TECHNISCHE UNIVERSITÄT DRESDEN

Computational Logic ∴ Group

Hannes Strass

Faculty of Computer Science, Institute of Artificial Intelligence, Computational Logic Group

Games with Missing Information: Modelling

Lecture 6, 27th May 2024 // Algorithmic Game Theory, SS 2024

Previously . . .

- **Monte Carlo Tree Search** uses random playouts to evaluate moves and keeps statistics on which moves led to which payoffs how many times.
- A **selection policy** balances exploitation and exploration.
- **UCT** is an effective selection policy that applies UCB1 to trees.
- A **playout policy** steers playout simulations towards realistic play.
- MCTS and deep reinforcement learning led to expert-level Go programs.

Type A

Alpha-Beta Tree Search

Type B

Monte Carlo Tree Search

Overview

Example: The Monty Hall Problem

Extensive-Form Games

Bayes' Theorem

Preview: Simplified Poker

Motivation: Missing Information

- So far, we have considered games with perfect information:
- In every state, all players know the full history of play so far, i.e. they know their (joint) position in the game tree.
- However, e.g. in card games, players typically do not know the cards of opponents.
- This form of incomplete knowledge can be formalised by sets of indistinguishable nodes in the game tree, typically called information sets.
- In this context, we also add another element to games: chance.
- This is modelled via moves by nature and can be used to formalise dealing cards or throwing dice.
- We will see that this also allows us to model games with incomplete information, where e.g. some of the payoffs may be uncertain.
- In principle, however, chance and imperfect information are unrelated and we could model either without the other.

Course Evaluation: Lecture

15min to fill out this form:

https://befragung.zqa.tu-dresden.de/uz/de/sl/uGo82Q2GBW5i

Example: The Monty Hall Problem

Example: The Monty Hall Problem

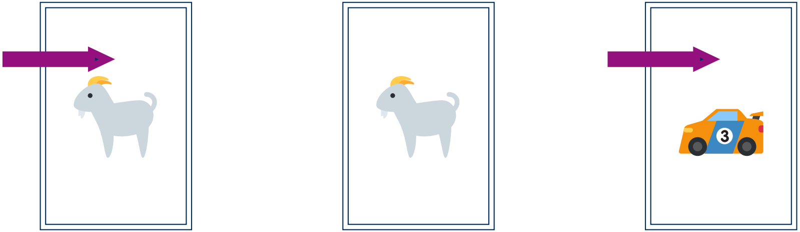

The Monty Hall Problem

A game show participant (`Guest`) is shown three doors behind which there are prizes. Behind one door, there is an expensive car, behind each of the other doors there is a goat. (The participant prefers the car over a goat.) The `Guest` is asked to `Choose` one of the doors. The game show `Host` (the other player) now opens one of the remaining doors that has a goat behind it. The `Guest` then gets their final move: `Stay` with the door they initially picked, or `Switch` to the other door. What should the participant do?

The Monty Hall Problem: Game Tree Sketch

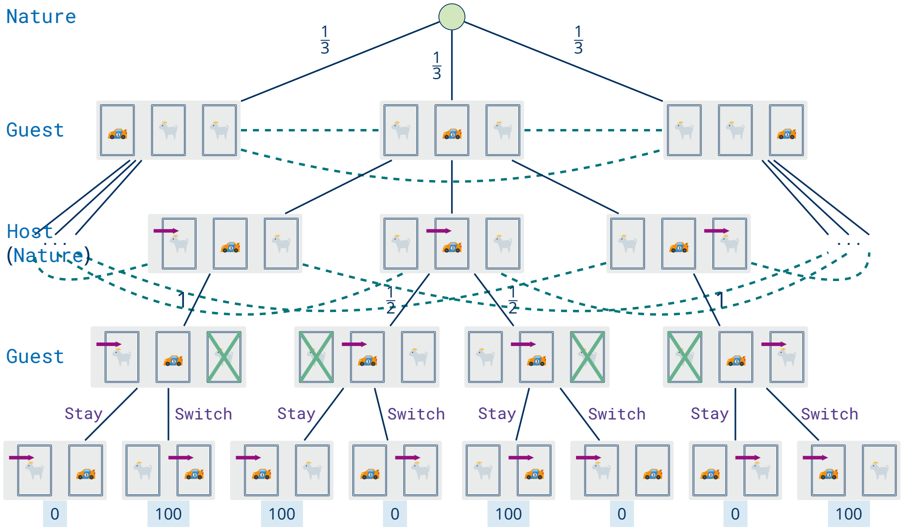

The Monty Hall Problem: Analysis

- Each of the possible states s_1, s_2, s_3 after $\texttt{Nature}$'s move has probability $\frac{1}{3}$.
- For each of these states, the ensuing game is symmetric.
- If $\texttt{Guest}$ chooses their door uniformly at random, then:
  - With probability $\frac{1}{3}$, their initial guess is correct; thus $\texttt{Switch}$ has a payoff of 0.
  - With probability $\frac{2}{3}$, their initial guess is wrong; thus $\texttt{Switch}$ has a payoff of 100.
- Thus in each s_i, $\texttt{Switch}$ has an expected payoff of $\frac{1}{3} \cdot 0 + \frac{2}{3} \cdot 100$.
- The overall payoff of $\texttt{Switch}$ is thus

$$u_{\texttt{Guest}}(\texttt{Switch}) = 3 \cdot \frac{1}{3} \cdot \left(\frac{1}{3} \cdot 0 + \frac{2}{3} \cdot 100 \right) = 66\frac{2}{3}$$

- Likewise, the overall payoff of $\texttt{Stay}$ is obtained as $u_{\texttt{Guest}}(\texttt{Stay}) = \frac{1}{3} \cdot 100 = 33\frac{1}{3}$.
- Therefore, a rational player should always choose $\texttt{Switch}$ over $\texttt{Stay}$.

Extensive-Form Games

Missing Information: Formalisation

Definition

An **extensive-form game** consists of the following:

1. A set $P = \{1, \ldots, n\}$ of at least two players, and possibly $\texttt{Nature}$.
2. An $n+1$-tuple $\mathbf{M} = (M_1, \ldots, M_n, M_{\texttt{Nature}})$ of sets M_i of **moves** for all players.
3. A set H of **histories**, sequences of moves $m_j \in M_1 \cup \ldots \cup M_n \cup M_{\texttt{Nature}}$.
4. A subset $Z \subseteq H$ of **terminal** histories.
5. A partition $\mathcal{I}_1 \dot\cup \ldots \dot\cup \mathcal{I}_k = H \setminus Z$ of non-terminal histories into **information sets** such that for all $1 \leq j \leq k$, all $h_1, h_2 \in \mathcal{I}_j$ have the same legal moves.
6. A **player function** $p \colon \{1, \ldots, k\} \to P \cup \{\texttt{Nature}\}$ (stating whose turn it is).
7. An n-tuple $\mathbf{u} = (u_1, \ldots, u_n)$ of utility functions $u_i \colon Z \to \mathbb{R}$.

Starting with the **empty history** $[]$, in each history $h = [m_1, \ldots, m_k] \in H \setminus Z$, player $i = p(h)$ chooses a move $m \in M_i$, leading to the history $[m_1, \ldots, m_k, m]$. Each $\mathcal{I}_j$ with $p(\mathcal{I}_j) = \texttt{Nature}$ has a probability distribution on possible moves.

Information Sets: Remarks

Intuition of information sets: The player (whose turn it is) does not have the information to distinguish between states in the set (but from other sets).

- $\mathcal{I} = \{\mathcal{I}_1, \ldots, \mathcal{I}_k\}$ being a **partition** means that:
  - for all $1 \leq j \leq k$, we have $\mathcal{I}_j \neq \emptyset$,
  - $\mathcal{I}_1 \cup \ldots \cup \mathcal{I}_k = H \setminus Z$, and
  - for all $1 \leq j, \ell \leq k$, we have $\mathcal{I}_j \cap \mathcal{I}_\ell = \emptyset$.
- Thus every $h \in H \setminus Z$ belongs to exactly one information set $\mathcal{I}_j \in \mathcal{I}$.
- For all $1 \leq j \leq k$ and $h \in \mathcal{I}_j$, we denote $p(h) := p(\mathcal{I}_j)$.
- $\mathcal{I}$ can also be represented by an **equivalence relation** $\sim^G$, where for any $h_1, h_2 \in H$, we have $h_1 \sim^G h_2$ iff there is a $\mathcal{I}_j \in \mathcal{I}$ such that $h_1, h_2 \in \mathcal{I}_j$.
- We graphically represent $\sim^G$ in game trees via dashed edges - - - - - .

Battleship

The initial placement of ships is private to the players and can be modelled via information sets. Some information may later be disclosed through hits.

Information Sets: Example

The Monty Hall Problem

- The (true) initial state is represented by the information set $\mathcal{I}_0 = \{[]\}$.
- The (seemingly) initial state for `Guest` is given by the information set

$$\mathcal{I}_1 = \{[\texttt{Car1}], [\texttt{Car2}], [\texttt{Car3}]\}.$$

- For each possible (initial) choice of door for `Guest`, there is one set:

$$\mathcal{I}_{\texttt{Choose1}} = \{[\texttt{Car1}, \texttt{Choose1}], [\texttt{Car2}, \texttt{Choose1}], [\texttt{Car3}, \texttt{Choose1}]\}$$
$$\mathcal{I}_{\texttt{Choose2}} = \{[\texttt{Car1}, \texttt{Choose2}], [\texttt{Car2}, \texttt{Choose2}], [\texttt{Car3}, \texttt{Choose2}]\}$$
$$\mathcal{I}_{\texttt{Choose3}} = \{[\texttt{Car1}, \texttt{Choose3}], [\texttt{Car2}, \texttt{Choose3}], [\texttt{Car3}, \texttt{Choose3}]\}$$

- Some information is disclosed by the host opening a door:

$$\mathcal{I}_{[\texttt{Choose1}, \texttt{Open2}]} = \{[\texttt{Car1}, \texttt{Choose1}], [\texttt{Car3}, \texttt{Choose1}]\}$$

Perfect Recall

Definition

Let $G = (P, \mathbf{M}, H, \mathcal{I}, p, \mathbf{u})$ be an extensive-form game.

- For every player $i \in P$ and history $h \in H$, define the sequence h_i of pairs $(\mathcal{I}_j, m)$ for $\mathcal{I}_j \in \mathcal{I}$ and $m \in M_i$ by induction:

$$[]_i := [] \qquad \text{and} \qquad [h; m]_i := \begin{cases} [h_i; (\mathcal{I}_h, m)] & \text{if } p(h) = i, \\ h_i & \text{otherwise.} \end{cases}$$

  where for $h = [m_1, \ldots, m_k]$ we denote $[h; m] := [m_1, \ldots, m_k, m]$.
- Player $i \in P$ has **perfect recall** in G iff for all $\mathcal{I}_j \in \mathcal{I}$, for every $h, h' \in \mathcal{I}_j$, it holds that $h_i = h'_i$.
- G has **perfect recall** iff every player $i \in P$ has perfect recall in G.

- h_i extracts all decision points and decisions (of player i) from history h.
- $h_i = h'_i$ means that i made the same moves *in the same information sets*.
- With perfect recall, players remember their trajectory through the game.

Perfect Recall: Examples (1)

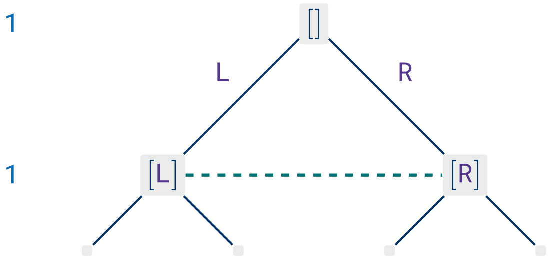

This game does not have perfect recall:

- Denote $\mathcal{I}_0 = \{[]\}$ and $\mathcal{I}_1 = \{[\mathsf{L}], [\mathsf{R}]\}$.
- We have $[\mathsf{L}] \in \mathcal{I}_1$ and $[\mathsf{R}] \in \mathcal{I}_1$, but:

$$[\mathsf{L}]_1 = [(\mathcal{I}_0, \mathsf{L})] \neq [(\mathcal{I}_0, \mathsf{R})] = [\mathsf{R}]_1$$

Perfect Recall: Examples (2)

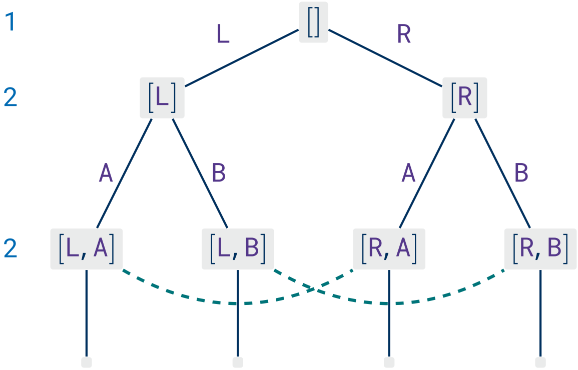

This game does not have perfect recall:

- Denote $\mathcal{I}_0 = \{[]\}$, $\mathcal{I}_1 = \{[L]\}$, $\mathcal{I}_2 = \{[R]\}$, $\mathcal{I}_3 = \{[L, A], [R, A]\}$, $\mathcal{I}_4 = \{[L, B], [R, B]\}$.
- Then $[L, A]_2 = [(\mathcal{I}_1, A)] \neq [(\mathcal{I}_2, A)] = [R, A]_2$.

Perfect Recall: Examples (3)

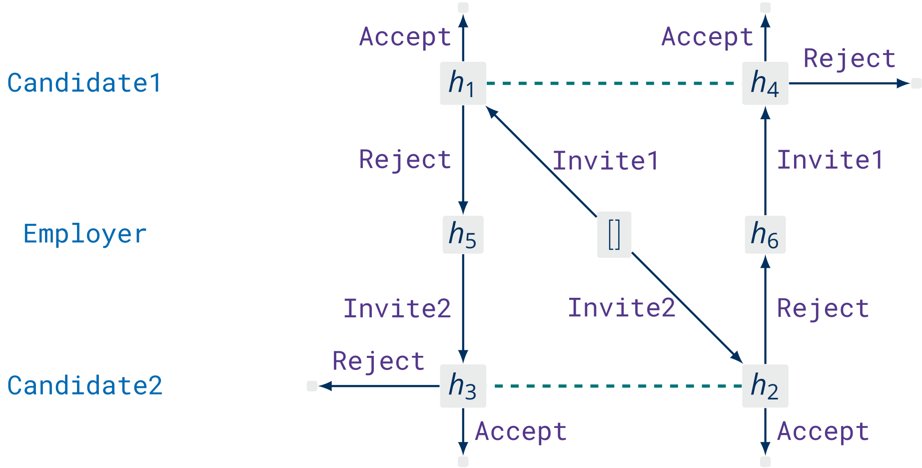

This game has perfect recall: e.g. $h \in \{h_2, h_3\}$ implies $h_{\texttt{Candidate2}} = []$.

Strategic Games and Imperfect Information

- Uncertainty induced by simultaneous moves can be modelled in extensive-form games (that seem to be sequential by definition).
- Main idea: Sequentialise moves, model uncertainty in information sets.

Example: Recall the game penalties. One extensive-form variant is:

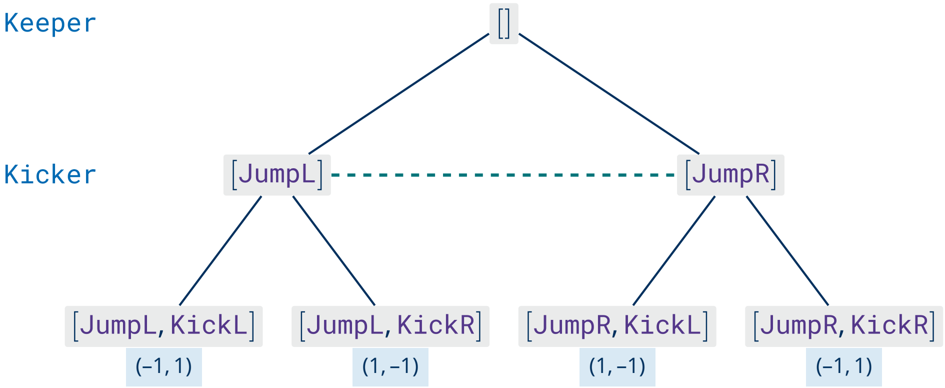

Chance Nodes (Moves by `Nature`)

Intuition of chance nodes: Something happens that is controlled by an entity with no strategic interest in the game's outcome.

Examples

- In card games, `Nature` controls the dealer's shuffling the cards.
- In games involving dice, `Nature` controls the dice throws.

- Probability distributions model uncertainty about effects of such actions.
- We typically use uniform distributions over possible atomic results.

$\rightsquigarrow$ We need some (more) probability theory to analyse games with chance . . .

Bayes' Theorem

Probabilities

Recall

- A **probability space** is a finite set $\mathcal{E} = \{e_1, \dots, e_k\}$ of **atomic events**.
- A **probability distribution** is a mapping $P \colon \mathcal{E} \to [0, 1]$, where atomic event e_i occurs with probability $P(e_i)$ and we have $\sum_{i=1}^{k} P(e_i) = 1$.
- An **event** $E \subseteq \mathcal{E}$ has (total) probability $P(E) = \sum_{e \in E} P(e)$.
- For all events $A, B \subseteq \mathcal{E}$ we have the following:

1. $0 \leq P(A) \leq 1$ with $P(\emptyset) = 0$ and $P(\mathcal{E}) = 1$.
2. $P(\overline{A}) = 1 - P(A)$ where $\overline{A} := \mathcal{E} \setminus A$ is the event **complementary** to A.
3. $P(A \cup B) = P(A) + P(B) - P(A \cap B)$.

Example

If all events $e_i \in \mathcal{E}$ have the same probability $\frac{1}{|\mathcal{E}|}$, we have a **uniform distribution**.

Conditional Probabilities

Definition

Let A and B be events with $P(B) > 0$.

1. The **conditional probability** for A to occur under the condition of B occurring is
$$P(A \mid B) := \frac{P(A \cap B)}{P(B)}$$
2. Events A and B are **independent** iff
$$P(A \cap B) = P(A) \cdot P(B)$$

That events A and B are independent is equivalently characterised by each of:

- $P(A \mid B) = P(A)$
- $P(A \mid B) = P(A \mid \overline{B})$
- $P(B \mid A) = P(B)$
- $P(B \mid A) = P(B \mid \overline{A})$

Bayes' Theorem

Theorem (Bayes)

1. If A and B are two events with $P(A) > 0$ and $P(B) > 0$, then
$$P(A) \cdot P(B \mid A) = P(B) \cdot P(A \mid B)$$
2. If A and $B_1, B_2, \ldots, B_\ell$ are events with $P(A) > 0$ and $P(B_i) > 0$ for all $1 \leq i \leq \ell$, where $\bigcup_{i=1}^{\ell} B_i = \mathcal{E}$ is a partition of $\mathcal{E}$, then for every $1 \leq i \leq \ell$:
$$P(B_i \mid A) = \frac{P(A \mid B_i) \cdot P(B_i)}{\sum_{j=1}^{\ell} \left(P(A \mid B_j) \cdot P(B_j)\right)} = \frac{P(A \mid B_i) \cdot P(B_i)}{P(A)}$$

In the second item of the theorem, the **law of total probability** is used:

$$P(A) = \sum_{j=1}^{\ell} P(B_j \cap A) = \sum_{j=1}^{\ell} \left(P(A \mid B_j) \cdot P(B_j)\right)$$

Note that $P(A) = P(\mathcal{E} \cap A) = P((\bigcup_{j=1}^{\ell} B_j) \cap A) = P(\bigcup_{j=1}^{\ell} (B_j \cap A)) = \sum_{j=1}^{\ell} P(B_j \cap A)$.

Solving the Monty Hall Problem (1)

Consider the following events:

A: The Guest wins the car.
B_1: The Guest initially chooses a goat door.
B_2: The Guest initially chooses the car door.

- If Guest chooses uniformly at random, then $P(B_1) = \frac{2}{3}$ and $P(B_2) = \frac{1}{3}$.
- Since every door has exactly one object, B_1 and B_2 are complementary, and the law of total probability yields $P(A) = P(A \mid B_1) \cdot P(B_1) + P(A \mid B_2) \cdot P(B_2)$.
- If Guest plays Stay, then clearly $P(A \mid B_1) = 0$ and $P(A \mid B_2) = 1$, whence

$$P(A) = P(A \mid B_1) \cdot P(B_1) + P(A \mid B_2) \cdot P(B_2) = 0 \cdot \frac{2}{3} + 1 \cdot \frac{1}{3} = \frac{1}{3}$$

- If Guest plays Switch, then $P(A \mid B_1) = 1$ and $P(A \mid B_2) = 0$, thus

$$P(A) = P(A \mid B_1) \cdot P(B_1) + P(A \mid B_2) \cdot P(B_2) = 1 \cdot \frac{2}{3} + 0 \cdot \frac{1}{3} = \frac{2}{3}$$

Solving the Monty Hall Problem (2)

Consider the following events:

A : The `Guest` wins the car.

B : The `Guest` initially chooses a goat door.

If the `Guest` plays `Switch`, then

- $P(B) = \frac{2}{3}$ as before,
- $P(A \mid B) = 1$ (initially choosing a goat door and switching win the car), and
- $P(B \mid A) = 1$ (initially choosing a goat door is the only way a `Switch` player can win the car).

According to Bayes' Theorem, we thus obtain

$$P(A) = \frac{P(A \mid B) \cdot P(B)}{P(B \mid A)} = \frac{2}{3}$$

Preview: Simplified Poker

Example: Simplified Poker

Binmore's Simplified Poker

- Two players, Ann and Bob, each put $1 into a jackpot.
- They then draw one card from a deck of three cards: $\{1, 2, 3\}$.
- Ann can either `check` (pass on), or `raise` (put another $1 into the jackpot).
- Next, Bob responds:
  - If Ann has `check`ed, then Bob must `call`, that is, a showdown happens: Both players show their cards and the player with the higher (number) card receives the jackpot.
  - If Ann has `raise`d, then Bob can decide between `fold` (withdraw from the game and let Ann get the jackpot) or `call` (put another $1 into the jackpot and then have a showdown).

Simplified Poker: Preliminary Analysis

Nature shuffles and deals the cards. There are six possible outcomes:

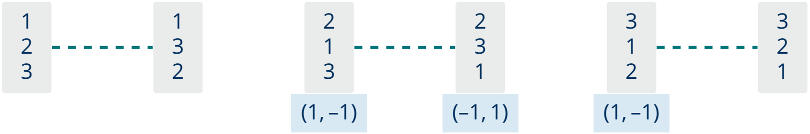

- If Ann draws a 3, she will raise; if Bob draws a 1, he will fold.
- If Bob draws a 3, he will call; if Ann draws a 2, she will check:
  Were she to raise, she would lose 2 if Bob has a 3 (as he would call), but still only win 1 if Bob has a 1 (as he would fold then).

What happens in the two remaining cases?

1. Should Ann raise (i.e. bluff) if she has a 1?
2. Should Bob call (the bluff) if he has a 2?

Conclusion

Summary

- In **complete information** games, players know the rules, possible outcomes and each other's preferences over outcomes.
- In **perfect information** games, moves are sequential and all players know all previous moves.
- In **extensive-form** games, information is not necessarily complete or perfect.
- Uncertainty of players (due to missing information) can be modelled by **information sets** and **chance nodes** (moves by `Nature`).
- **Bayes' Theorem** shows how to compute with conditional probabilities.
- The **law of total probability** relates marginal to conditional probabilities.