Algorithmic Game Theory and Applications

Lecture 10: Games in Extensive Form

Kousha Etessami

School of Informatics

the setting and motivation

- Most games we encounter in "real life" are not in "strategic form": players don't pick their entire strategies independently ("simultaneously").

  Instead, the game transpires over time, with players making "moves" to which other players react with their own "moves", etc.
  Examples: chess, poker, bargaining, dating, . . .

- A "game tree" looks something like this:

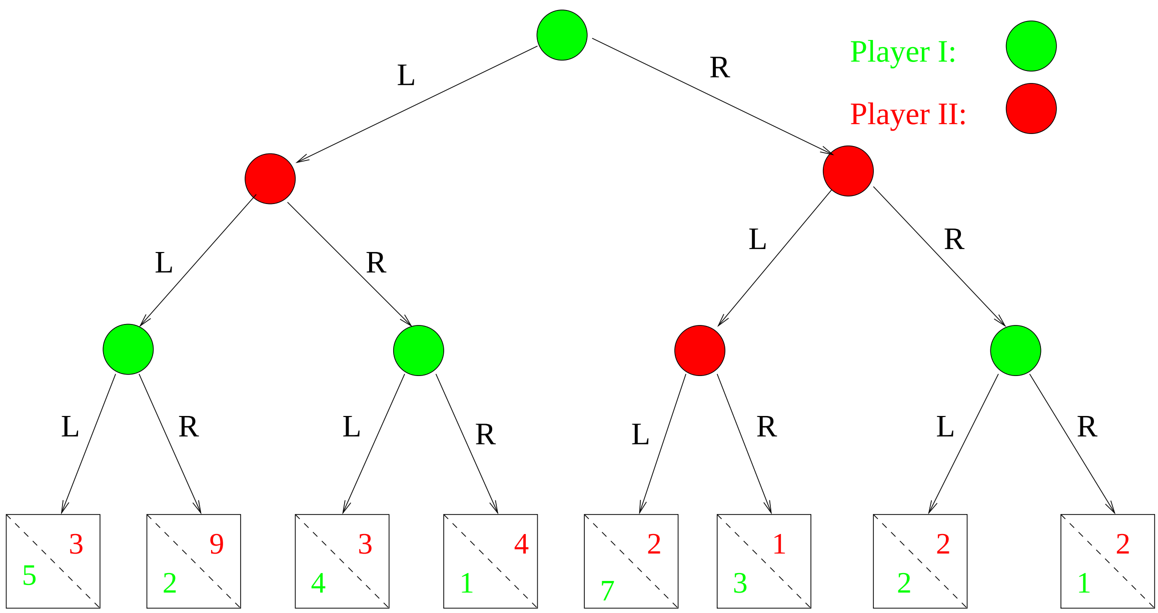

- But we may also need some other "features".

chance, information, etc.

Some tree nodes may be *chance* (probabilistic) nodes, controlled by no player (or, as is often said, controlled by "nature"). (Poker, Backgammon.)

Also, a player may not be able to distinguish between several of its "positions" or "nodes", because not all *information* is available to it. (Think Poker, with opponent's cards hidden.) Whatever move a player employs at a node must be employed at all nodes in the same "information set".

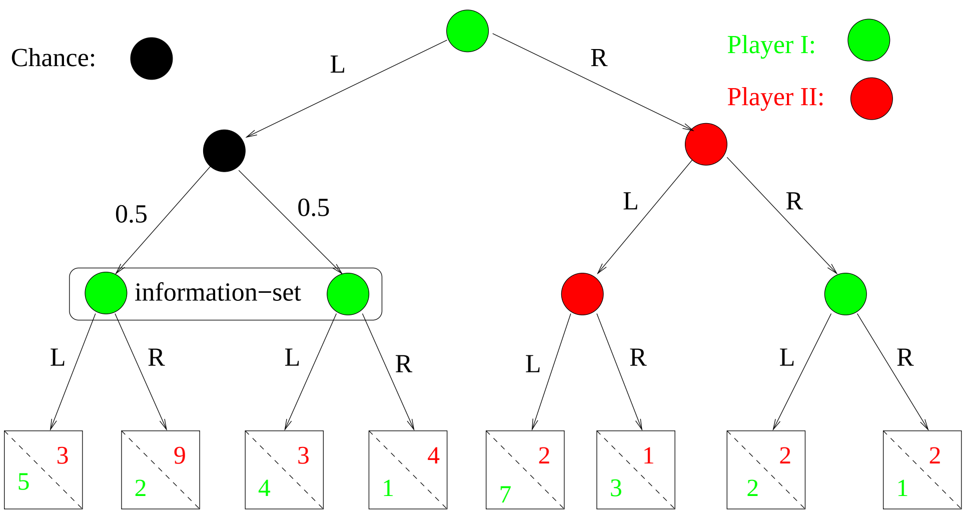

To define extensive form games, we have to formalize all these notions: game trees, whose turn it is to move, chance nodes, information sets, etc., etc., . . . So don't be annoyed at the abundance of notation......... it's all simple.

Trees: a formal definition

- Let $\Sigma = \{a_1, a_2, \ldots, a_k\}$ be an alphabet. A tree over Σ is a set $T \subseteq \Sigma^*$, of nodes $w \in \Sigma^*$ such that: if $w = w'a \in T$, then $w' \in T$.
- For a node $w \in T$, the children of w are $ch(w) = \{w' \in T \mid w' = wa \text{ , for some } a \in \Sigma\}$.
  For $w \in T$, let $Act(w) = \{a \in \Sigma \mid wa \in T\}$ be the "actions" available at w.
- A leaf (or terminal) node $w \in T$ is one where $ch(w) = \emptyset$. Let $L_T = \{w \in T \mid w \text{ a leaf}\}$.
- A (finite or infinite) path π in T is a sequence $\pi = w_0, w_1, w_2, \ldots$ of nodes $w_i \in T$, where if $w_{i+1} \in T$ then $w_{i+1} = w_i a$, for some $a \in \Sigma$. It is a complete path (or play) if $w_0 = \epsilon$ and every non-leaf node in π has a child in π.
  Let Ψ_T denote the set of plays of T.

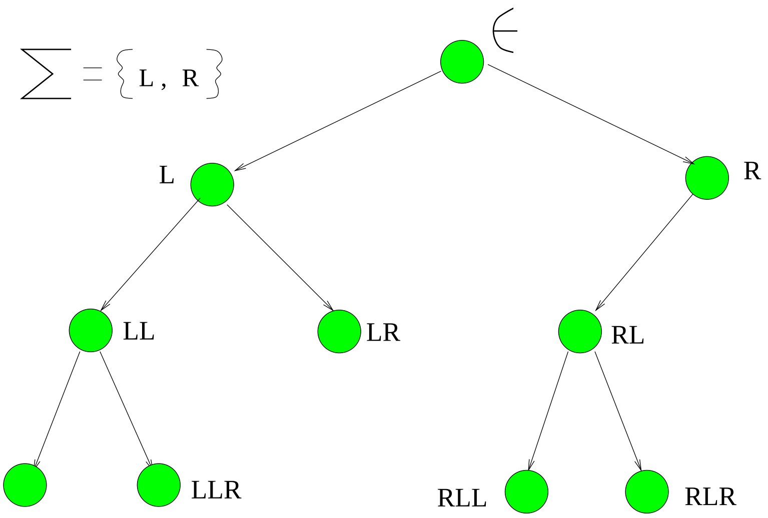

games in extensive form

A **Game in Extensive Form**, $\mathcal{G}$, consists of

1. A set $N = \{1, \ldots, n\}$ of players.
2. A tree T, called the game tree, over some Σ.
3. A map $pl : T \mapsto N \cup \{0\}$ from each $w \in T$ to the player $pl(w)$ whose "move" it is at w. (If $pl(w) = 0$ then it's "nature's move".) Let $Pl_i = pl^{-1}(i)$ be the nodes where it's player i's turn to move.
4. For each "nature" node, $w \in Pl_0$, a probability distribution $q_w : Act(w) \mapsto \mathbb{R}$ over its actions. (I.e., $q_w(a) \geq 0$, and $\sum_{a \in Act(w)} q_w(a) = 1$.)
5. For each player i, a map $info_i : Pl_i \mapsto \mathbb{N}$, which assigns to each $w \in Pl_i$ an index $info_i(w)$ for an information set. Let $Info_{i,j} = info_i^{-1}(j)$ be the set of nodes in the j'th information set for player i.

   Furthermore, for any i, j, & all nodes $w, w' \in Info_{i,j}$, $Act(w) = Act(w')$. (I.e., the set of possible "actions" from all nodes in the same information set is the same.)
6. For each player i, a function $u_i : \Psi_T \mapsto \mathbb{R}$, from (complete) plays to the payoff for player i.

explanation and comments

- Question: Why associate payoffs to "plays" rather than to leaves at the "end" of play?
  Answer: We in general allow infinite trees. We will later consider "infinite horizon" games in which play can go on for ever. Payoffs are then determined by the entire history of play.
  For "finite horizon" games, where tree T is finite, it suffices to associate payoffs to the leaves, i.e., $u_i : L_T \mapsto \mathbb{R}$.

- We defined our alphabet of possible actions Σ to be finite, which is generally sufficient for our purposes. In other words, the tree is finitely branching. In more general settings, even the set of possible actions at a given node can be infinite.

- In subsequent lectures, we will mainly focus on the following class of games:

  Definition An extensive form game $\mathcal{G}$ is called a game of **perfect information**, if every information set *Info*$_{i,j}$ has only 1 node.

pure strategies

- A pure strategy s_i for player i in an extensive game $\mathcal{G}$ is a function $s_i : \mathit{Pl}_i \mapsto \Sigma$ that assigns actions to each of player i's nodes, such that $s_i(w) \in Act(w)$, & such that if $w, w' \in \mathit{Info}_{i,j}$, then $s_i(w) = s_i(w')$.
  Let S_i be the set of pure strategies for player i.
- Given pure profile $s = (s_1, \ldots, s_n) \in S_1 \times \ldots \times S_n$,

  if there are no chance nodes (i.e., $\mathit{Pl}_0 = \emptyset$) then s uniquely determines a play π_s of the game: players move according their strategies:

  – Initialize $j := 0$, and $w_0 := \epsilon$;
  – While (w_j is not at a terminal node)
    If $w_j \in \mathit{Pl}_i$, then $w_{j+1} := w_j\ s_i(w_j)$;
    $j := j + 1$;
  – $\pi_s = w_0, w_1, \ldots$

- What if there are chance nodes?

pure strategies and chance

If there are chance nodes, then $s \in S$ determines a probability distribution over plays π of the game.

For finite extensive games, where T is finite, we can explicitly calculate the probability $p_s(\pi)$ of each play π using the probabilities $q_w(a)$:

Suppose $\pi = w_0, \ldots, w_m$, is a play of T.
Suppose further that for each $j < m$, if $w_j \in Pl_i$, then $w_{j+1} = w_j\ s_i(w_j)$. Otherwise, let $p_s(\pi) = 0$.

Let $w_{j_1}, \ldots, w_{j_r}$ be the chance nodes in π, and suppose, for each $k = 1, \ldots, r$, $w_{j_k+1} = w_{j_k} a_{j_k}$, i.e., the required action to get from node w_{j_k} to node w_{j_k+1} is a_{j_k}. Then

$$p_s(\pi) := \prod_{k=1}^{r} q_{w_{j_k}}(a_{j_k})$$

For infinite extensive games, defining these distributions in general requires much more elaborate definitions of the probability spaces, distributions, and densities (proper "measure theoretic" probability). (To even be able to define a distribution we would at least need a "finitistic" description of T's structure!)

We will avoid the heavy stuff as much as possible.

chance and expected payoffs

For a finite extensive game, given pure profile $s = (s_1, \ldots, s_n) \in S_1 \times \ldots \times S_n$, we can, define the "expected payoff" for player i under s as:

$$h_i(s) := \sum_{\pi \in \Psi_t} p_s(\pi) * u_i(\pi)$$

Again, for infinite games, much more elaborate definitions of "expected payoffs" would be required.

Note: This "expected payoff" does not arise because any player is mixing its strategies. It arises because the game itself contains randomness.

We can also combine both: players may also randomize amongst their strategies, and we could then define the overall expected payoff.

from strategic games to extensive games

Every finite strategic game Γ can be encoded easily and concisely as an extensive game $\mathcal{G}_\Gamma$. We illustrate this via the Rock-Paper-Scissor 2-player game, and leave the general n-player case as "homework".
(To encode infinite strategic games as extensive games, we would need an infinite action alphabet.)

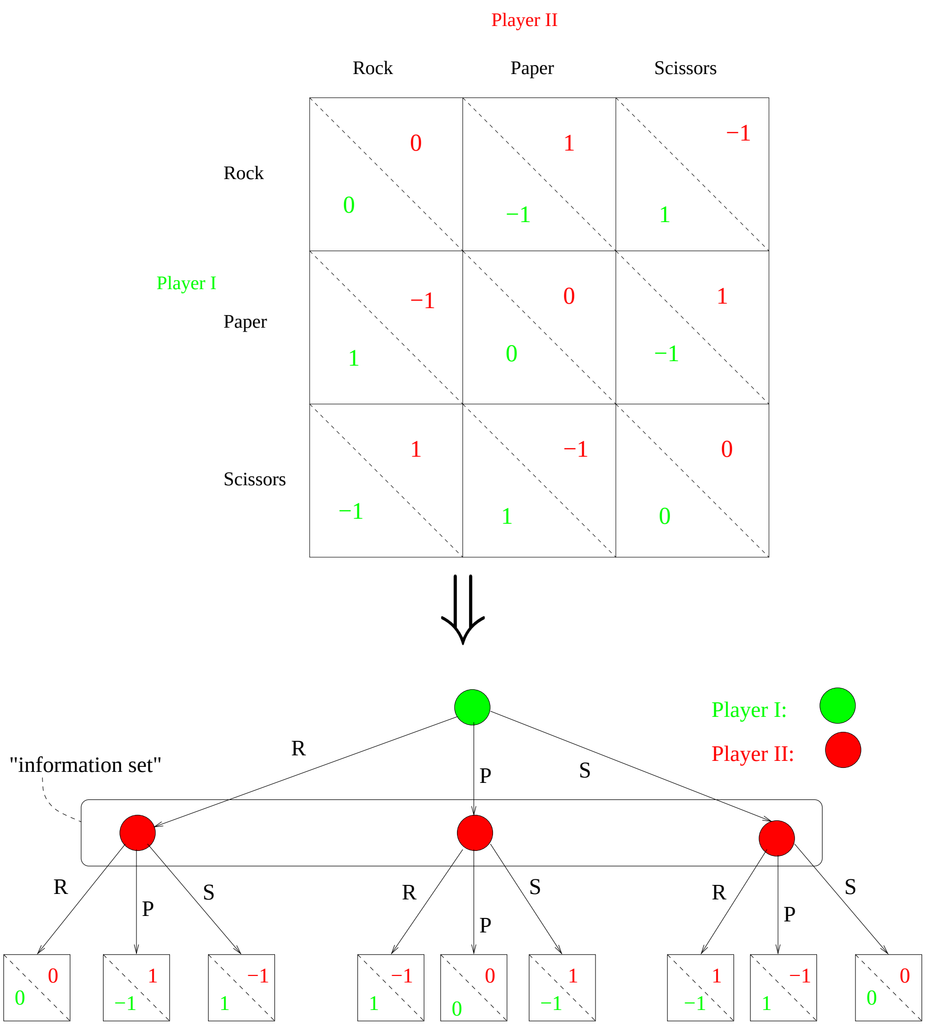

from extensive games to strategic games

Every extensive game $\mathcal{G}$ can be viewed as a strategic game $\Gamma_{\mathcal{G}}$:

- In $\Gamma_{\mathcal{G}}$, the strategies for player i are the mappings $s_i \in S_i$.
- In $\Gamma_{\mathcal{G}}$, we define payoff $u_i(s) := h_i(s)$, for every pure profile s.

(For an infinite game, we would need the expectations $h_i(s)$ to somehow be defined!)

If the extensive game $\mathcal{G}$ is finite, i.e., tree T is finite, then the strategic game $\Gamma_{\mathcal{G}}$ is also finite.
Thus, all the theory we developed for finite strategic games also applies to finite extensive games.

Unfortunately, the strategic game $\Gamma_{\mathcal{G}}$ is generally exponentially bigger than $\mathcal{G}$. Note that the number of pure strategies for a player i with $|Pl_i| = m$ nodes in the tree, is in the worst case $|\Sigma|^m$.

So it is often unwise to naively translate a game from extensive to strategic form in order to "solve" it.
If we can find a way to avoid this blow-up, we should.

solving games of imperfect info.

For finite extensive games of "imperfect information" (imp-info) there are some ways to mitigate the blow-up, but things are generally more difficult. We only briefly mention algorithms for imp-inf games. (See, e.g., [Koller-Megiddo-von Stengel'94].)

- In strategic 2-player zero-sum games we can find minimax solutions efficiently (P-time) via LP. For 2-player zero-sum extensive imp-info games, finding a minimax solution is "NP-hard". NE's of 2-player extensive games can be found, by cleverer exhaustive search, in exponential time.
- The situation is better for particular classes of games, e.g., games of "perfect recall". Intuitively, an imp-info game is of "perfect recall" if each player i never "forgets" its own actions: if it made different choices in the history leading to two of its nodes w and w', then $\mathit{info}_i(w) \neq \mathit{info}_i(w')$. 2-player zero-sum imp-info games of perfect recall can be solved in P-time, via LP, and 2-player NE's for arbitrary perfect recall games can be found in exponential time using an LH-style algorithm.

Our main priority will be games of perfect information. There the situation is much easier.

games of perfect information

Recall, a game of perfect information has only 1 node per information set. So, for these we can forget about information sets.

Examples: Chess, Backgammon, ...

counter-Examples: Poker, Bridge, ...

Theorem([Zermelo'1912,Kuhn'53]) Every finite extensive game of perfect information, $\mathcal{G}$, has a NE is pure strategies.

In other words, there is a pure profile $(s_1, \ldots, s_n) \in S$ that is a Nash Equilibrium.

Our proof will actually provide an easy algorithm to efficiently compute such a pure profile given $\mathcal{G}$, using "backward induction".

A special case of this theorem says the following:

Proposition([Zermelo'1912]) In Chess, either

1. White has a "winning strategy", or
2. Black has a "winning strategy", or
3. Both players have strategies to force a draw.

Next time, we continue with perfect information games.