



Artificial Intelligence 1: game playing

Lecturer: Tom Lenaerts

Institut de Recherches Interdisciplinaires et de
Développements en Intelligence Artificielle
(IRIDIA)

Université Libre de Bruxelles



Outline

- What are games?
- Optimal decisions in games
 - **Which strategy leads to success?**
- α - β pruning
- Games of imperfect information
- Games that include an element of chance

What are and why study games?

- Games are a form of *multi-agent environment*
 - **What do other agents do and how do they affect our success?**
 - **Cooperative vs. competitive multi-agent environments.**
 - **Competitive multi-agent environments give rise to adversarial problems a.k.a. *games***
- Why study games?
 - **Fun; historically entertaining**
 - **Interesting subject of study because they are hard**
 - **Easy to represent and agents restricted to small number of actions**

November 2, 2004

TLo (IRIDIA)

3

Relation of Games to Search

- Search – no adversary
 - **Solution is (heuristic) method for finding goal**
 - **Heuristics and CSP techniques can find *optimal* solution**
 - **Evaluation function: estimate of cost from start to goal through given node**
 - **Examples: path planning, scheduling activities**
- Games – adversary
 - **Solution is strategy (strategy specifies move for every possible opponent reply).**
 - **Time limits force an *approximate* solution**
 - **Evaluation function: evaluate “goodness” of game position**
 - **Examples: chess, checkers, Othello, backgammon**

November 2, 2004

TLo (IRIDIA)

4

Types of Games

	deterministic	chance
perfect information	chess, checkers, go, othello	backgammon monopoly
imperfect information		bridge, poker, scrabble nuclear war

November 2, 2004

TLo (IRIDIA)

5

Game setup


- Two players: MAX and MIN
- MAX moves first and they take turns until the game is over. Winner gets award, loser gets penalty.
- Games as search:
 - **Initial state: e.g. board configuration of chess**
 - **Successor function: list of (move,state) pairs specifying legal moves.**
 - **Terminal test: Is the game finished?**
 - **Utility function: Gives numerical value of terminal states. E.g. win (+1), loose (-1) and draw (0) in tic-tac-toe (next)**
- MAX uses search tree to determine next move.

November 2, 2004

TLo (IRIDIA)

6

Partial Game Tree for Tic-Tac-Toe



November 2, 2004

o (IRIDIA)

7

Optimal strategies

- Find the contingent *strategy* for MAX assuming an infallible MIN opponent.
- Assumption: Both players play optimally !!
- Given a game tree, the optimal strategy can be determined by using the minimax value of each node:

MINIMAX-VALUE(n) =

UTILITY(n)

$\max_{s \in \text{successors}(n)} \text{MINIMAX-VALUE}(s)$

$\min_{s \in \text{successors}(n)} \text{MINIMAX-VALUE}(s)$

If n is a terminal

If n is a max node


If n is a min node

November 2, 2004

TLo (IRIDIA)

8

Two-Ply Game Tree




November 2, 2004

TLo (IRIDIA)

9

Two-Ply Game Tree




November 2, 2004

TLo (IRIDIA)

10

Two-Ply Game Tree

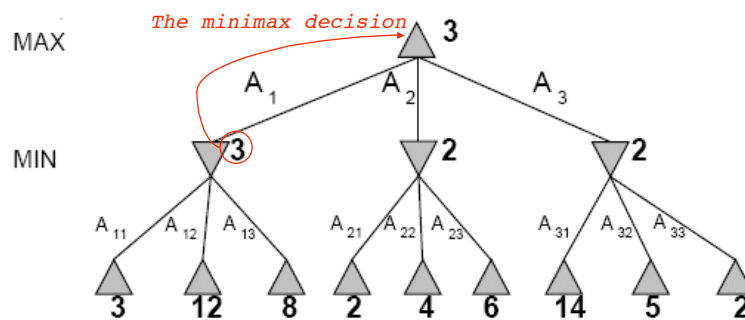


November 2, 2004

TLo (IRIDIA)

11

Two-Ply Game Tree



Minimax maximizes the worst-case outcome for max.

November 2, 2004

TLo (IRIDIA)

12

What if MIN does not play optimally?

- Definition of optimal play for MAX assumes MIN plays optimally: maximizes worst-case outcome for MAX.
- But if MIN does not play optimally, MAX will do even better. [Can be proved.]

Minimax Algorithm

function MINIMAX-DECISION(*state*) *returns an action*

inputs: *state*, current state in game

$v \leftarrow \text{MAX-VALUE}(state)$

return the *action* in SUCCESSORS(*state*) with value v

function MAX-VALUE(*state*) *returns a utility value*

if TERMINAL-TEST(*state*) **then return** UTILITY(*state*)

$v \leftarrow \infty$

for a, s in SUCCESSORS(*state*) **do**

$v \leftarrow \text{MAX}(v, \text{MIN-VALUE}(s))$

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**

$v \leftarrow \text{MIN}(v, \text{MAX-VALUE}(s))$

return v

Properties of Minimax

Criterion	Minimax
Complete?	Yes ☺
Time	$O(b^m)$ ☹
Space	$O(bm)$ ☺
Optimal?	Yes ☹


November 2, 2004

TLo (IRIDIA)

15

Multiplayer games

- Games allow more than two players
- Single minimax values becomes vector



November 2, 2004

TLo (IRIDIA)

16

Problem of minimax search

- Number of games states is exponential to the number of moves.
 - **Solution: Do not examine every node**
 - **==> Alpha-beta pruning**
 - Alpha = value of best choice found so far at any choice point along the MAX path
 - Beta = value of best choice found so far at any choice point along the MIN path
- Revisit example ...


November 2, 2004

TLo (IRIDIA)

17

Alpha-Beta Example

Do DF-search until first leaf




November 2, 2004

TLo (IRIDIA)

18

Alpha-Beta Example (continued)




November 2, 2004

TLo (IRIDIA)

19

Alpha-Beta Example (continued)




November 2, 2004

TLo (IRIDIA)

20

Alpha-Beta Example (continued)




November 2, 2004

TLo (IRIDIA)

21

Alpha-Beta Example (continued)




November 2, 2004

TLo (IRIDIA)

22

Alpha-Beta Example (continued)




November 2, 2004

TLo (IRIDIA)

23

Alpha-Beta Example (continued)




November 2, 2004

TLo (IRIDIA)

24

Alpha-Beta Example (continued)




November 2, 2004

TLo (IRIDIA)

25

Alpha-Beta Example (continued)



November 2, 2004

TLo (IRIDIA)

26

Alpha-Beta Algorithm

function ALPHA-BETA-SEARCH(*state*) *returns an action*
inputs: *state*. current state in game
 $v \leftarrow \text{MAX-VALUE}(\textit{state}, -\infty, +\infty)$
return the *action* in SUCCESSORS(*state*) with value *v*


function MAX-VALUE(*state*, α , β) *returns a utility value*
if TERMINAL-TEST(*state*) **then return** UTILITY(*state*)
 $v \leftarrow -\infty$
for *a,s* in SUCCESSORS(*state*) **do**
 $v \leftarrow \text{MAX}(v, \text{MIN-VALUE}(s), \alpha, \beta)$
if $v \geq \beta$ **then return** *v*
 $\alpha \leftarrow \text{MAX}(\alpha, v)$
return *v*

Alpha-Beta Algorithm

function MIN-VALUE(*state*, α , β) *returns a utility value*
if TERMINAL-TEST(*state*) **then return** UTILITY(*state*)
 $v \leftarrow +\infty$
for *a,s* in SUCCESSORS(*state*) **do**
 $v \leftarrow \text{MIN}(v, \text{MAX-VALUE}(s), \alpha, \beta)$
if $v \leq \alpha$ **then return** *v*
 $\beta \leftarrow \text{MIN}(\beta, v)$
return *v*

General alpha-beta pruning

- Consider a node n somewhere in the tree
- If player has a better choice at
 - **Parent node of n**
 - **Or any choice point further up**
- n will **never** be reached in actual play.
- Hence when enough is known about n , it can be pruned.



November 2, 2004

TLo (IRIDIA)

29

Final Comments about Alpha-Beta Pruning

- Pruning does not affect final results
- Entire subtrees can be pruned.
- Good move *ordering* improves effectiveness of pruning
- With “perfect ordering,” time complexity is $O(b^{m/2})$
 - **Branching factor of \sqrt{b} !!**
 - **Alpha-beta pruning can look twice as far as minimax in the same amount of time**
- Repeated states are again possible.
 - **Store them in memory = transposition table**

November 2, 2004

TLo (IRIDIA)

30

Games of imperfect information

- Minimax and alpha-beta pruning require too much leaf-node evaluations.
- May be impractical within a reasonable amount of time.
- SHANNON (1950):
 - **Cut off search earlier (replace TERMINAL-TEST by CUTOFF-TEST)**
 - **Apply heuristic evaluation function EVAL (replacing utility function of alpha-beta)**


Cutting off search

- Change:
 - **If TERMINAL-TEST(*state*) then return UTILITY(*state*)**into
 - **If CUTOFF-TEST(*state*,*depth*) then return EVAL(*state*)**
- Introduces a fixed-depth limit *depth*
 - **Is selected in so that the amount of time will not exceed what the rules of the game allow.**
- When cutoff occurs, the evaluation is performed.


Heuristic EVAL

- Idea: produce an estimate of the expected utility of the game from a given position.
- Performance depends on quality of EVAL.
- Requirements:
 - **EVAL should order terminal-nodes in the same way as UTILITY.**
 - **Computation may not take too long.**
 - **For non-terminal states the EVAL should be strongly correlated with the actual chance of winning.**
- Only useful for quiescent (no wild swings in value in near future) states


Heuristic EVAL example




(a) White to move
Fairly even



(b) Black to move
White slightly better




(c) White to move
Black winning




(d) Black to move
White about to lose

$$Eval(s) = w_1 f_1(s) + w_2 f_2(s) + \dots + w_n f_n(s)$$

Heuristic EVAL example




(a) White to move
Fairly even



(b) Black to move
White slightly better



(c) White to move
Black winning



(d) Black to move
White about to lose

Addition assumes independence

$$Eval(s) = w_1 f_1(s) + w_2 f_2(s) + \dots + w_n f_n(s)$$

Heuristic difficulties

Heuristic counts pieces won




(a) White to move



(b) White to move

Horizon effect

Fixed depth search
thinks it can avoid
the queening move





November 2, 2004

TLo (IRIDIA)

37

Games that include chance





- Possible moves (5-10,5-11), (5-11,19-24),(5-10,10-16) and (5-11,11-16)

November 2, 2004

TLo (IRIDIA)

38

Games that include chance





- Possible moves (5-10,5-11), (5-11,19-24),(5-10,10-16) and (5-11,11-16)
- [1,1], [6,6] chance 1/36, all other chance 1/18

November 2, 2004

TLo (IRIDIA)

39

Games that include chance



- [1,1], [6,6] chance 1/36, all other chance 1/18
- Can not calculate definite minimax value, only *expected* value

November 2, 2004

TLo (IRIDIA)

40


Expected minimax value

EXPECTED-MINIMAX-VALUE(n)=

$$\begin{aligned}
 & \text{UTILITY}(n) && \text{If } n \text{ is a terminal} \\
 & \max_{s \in \text{successors}(n)} \text{MINIMAX-VALUE}(s) && \text{If } n \text{ is a max node} \\
 & \min_{s \in \text{successors}(n)} \text{MINIMAX-VALUE}(s) && \text{If } n \text{ is a max node} \\
 & \sum_{s \in \text{successors}(n)} P(s) \cdot \text{EXPECTEDMINIMAX}(s) && \text{If } n \text{ is a chance node}
 \end{aligned}$$

These equations can be backed-up recursively all the way to the root of the game tree.


Position evaluation with chance nodes



- Left, A1 wins
- Right A2 wins
- Outcome of evaluation function may not change when values are scaled differently.
- Behavior is preserved only by a *positive linear* transformation of EVAL.

Discussion

- Examine section on state-of-the-art games yourself
- Minimax assumes right tree is better than left, yet ...
 - **Return probability distribution over possible values**
 - **Yet expensive calculation**




November 2, 2004

TLo (IRIDIA)

43

Discussion

- Utility of node expansion
 - **Only expand those nodes which lead to significantly better moves**
- Both suggestions require meta-reasoning



November 2, 2004

TLo (IRIDIA)

44

Summary

- Games are fun (and dangerous)
- They illustrate several important points about AI
 - **Perfection is unattainable -> approximation**
 - **Good idea what to think about**
 - **Uncertainty constrains the assignment of values to states**
- Games are to AI as grand prix racing is to automobile design.