

## **Outline**



#### **Types of Games**

- > Formulation of games
- Perfect-Information Games
- 202 Minimax and Negamax search
  - $\succ \alpha$ - $\beta$  Pruning
  - > Pruning more
  - Imperfect decision  $\succ$

#### **Stochastic Games**

- > EXPECTIMINIMAX
- **Monte Carlo simulation**
- Partially Observable Games
  - Nash equilibrium







### **Types of Games**

- Adversarial search considers multi-agent and competitive environments.
  - Game theory consider both competitive and cooperative environments.
  - Most common games are deterministic, turn-taking, two-player, zero-sum games with perfect information.
    - > Let's focus on this type of games for a while until told otherwise.

	Beterministie	Stuthastic
Perfect Information (	Chess, Checkers,	Backgammon,
	Go, Othello	Monopoly
Imperfect Information [	Battleships, Bingo	Bridge, Poker



### **Types of Games**



- $S_0$ : Initial state.
- PLAYER(s): The player in state s.
- ACTION(s): Returns the set of legal moves in state s.
- RESULT(s,a): The transition model, which returns the resulting state of a move a in state s.
- TERMINAL-TEST(s): TRUE/FALSE. States where the game has terminated are called terminal states.
- UTILITY(s,p): A utility function (also called objective or payoff).
  > UTILITY(s) for 2-player, zero-sum games.
  - > Reason: UTILITY(s,p1) = -UTILITY(s,p2)



JV.









# MinMax Search



**inputs**: *state*, current state in game

**return** the *a* in ACTIONS(*state*) maximizing MIN-VALUE(RESULT(*a*, *state*))

function MAX-VALUE(state) returns a atility value if TERMINAL-TEST(state) then return UTILITY(state)

 $v \leftarrow -\infty$ 

for a, s in SUCCESSORS(state) do  $v \leftarrow Max(v, MIN-VALUE(s))$ 

return v

function MIN-VALUE(state) returns a utility value if TERMINAL-TEST(state) then return UTILITY(state)  $v \leftarrow \infty$ 

for a, s in SUCCESSORS(*state*) do  $v \leftarrow MIN(v, MAX-VALUE(s))$ return v

### Negamax Search

$$\min(a_0,\ldots,a_n) = -\max\{-a_0,\ldots,-a_n\}$$



Such simplified implementation of *MINIMAX* is called *NEGAMAX*.
 Copying the whole state (line 5) is memory consuming. Practical implementation usually adopts *s* = *BACKTRACK*(s', *a*).

福IFFPDF编

#### NEGAMAX(*s*)

- 1 if TERMINAL-TEST(s)
- 2 return UTILITY(s, p)
- 3 result =  $-\infty$
- 4 for each  $a \in ACTION(s)$

5 
$$s' = \text{RESULT}(s, a)$$
  
6  $result = \max(result)$ 

$$\textit{result} = \max(\textit{result}, -\operatorname{NeGAMAX}(s'))$$

福田FPDF组

7 return result

### **Properties of Minimax(Negamax) Search**



- Optimality: Yes, against an optimal opponent. Otherwise?
  - Risky moves that leads to complicated variations might be better to revert unfavored situations.
- **Time Complexity:**  $O(b^m)$ .
- Space Complexity: 0(bm) (DFS); or 0(m) if the algorithm generates actions one at a time.

For chess,  $b \approx 35$ ,  $m \approx 100 = >$  optimal decision is practically intractable. Do we need to explore every path?







# $\alpha$ - $\beta$ Pruning



- Keeping α (maximum lower bound) for the maximum utility for player MAX, initialized to -∞.
- Keeping β (minimum upper bound) for the minimum utility for player MIN, initialized to ∞.
- Only the moves within the [α, β] window are expanded; otherwise its branches are pruned.
- The pruning does NOT compromise solution quality.

## Implementation of α-β Pruning

#### ALPHABETA( $\boldsymbol{s}, \alpha, \beta$ ) if TERMINAL-TEST(s) 1 return UTILITY(s) 2 if PLAYER(s) == MAX3 $v = -\infty$ 4 福昕PDF编辑器 for each $a \in ACTION(s)$ 5 $v = \max(v, \text{AlphaBeta}(\text{Result}(s, a), \alpha, \beta))$ 6 if $v \geq \beta$ return v 7 $\alpha = \max(\alpha, v)$ 8 9 return v else 10 11 for each $a \in ACTION(s)$ 12 $v = \min(v, \text{ALPHABETA}(\text{RESULT}(s, a), \alpha, \beta))$ 13 if $\alpha \geq v$ return v14 $= \min(\beta, v)$ 15 16 return v









# Implementation of α-β Pruning

■ NEGAMAX + ALPHABATE = AB-NEGAMAX.

#### AB-NEGAMAX( $\boldsymbol{s}, \alpha, \beta$ )



# **Efficiency of α-β Pruning**



Highly depends on the order of moves.

- Good move ordering improves effectiveness of pruning.
- Worst case: no pruning  $\rightarrow O(b^m)$ .
- Best case: Always check the best move first.
  - > Still need to check every move for the first player.
  - > Only need to check one move for the second player.
  - $\stackrel{>}{\succ} O(b \times 1 \times b \times 1...) = O(b^{\frac{m}{2}})$
- Average case:  $O(b^{\frac{3m}{4}})$  Very simple ordering usually achieves  $O(b^{\frac{m}{2}})$ 
  - > Another good reason to adopt iterative deepening.
  - > Reduce the effective branch factor to  $\sqrt{b}$ .
  - > Make the search twice as deep as before.



### **Imperfect Real-Time Decisions**

#### Use CUTOFF-TEST instead of TERMINAL-TEST

- Can be as simple as depth limit.
- Can adopt quiescence search to conquer the horizon effect.
- Yet another good reason to adopt iterative deepening. Return the current best move when time's up.

#### Use EVAL instead of UTILITY

- Usually maps state s into feature space  $f_i$
- Typically use linear combination of features:  $E_{VAL}(s) = \sum_i w_i f_i(s)$
- Need to fine tune weights for strong play.

$$\begin{split} \text{MINIMAX}(s) &= \\ \begin{cases} \text{UTILITY}(s) & \text{if Terminal-Test}(s) \\ \max_{a \in Actions(s)} \text{MINIMAX}(\text{RESULT}(s, a)) & \text{if PLAYER}(s) = \text{MAX} \\ \min_{a \in Actions(s)} \text{MINIMAX}(\text{RESULT}(s, a)) & \text{if PLAYER}(s) = \text{MIN} \end{cases} \end{split}$$

# **Cutting off search**



#### MinimaxCutoff is identical to MinimaxValue except

- *1. Terminal?* is replaced by *Cutoff?*
- 2. Utility is replaced by Eval
- **Does it work in practice?**

b<sup>m</sup> = 10<sup>6</sup>, b=35 → m=4

4-ply lookahead is a hopeless chess player!

4-ply ≈ human novice

- 8-ply ≈ typical PC, human master
- 12-ply ≈ Deep Blue, Kasparov









#### Heuristic Where Minimax May Go Wrong



- MINIMAX chooses the right branch.
- **EVAL** with an error with zero mean and standard deviation  $\sigma$ .
- $\sigma$  = 3, the left branch is better 54.6% of the time.
- $\sigma$  = 5, the left branch is better 64.4% of the time.

## **Forward Pruning**



- Forward pruning does compromise solution quality (so is using EVAL)
  - > Some moves are pruned immediately without further consideration.
- Beam search is one way to forward pruning. Dangerous since the best move might be pruned.
- **PROBCUT** uses the scores from previous searches to estimate the probability that a node is outside the  $[\alpha, \beta]$  window

### Search vs. Table Lookup



- For many games, deep search usually helps little at the beginning.
- Instead, fast table looking up, huge databases, and statistical analysis help more.
- Table lookup also helps a lot toward the end of games.
  - The KBNK (king, bishop, knight vs. king) lookup: 462x62x61x2 = 3, 494, 568 possibilities.
  - Bourzutschky (2006) solved all pawn-less 6-piece endgames and some 7-piece endgames.

A KQNKRBN endgame requires 517 moves!

Finally, early exchange favors computers than humans -> deeper search and more probable falls in lookup.











### **Sensitivity to Heuristic**

- As mentioned before, we rarely can actually use UTILITY in EXPECTMINIMAX.
- Instead, we use heuristic.
- However, unlike in MINIMAX, actual values of heuristic matter now. (In MINIMAX, only relative order matters.)





### **Performance of EXPECTIMINIMAX**

- $\alpha$ - $\beta$  pruning now does not apply to *MAX/MIN* nodes (why?).
- $\alpha$ - $\beta$  pruning now still applies to *CHANCE* nodes (why?).
- **Time Complexity:**  $O(b^m) \rightarrow O(b^m n^m)$ , where *n* is the number of distinct dice rolls.
- Causes *EXPECTIMINIMAX* impractical in many cases.
- Solution: Instead of checking every MAX/MIN node, adopts Monte Carlo simulation at chance nodes.
- Using random dice rolls to check only a certain number (decided by quality/time limit) of paths.

argmax 
$$\sum_{s} p(s)MINMAX(RESULT(s, a))$$
  
Nonte Carlo simulation  $\arg \max_{a} \frac{1}{N} \sum_{i=1}^{N} MINIMAX(RESULT(s_i, a))$ 



 $Eval = 1/10^{*}[(1) + (3) + (50) + (50) + (50) + (-50) + (-50) + (50) + (15) + (-5)] = 11.4$ 



numerical results.



Monte Carlo Method: More sampling, more closer to the optimal solution Las Vegas Method: More sampling, higher opportunity to optimal solution.



- Different from stochastic games, unobservable parts are usually controlled by opponents, not probability.
- Examples: Cards held by other player in bridge, folded cards in poker, fogs in star craft.
- Different strategies may be applied and may all considered optimum against different opponents.
- If equilibrium exists, it' s usually considered as optimum strategy.





# Nash Equilibrium

By John Nash — check out "A Beautiful Mind (2001)" if you want an informal introduction to him.

#### Information Definition

A set of strategies is a Nash equilibrium if no player can do better by unilaterally changing his or her strategy.

#### Prisoners' dilemma

福BlfpDF编辑器	B stays silent	B confesses
A stays silent	Each serves 1 yr	A: 5 yrs; B: free
A confesses	A: free; B: 5 yrs	Each serves 3 yrs

Read Chapter 17 if you want to know more.

# Summary



- MINIMAX search for zero-sum two-player games.
- Pruning techniques enable to search deeper.
- Due to time limit, heuristics are used to evaluate the "goodness" for a player.
- For stochastic games, we need to introduce chance nodes and search for expected maximum/minimum values.
- For stochastic games, α-β pruning is much less efficient, Monte Carlo simulations are often adopted to speed up the search.
- With limited observation, optimality is usually not well defined. If equilibrium exists, strategies in equilibrium are often considered optimum.



