




# More on Games

## Chapter 6

# Overview

- Stochastic games
- Other issues
- General game playing


# Stochastic Games



- In real life, unpredictable external events can put us into unforeseen situations
- Many games introduce unpredictability through a random element, such as the throwing of dice
- These offer simple scenarios for problem solving with adversaries and uncertainty

# Example: Backgammon

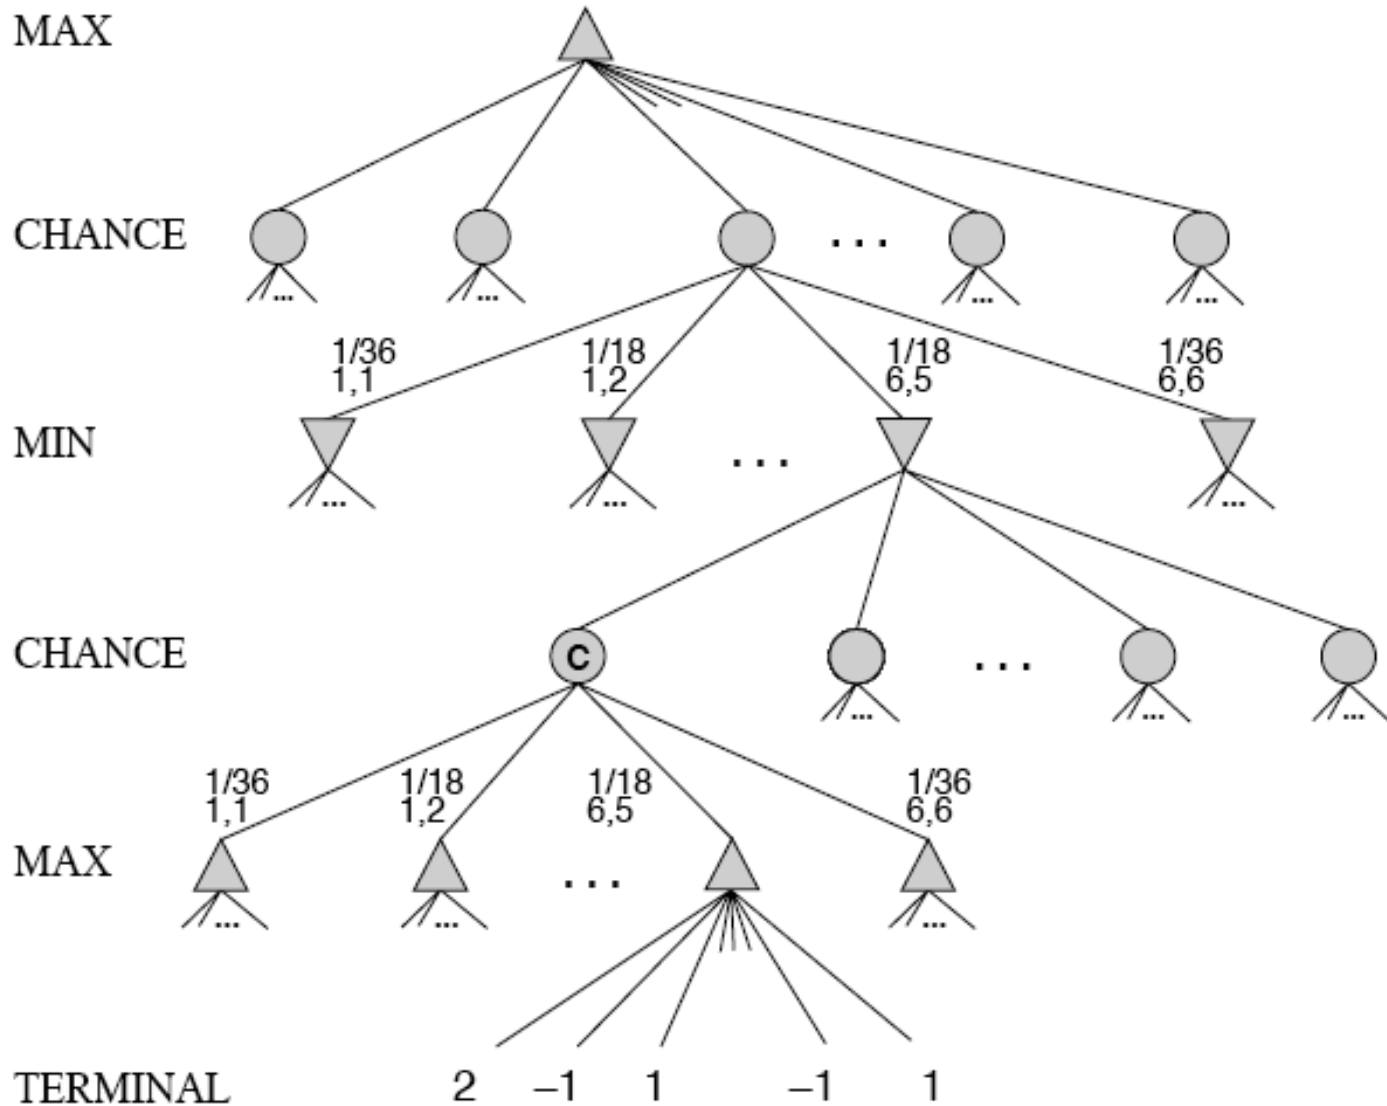
- Backgammon is a two-player game with **uncertainty**.
- Players roll dice to determine what moves to make.
- White has just rolled *5 and 6* and has four legal moves:
  - 5-10, 5-11
  - 5-11, 19-24
  - 5-10, 10-16
  - 5-11, 11-16
- Such games are good for exploring decision making in adversarial problems involving skill and luck




# Why can't we use MiniMax?

- Before a player chooses her move, she rolls dice and then knows exactly what they are
- And the immediate outcome of each move is also known
- But she does not know what moves her opponent will have available in the future
- Need to adapt MiniMax to handle this

# MiniMax trees with Chance Nodes



# Understanding the notation



Board state includes chance outcome  
determining what moves are available

# Game trees with chance nodes

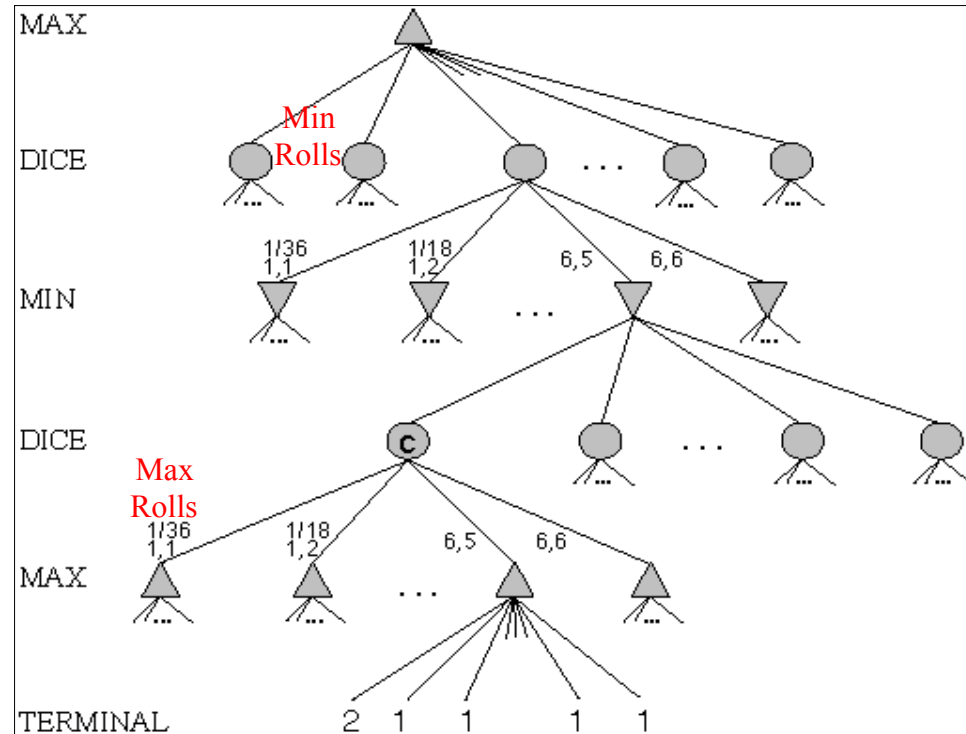
- **Chance nodes** (circles) represent random events
- For a random event with N outcomes, chance node has N children, each with a probability
- 2 dice: 21 distinct outcomes
- Use minimax to compute values for MAX and MIN nodes
- Use **expected values** for chance nodes

- Chance nodes over max node:

$$\text{expectimax}(C) = \sum_i (P(d_i) * \text{maxval}(i))$$

- Chance nodes over min node:

$$\text{expectimin}(C) = \sum_i (P(d_i) * \text{minval}(i))$$






# Impact on Lookahead

- Dice rolls increase branching factor
  - There are 21 possible rolls with two dice
- Backgammon: ~20 legal moves for given roll  
~6K with 1-1 roll (*get to roll again!*)
- At depth 4:  $20 * (21 * 20)**3 \approx 1.2\text{B}$  boards
- As depth increases, probability of reaching a given node shrinks
  - lookahead's value diminished and alpha-beta pruning is much less effective
- [TDGammon](#) used depth-2 search + good static evaluator to achieve world-champion level

# Meaning of the evaluation function



- With probabilities and expected values we must be careful about the *meaning* of values returned by static evaluator
- *relative-order preserving* change of static evaluation values doesn't change minimax decision, but could here
- Linear transformations are OK

# Games of imperfect information

- E.g. card games where opponent's initial hand unknown
  - Can calculate probability for each possible deal
  - Like having one big dice roll at beginning of game
- Possible approach: minimax over each action in each deal; choose action with highest expected value over all deals
- Special case: if action optimal for all deals, it's optimal
- GIB bridge program, approximates this idea by
  1. Generating 100 deals consistent with bidding
  2. Picking action that wins most tricks on average

# High-Performance Game Programs

- Many game programs based on alpha-beta + iterative deepening + extended/singular search + transposition tables + huge databases + ...
- For instance, [Chinook](#) searched all checkers configurations with  $\leq 8$  pieces to create endgame database of 444 billion board configurations
- Methods are general, but implementations improved by many specifically tuned-up enhancements (e.g., the evaluation functions)

# Other Issues

- Multi-player games, no alliances
  - E.g., many card games like Hearts
- Multi-player games with alliances
  - E.g., Risk
  - More on this when we discuss game theory
  - Good model for a social animal like humans, where we are always balancing cooperation and competition

# AI and Games II

- AI also of interest to the video game industry
- Many games include ‘agents’ controlled by the game program that could be
  - Adversaries, in *first person shooter* games
  - Collaborators, in a virtual reality game
- Some game environments used as AI challenges
  - [2009 Mario AI Competition](#)
  - [Unreal Tournament bots](#)

# General Game Playing

- [General Game Playing](#) is an idea developed by Professor Michael Genesereth of Stanford
- See his [site](#) for more information
- Goal: don't develop specialized systems to play specific games (e.g., Checkers) well
- Goal: design AI programs to be able to play more than one game successfully
- Work from a description of a novel game



# General Game Playing

- Stanford's GGP is a Web-based system
- Complete, logical specification of many different games in terms of:
  - relational descriptions of states
  - legal moves and their effects
  - goal relations and their payoffs
- Management of matches between automated players and of competitions involving many players and games



# GGP

- Input: logical description of a game in a custom game description language
- Game bots must
  - Learn how to play legally from description
  - Play well using general problem solving strategies
  - Improve using general machine learning techniques
- Regular completions since 2005, \$10K prize
- Java General Game Playing Base Package

# GGP Peg Jumping Game

; <http://games.stanford.edu/gamemaster/games-debug/peg.kif>

(init (hole a c3 peg))

(init (hole a c4 peg))

...

(init (hole d c4 empty))

...

(<= (next (pegs ?x)) (does jumper (jump ?sr ?sc ?dr ?dc)) (true (pegs ?y))  
(succ ?x ?y)) (<= (next (hole ?sr ?sc empty)) (does jumper (jump ?sr ?sc ?dr ?dc))))

...

(<= (legal jumper (jump ?sr ?sc ?dr ?dc)) (true (hole ?sr ?sc peg))  
(true (hole ?dr ?dc empty)) (middle ?sr ?sc ?or ?oc ?dr ?dc) (true (hole ?or ?oc peg))))






...

(<= (goal jumper 100) (true (hole a c3 empty)) (true (hole a c4 empty))  
(true (hole a c5 empty)))

...

(succ s1 s2)

(succ s2 s3)



# Tic-Tac-Toe in GDL

```
(role xplayer)           ...
(role oplayer)           ...
;; Initial State
(init (cell 1 1 b))      (<= (next (control xplayer))
                          (true (control oplayer)))
(init (cell 1 2 b))      (<= (legal ?w (mark ?x ?y))
                          (true (cell ?x ?y b))
                          (true (control ?w)))
...
(init (control xplayer)) (<= (legal xplayer noop)
                          (true (control oplayer)))
;; Dynamic Components
(<= (next (cell ?m ?n x)) (<= (legal oplayer noop)
                          (true (cell ?m ?n b)))
  (does xplayer (mark ?m ?n))
  (true (cell ?m ?n b)))
...
(<= (next (cell ?m ?n o)) (<= (goal xplayer 100) (line x))
  (does oplayer (mark ?m ?n))
  (true (cell ?m ?n b)))
...
(<= (goal oplayer 0) (line x))
...
```

See ggp-base repository: <http://bit.ly/RB49q5>

# A example of General Intelligence

- [Artificial General Intelligence](#) describes research that aims to create machines capable of general intelligent action
- Harkens back to early visions of AI, like McCarthy's [Advise Taker](#)
  - See [Programs with Common Sense](#) (1959)
- A response to frustration with narrow specialists, often seen as “hacks”
  - See [On Chomsky and the Two Cultures of Statistical Learning](#)

# Why study games?

- Interesting, hard problems that require minimal “initial structure”
- Clear criteria for success
- A way to study problems involving {hostile, adversarial, competing} agents and the uncertainty of interacting with the natural world
- People have used them to assess their intelligence
- Fun, good, easy to understand, PR potential
- Games often define very large search spaces
  - chess  $35^{100}$  nodes in search tree,  $10^{40}$  legal states

# Perspective on Games: **Con** and **Pro**

“Chess is the *Drosophila* of artificial intelligence. However, computer chess has developed much as genetics might have if the geneticists had concentrated their efforts starting in 1910 on breeding racing *Drosophila*. We would have some science, but mainly we would have very fast fruit flies.”

John McCarthy, Stanford

“Saying Deep Blue doesn’t really think about chess is like saying an airplane doesn't really fly because it doesn't flap its wings.”

Drew McDermott, Yale