Decision Methods and Models

Master's Degree in Computer Science

Roberto Cordone

DI - Università degli Studi di Milano

Schedule:	Thursday 16.30 - 18.30 in Aula Magna (CS department)
	Friday 12.30 - 14.30 in classroom 301
Office hours:	on appointment
E-mail:	roberto.cordone@unimi.it
Web page:	https://homes.di.unimi.it/cordone/courses/2024-mmd/2024-mmd.html
Ariel site:	https://myariel.unimi.it/course/view.php?id=4467

Multiple decision-makers

We assume

- multiple decision-makers: $|D| > 1$
- preference relations Π_d that are weak orders,
  possibly with a known consistent value function $u^{(d)}(f)$
- a certain environment: $|\Omega| = 1 \Rightarrow f(x, \bar{\omega})$ reduces to $f(x)$

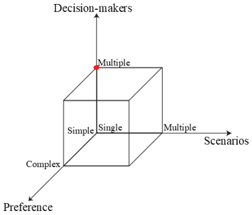

We consider only the two main cases

- game theory: the decision-makers make independent choices
- group decisions: the decision-makers make one coordinated choice

Game theory versus group decisions

In both cases each decision-maker $d \in D$ has a preference relation Π_d

In group decisions

- all relations Π_d are aggregated into a group preference Π
- the overall impact f depends on x
- Π determines the choice of a solution $x \in X$ or a ranking on X

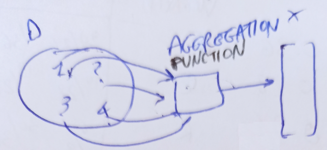

In game theory

- the overall impact f depends on all subvectors $x^{(d)}$ that form x
- Π_d determines only the choice of $x^{(d)} \in X^{(d)}$

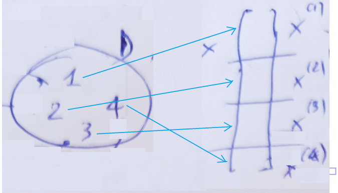

Given

- $X = X^{(1)} \times \ldots \times X^{(|D|)}$, that is $x = \left[\, x^{(1)} \; \ldots \; x^{(|D|)} \, \right]$
- $f = \left[\, f^{(1)} \; \ldots \; f^{(|D|)} \, \right]$ with $f^{(d)} \in \mathbb{R}$
- $\Pi^{(d)} = \left\{ \left(f^{(d)}, f'^{(d)} \right) : f^{(d)} \geq f'^{(d)} \right\}$,
  that is each $f^{(d)}$ is a benefit for decision-maker d

the decision problem can be reformulated as

$$\max f^{(d)}(x^{(1)}, \ldots, x^{(|D|)})$$
$$x^{(d)} \in X^{(d)}$$

It looks like a set of $|D|$ optimisation problems,
but each $f^{(d)}$ depends on all $x^{(d)}$!

We shall see that in general the concept of "optimal solution" is ill-posed

Terminology

Game theory has its own set of words for the usual basic concepts

- a decision problem is named a game
- a decision-maker is named a player
- an impact is named a payoff (it must be maximised)
- a solution x is named a strategy profile (*we shall see why*)
  - a (pure) strategy is the subvector $x^{(d)}$ associated with each player
  - $x^{(d)}$ consists of sub-subvectors associated with ordered time instants named moves

Classification

Games can be classified from different points of view

- with respect to the relation between players:
  - noncooperative: each player is independent (their choice is based only on the data)
  - cooperative: the players can agree to share payoffs
- with respect to the information on the data:
  - complete: all players know the whole of X and f
  - incomplete: each player d knows only $X^{(d)}$ and $f^{(d)}$
- with respect to the information on the moves:
  - perfect: all players knows all past moves
  - imperfect: player d knows his own past moves

The three classifications are independent

Consider noncooperative games with complete and perfect information

Game representations

All games can be represented in two ways

1. in extended form, that is as a tree
2. in strategic form, that is as a matrix

Usually, for each game one of them is clearly more natural than the other

Extended form

The extended form adopts the game tree:

- the nodes correspond to game states (the root to the starting one)
- typically, a turn of the game corresponds to $|D|$ consecutive levels
- all nodes on a level are associated with a player
- the levels are in chronological order
- the outgoing arcs represent the possible moves of the current player
- the leaves are associated with the payoffs (possibly in different levels)

Example: Tic-tac-toe

A 3x3 board is filled with crosses and noughts, 3 symbols in a line win

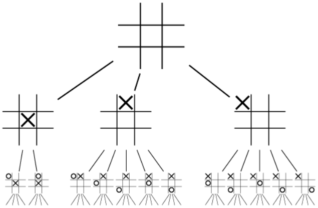

Reduce the combinatorial explosion exploiting symmetries!

Extended form

The game tree is very similar to the decision tree, but

- deterministic games have no level for the scenarios
- each stage has $|D|$ levels, instead of two

The game can be solved by backward induction:

- start from the leaves, where all payoffs are known
- in each internal node, consider the labels of the children nodes
- take the best choice for the current player (based on the level) and copy the payoff

If the payoff of the games is simply a win, loss or tie:

- if at least one children is a win, then choose it and win
- if no children wins and at least one is a tie, then choose it and tie
- if all children are a loss, then choose any and get a loss

Example: the Marienbad game

In its general form

- the game considers r rows with m_r matches on row r
- there are two players who move alternatively
- a move takes away from any single row any number of matches
- the player who takes the last match loses

(*the Nim game applies the opposite rule*)

We will determine a strategy for the specific case $r = m_1 = m_2 = 2$

(*a general strategy exists*)

To reduce the combinatorial explosion, at each move we

- sort the rows by nondecreasing number of matches

  ($(2, 1)$ *becomes* $(1, 2)$)
- merge all nodes with the same cardinalities and current player
  ($(1, 2)$ *for* P_1 *merges with* $(2, 1)$ *for* P_1, *but not with* $(1, 2)$ *for* P_2)
- put the end game nodes on a level on their own

This will turn the game tree into an acyclic directed graph

Example: the Marienbad game

The resulting analysis is

1. player P_1 moves from $A = (2,2)$ to $B = (1,2)$ or $C = (0,2)$
   ($nodes$ $(2,1)$ and $(2,0)$ $merge$ $with$ B and C)
2. player P_2 moves from $B = (1,2)$
   to $D = (0,2)$, $E = (1,1)$ or $F = (0,1)$, that is originally $(1,0)$
   (D *does not merge with* C *because they have different players*)
3. player P_2 moves from $C = (0,2)$ to $F = (0,1)$ or $G = (0,0)$
   (G *is and end game node*)
4. player P_1 moves from $D = (0,2)$ to $H = (0,1)$ or $I = (0,0)$
   (I *is and end game node with a payoff different from* G)
5. player P_1 moves from $E = (1,1)$ to $H = (0,1)$
6. player P_1 moves from $F = (0,1)$ to $I = (0,0)$
7. player P_2 moves from $H = (0,1)$ to $G = (0,0)$

There are only two end game nodes

1. G is a win for player P_1
2. I is a win for player P_2

Example: the Marienbad game

Backward induction is now straightforward

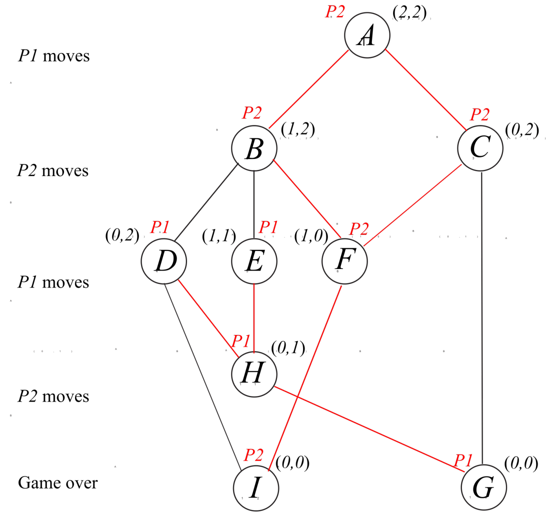

Example: the Marienbad game

The result is that

- H necessarily leads to G, and P_1 wins
- F necessarily leads to I, and P_2 wins
- E necessarily leads to H, and P_1 wins
- D can choose between H and I;
  since P_1 moves, the choice is on H, and P_1 wins
- C can choose between F and G;
  since P_2 moves, the choice is on F, and P_2 wins
- B can choose between D, E and F;
  since P_2 moves, the choice is on F, and P_2 wins
- A can choose between B and C, but both lead to a win for P_2

P_2 always wins, unless (playing irrationally) a way out is offered to P_1

A strategy indicates the move a player should make in each possible state

Therefore, it corresponds to a subset of arcs of the game tree that is

- consistent, including at most one arc for every node of the player
- complete, including at least one arc for every node of the player

In the Marienbad example, the optimal strategies for player P_1 are

Player P_1	A	D	E	F
$x^{(1)}$	(A, B)	(D, H)	(E, H)	(F, I)
$x'^{(1)}$	(A, C)	(D, H)	(E, H)	(F, I)

They are both losing, but allow to exploit possible mistakes of P_2

The optimal strategy for player P_2 is

Player P_2	B	C	H
x	(B, F)	(C, F)	(H, G)

Followed consistently, it guarantees to win

Example: rock, paper and scissors

A more troublesome example is given by the rock paper scissors game

- there are two players who move simultaneously

  The levels are not in strict chronological order

- they choose one option out of { rock, paper, scissors }
- if the players make identical choices the result is a tie
- if the players make different choices, one wins and the other loses:
  - rock breaks scissors
  - scissors cut paper
  - paper folds rock
- conventionally, the payoff is $+1$ for win, 0 for tie, -1 for loss

Let us build a game tree with an arbitrary play order

Example: rock, paper and scissors

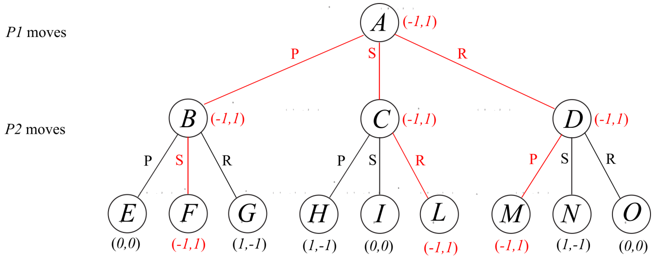

The backward induction method gives a winning strategy for P_2

- if P_1 plays paper, choose scissors
- if P_1 plays scissors, choose rock
- if P_1 plays rock, choose paper

This is of course impossible, because it requires unavailable information

Example: rock, paper and scissors

The player does not know the exact position in the tree!

We then partition the nodes of the game tree into information sets, i. e. minimal subsets of nodes of a player where the player knows to be, though the exact node is unknown

This is a more precise description of the state

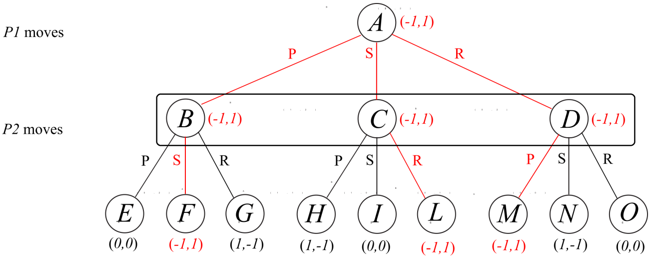

A strategy must be expressed in terms of information sets: information set $\{B, C, D\}$ must have a single move

With this limitation, an optimal solution can be impossible to find

Strategic representation

All games can be represented with a game matrix:

- the rows correspond to the strategies of the first player
- the columns correspond to the strategies of the second player
- lines in higher dimensions correspond to the strategies of the following players
- the matrix entries are associated with the payoffs

This clearly corresponds to an evaluation matrix

Example: Rock paper scissors

	R	P	S
R	(0,0)	(-1,1)	(1,-1)
P	(1,-1)	(0,0)	(-1,1)
S	(-1,1)	(1,-1)	(0,0)

Example: the Marienbad game

Player P_1 has 4 possible strategies (two choices in A and two in D)

1. (A, B), (D, H), (E, H), (F, I)
2. (A, B), (D, I), (E, H), (F, I)
3. (A, C), (D, H), (E, H), (F, I)
4. (A, C), (D, I), (E, H), (F, I)

We can further simplify the list by

- keeping the forced choices implicit   *(E, H) and (F, I) are irrelevant*
- merging the strategies that differ for choices in unreachable nodes
  (A, C) makes D unreachable: (D, H) and (D, I) can be merged

This yields 3 strategies

1. (A, B), (D, H)
2. (A, B), (D, I)
3. (A, C)

Example: the Marienbad game

Player P_2 has 6 possible strategies (three choices in B and two in C)

1. (B, D), (C, F), (H, G)
2. (B, D), (C, G), (H, G)
3. (B, E), (C, F), (H, G)
4. (B, E), (C, G), (H, G)
5. (B, F), (C, F), (H, G)
6. (B, F), (C, G), (H, G)

that we simplify as

1. (B, D), (C, F)
2. (B, D), (C, G)
3. (B, E), (C, F)
4. (B, E), (C, G)
5. (B, F), (C, F)
6. (B, F), (C, G)

Now we can build the matrix

Example: the Marienbad game

	$\{(B,D),(C,F)\}$	$\{(B,D),(C,G)\}$	$\{(B,E),(C,F)\}$	$\{(B,E),(C,G)\}$	$\{(B,F),(C,F)\}$	$\{(B,F),(C,G)\}$
$\{(A,B),(D,H)\}$	(1,-1)	(1,-1)	(1,-1)	(1,-1)	(-1,1)	(-1,1)
$\{(A,B),(D,I)\}$	(-1,1)	(-1,1)	(1,-1)	(1,-1)	(-1,1)	(-1,1)
$\{(A,C)\}$	(-1,1)	(1,-1)	(-1,1)	(1,-1)	(-1,1)	(1,-1)

The strategic form

- allows to compare alternative strategies at a glance
- potentially implies an exponential number of rows and columns

Player P_2 always wins with the optimal strategy $\{(B,E),(C,G)\}$

Dominance

For a given player $d \in D$, strategy $x^{(d)}$ dominates strategy $x'^{(d)}$ when

$$x^{(d)} \preceq x'^{(d)} \Leftrightarrow f(\ldots, x^{(d)}, \ldots) \geq f(\ldots, x'^{(d)}, \ldots)$$

for all $x^{(j)} \in X^{(j)}$ and all $j \in D \setminus d$

This is formally equivalent to

- Paretian preference, replacing the indicators l
- strong stochastic preference, replacing the scenarios ω

with the strategies of all other players

The effect is the same: strictly dominated strategies can be removed obtaining as a result

- an optimal solution (a rare case: e. g., Marienbad)
- an irreducible core (more frequent: e. g., rock, paper and scissors)

Example

	1	2	3
1	(4,5)	(5,0)	(5,2)
2	(2,6)	(9,1)	(3,2)
3	(3,2)	(2,8)	(6,0)

With complete information, both players know the dominated strategies

When removing a dominated strategy reveals other hidden dominances reduction cascades can occur

1. the column player removes column 3, that is dominated by 1
2. the row player removes row 3, that is now dominated by 1
   (a hidden dominance if the column payoffs are unknown)
3. the column player removes column 2
4. the row player removes row 2

In the end, only a single strategy profile (row 1 and column 1) remains: the corresponding strategies are optimal for the two players

Worst-case strategy

If a game requires simultaneous moves (backward induction is ruled out) and it does not reduce to a single nondominated strategy profile, can it be "solved"?

We can apply choice criteria, treating the other players as scenarios

The worst-case strategy assumes the worst payoff for each strategy

$$\max_{x^{(d)} \in X^{(d)}} \min_{x^{D \setminus (d)} \in X^{D \setminus (d)}} f(x)$$

Value of the game for player d is the minimum guaranteed payoff, that is the cost the player is willing to pay to join the game

In a two-player game

- $u^r = \max_{x^r \in X^r} \min_{x^c \in X^c} f^r(x^r, x^c)$
- $u^c = \max_{x^c \in X^c} \min_{x^r \in X^r} f^c(x^r, x^c)$

In general, the two values are totally unrelated

(*functions* $f^r(x)$ *and* $f^c(x)$ *are fully independent*)

Example

	1	2	3
1	(4,4)	(5,1)	(6,3)
2	(2,1)	(8,4)	(3,2)
3	(5,9)	(9,6)	(2,8)

u^r
4
2
2

u^c	1	1	2

Therefore

- the row player gains at least $u^r = 4$ with strategy 1
- the column player gains at least $u^c = 2$ with strategy 3

If both apply the worst-case criterium, the payoffs are $(6, 3)$

- the row player gains 6, that is actually $> u^r$
- the column player gains 3, that is actually $> u^c$
- better strategy profiles for both players exist (e. g., (9,6)), but require coordination and trust between the two players

Nash equilibrium

A concept recalling the idea of "solving" a game is the Nash equilibrium

This is a strategy profile $\left(x^{*(1)}, \ldots, x^{*|D|}\right) \in X$ such that

$$f^{(d)}(x^{*(1)}, \ldots, x^{*(d)}, \ldots, x^{*(|D|)}) \geq f^{(d)}(x^{*(1)}, \ldots, x^{(d)}, \ldots, x^{*(|D|)})$$

for all $d \in D, x^{(d)} \in X^{(d)}$

It looks similar to dominance, but it

- concerns a strategy profile x^*, instead of two strategies $x^{(d)}$ and $x'^{(d)}$
- compares a strategy $x^{*(d)}$ to all other strategies $x^{(d)}$ of the player
- fixes a strategy for all other players, instead of considering all possible ones

For two players, the contrast is easier:

- dominance compares corresponding entries in two rows or columns
- a Nash equilibrium compares a single entry
  with all other entries in the same row and in the same column

Nash equilibrium

The intuitive idea is that

- the game is repeated
- the last strategy profile is known to all players
- each player assumes that the other players will not change strategy

Is it profitable for the player to change strategy?

Finding the Nash equilibria

The best response method

- scans all players
- for each player d scans all strategies of the other players
- for each one, marks the best strategy of player d
- a strategy profile marked for all players is an equilibrium

In short, in two-player games we mark

- the best row payoff in each column
- the best column payoff in each row

Example: finding the Nash equilibria

	1	2	3
1	$(4, \bar{4})$	$(5, 1)$	$(\bar{6}, 3)$
2	$(2, 1)$	$(8, \bar{4})$	$(3, 2)$
3	$(\bar{5}, \bar{9})$	$(\bar{9}, 6)$	$(2, 8)$

- Considering the row player
  - the best response to column 1 is row 3: mark the entry $(3, 1)$
  - the best response to column 2 is row 3: mark the entry $(3, 2)$
  - the best response to column 3 is row 1: mark the entry $(1, 3)$
- Considering the column player
  - the best response to row 1 is column 1: mark the entry $(1, 1)$
  - the best response to row 2 is column 2: mark the entry $(2, 2)$
  - the best response to row 3 is column 1: mark the entry $(3, 1)$
- The entry $(3, 1)$ is marked for both players: it is a Nash equilibrium

Is this an "optimal" solution?
The row payoff could improve, but only if the column player cooperates, and the game is noncooperative

Example: finding the Nash equilibria

The rock paper scissor game has no Nash equilibrium

	R	P	S
R	$(0, 0)$	$(-1, \bar{1})$	$(\bar{1}, -1)$
P	$(\bar{1}, -1)$	$(0, 0)$	$(-1, \bar{1})$
S	$(-1, \bar{1})$	$(\bar{1}, -1)$	$(0, 0)$

Every player in every situation has an incentive to change choice

Example: finding the Nash equilibria

Finally, several Nash equilibria can coexist

	1	2	3
1	$(\bar{15}, \bar{11})$	$(4, 1)$	$(\bar{6}, 5)$
2	$(10, \bar{4})$	$(7, 2)$	$(3, 1)$
3	$(3, 6)$	$(\bar{12}, \bar{8})$	$(5, 7)$

If the last strategy profile was $(3, 2)$, moving to $(1, 1)$ would pay for both, but any single uncoordinated strategy change leads to worse payoffs