to full observability in MDPs

Strategies and Equilibria

A pure strategy for player i is any function that assigns an action $a_i \in A(h)$ to any node h owned by i.

A Nash equilibrium is any strategy profile σ such that
$u_i(\sigma_i \circ \sigma_{-i}) \geq u_i({\sigma_i}' \circ \sigma_{-i})$ for all players i and all ${\sigma_i}'$

Conversion to normal form

Consider each pure strategy in tree as an action

Equilibria are the same (by definition)

Example:

	(2,0)	(1,1)	(0,2)
000	(2,0)	(1,1)	(0,2)
00N	(2,0)	(1,1)	(0,0)
0N0	(2,0)	(0,0)	(0,2)
0NN	(2,0)	(0,0)	(0,0)
N00	(0,0)	(1,1)	(0,2)
N0N	(0,0)	(1,1)	(0,0)
NN0	(0,0)	(0,0)	(0,2)
NNN	(0,0)	(0,0)	(0,0)

Imperfect information

Matching pennies cannot be modeled as perfect information extensive form game, since simultaneity of moves cannot be expressed

Imperfect information games allow for expressing this:

information set for player i: a subset of nodes owned by player i
reflects histories that cannot be distinguished by i
set of moves at each node of an information set must be the same, and player's strategy is restricted to choosing the same move

Example: matching pennies

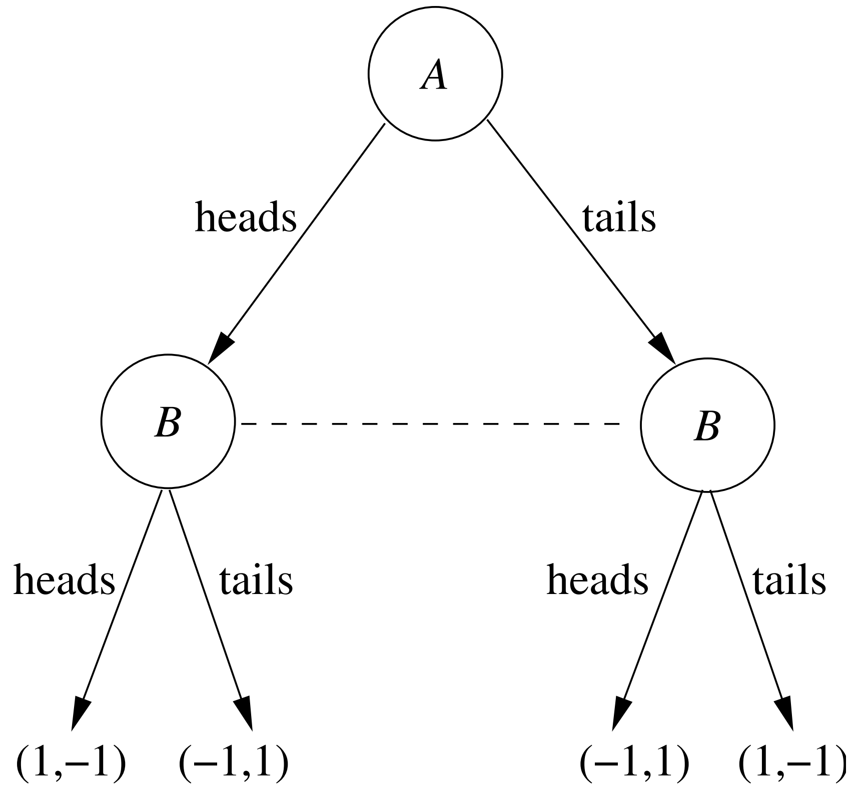

dashed line connects nodes in the same information set

Mechanism design / inverse game theory

So far: “Given a game, what is a rational strategy?”

Mechanism design / inverse game theory:

“Given that agents are rational, what game should we design?”
design a game such that for each agent pursuing a rational strategy means maximizing some global utility function

Example 1: Design protocols of internet traffic routers so that each router acts in a way such that global throughput is maximized.

Example 2: Intelligent multi-agent systems to solve complex problems in a distributed fashion without each agent knowing about the whole problem.