

Artificial Intelligence

Adversarial search Chapter 6, AIMA



This presentation owes a lot to V. Pavlovic @ Rutgers, who borrowed from J. D. Skrentny, who in turn borrowed from C. Dyer,...

Adversarial search

- At least two agents and a competitive environment: Games, economies.
- Games and AI:
 - Generally considered to require intelligence (to win)
 - Have to evolve in real-time
 - Well-defined and limited environment

Board games

© Thierry Dichtenmuller



Games & AI

	Deterministic	Chance
perfect info	Checkers, Chess, Go, Othello	Backgammon, Monopoly
imperfect info		Bridge, Poker, Scrabble

Games and search

Traditional search: single agent, searches for its well-being, unobstructed

Games: search against an opponent

Example: two player board game (chess, checkers, tic-tac-toe,...)



Board configuration: unique arrangement of "pieces"

Representing board games as goal-directed search problem (States = board configurations):

- **Initial state**: Current board configuration
- **Successor function**: Legal moves
- **Goal state**: Winning/terminal board configuration
- **Utility function**: Current board configuration

Example: Tic-tac-toe

- **Initial state**: 3x3 empty table.
- **Successor function**: Players take turns marking X or O in the table cells.
- **Goal state**: When all the table cells are filled or when either player has three symbols in a row.
- **Utility function**: +1 for three in a row, -1 if the opponent has three in a row, 0 if the table is filled and no-one has three symbols in a row.




The minimax principle

Assume the opponent plays to win and always makes the best possible move.

The **minimax value** for a node = the utility for you of being in that state, assuming that both players (you and the opponent) play optimally from there on to the end.

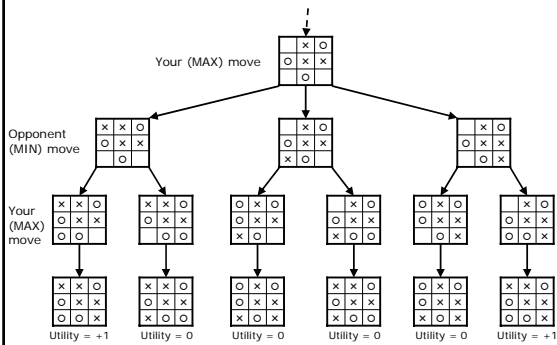
Terminology:
MAX = you, MIN = the opponent.

Example: Tic-tac-toe

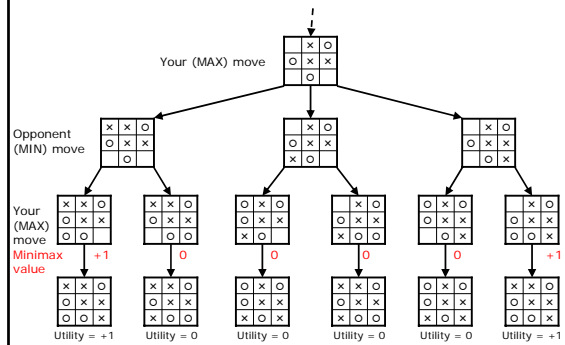


Assignment: Expand this tree to the end of the game.

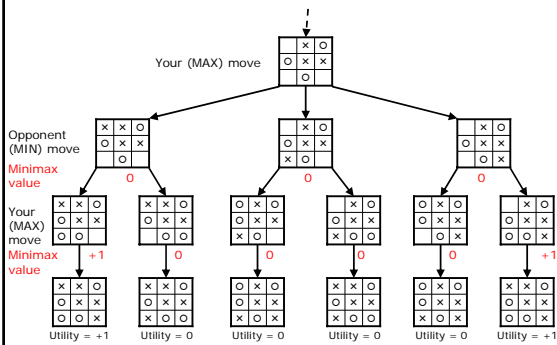
Example: Tic-tac-toe



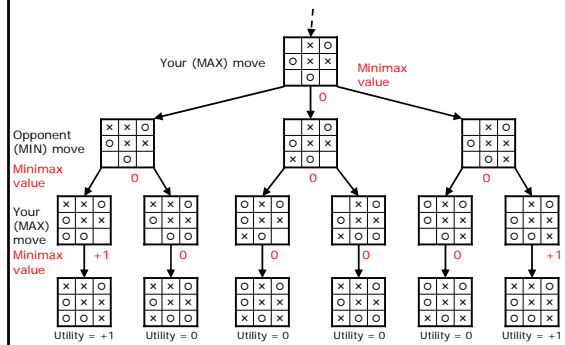
Example: Tic-tac-toe



Example: Tic-tac-toe



Example: Tic-tac-toe



The minimax value

Minimax value for node n =

{	Utility(n)	If n is a terminal node
	Max(Minimax-values of successors)	If n is a MAX node
	Min(Minimax-values of successors)	If n is a MIN node

High utility favours you (MAX), therefore choose move with highest utility

Low utility favours the opponent (MIN), therefore choose move with lowest utility

The minimax algorithm

1. Start with utilities of terminal nodes
2. Propagate them back to root node by choosing the **minimax** strategy




Figure borrowed from V. Pavlovic

The minimax algorithm

1. Start with utilities of terminal nodes
2. Propagate them back to root node by choosing the **minimax** strategy




Figure borrowed from V. Pavlovic

The minimax algorithm

1. Start with utilities of terminal nodes
2. Propagate them back to root node by choosing the **minimax** strategy




Figure borrowed from V. Pavlovic


Complexity of minimax algorithm

- A depth-first search
 - Time complexity $O(b^d)$
 - Space complexity $O(bd)$
- Time complexity impossible in real games (with time constraints) except in very simple games (e.g. tic-tac-toe)


Strategies to improve minimax

1. Remove redundant search paths
– **symmetries**
2. Remove uninteresting search paths
– **alpha-beta pruning**
3. Cut the search short before goal
– **Evaluation functions**
4. Book moves

1. Remove redundant paths



Tic-tac-toe has mirror symmetries
and rotational symmetries



First three levels of the tic-tac-toe state space reduced by symmetry




Image from G. F. Luger, "Artificial Intelligence", 2002


2. Remove uninteresting paths

If the player has a better choice m at n 's parent node, or at any node further up, then node n will never be reached.

Prune the entire path below node m 's parent node (except for the path that m belongs to, and paths that are equal to this path).

Minimax is depth-first \rightarrow keep track of highest (α) and lowest (β) values so far.


Called alpha-beta pruning.



Alpha-Beta Example

$\text{minimax}(A, 0, 4)$


$\text{minimax}(\text{node}, \text{level}, \text{depth limit})$



Slide adapted from V. Pavlovic

Alpha-Beta Example


$\text{minimax}(B, 1, 4)$



Slide adapted from V. Pavlovic

Alpha-Beta Example


$\text{minimax}(F, 2, 4)$



Slide adapted from V. Pavlovic

Alpha-Beta Example

minimax(N, 3, 4)




Slide adapted from V. Pavlovic

Alpha-Beta Example

minimax(F, 2, 4) is returned to


alpha = 4, maximum seen so far



Slide adapted from V. Pavlovic

Alpha-Beta Example


minimax(O, 3, 4)



Slide adapted from V. Pavlovic

Alpha-Beta Example

minimax(W, 4, 4)




Slide adapted from V. Pavlovic

Alpha-Beta Example

minimax(O, 3, 4) is returned to


beta = -3, minimum seen so far



Slide adapted from V. Pavlovic

Alpha-Beta Example

O's beta (-3) < F's alpha (4): Stop expanding O (alpha cut-off)




Slide adapted from V. Pavlovic

Alpha-Beta Example

Why?


Smart opponent selects W or worse \rightarrow O's upper bound is -3
So MAX shouldn't select O: -3 since N: 4 is better



Alpha-Beta Example

$\text{minimax}(F, 2, 4)$ is returned to


alpha not changed (maximizing)



Alpha-Beta Example


$\text{minimax}(B, 1, 4)$ is returned to

beta = 4, minimum seen so far



Alpha-Beta Example


$\text{minimax}(G, 2, 4)$



Alpha-Beta Example

$\text{minimax}(B, 1, 4)$ is returned to


beta = -5 , minimum seen so far



Alpha-Beta Example


$\text{minimax}(A, 0, 4)$ is returned to

alpha = -5 , maximum seen so far



Alpha-Beta Example


minimax(C, 1, 4)



Slide adapted from V. Pavlovic

Alpha-Beta Example

minimax(H, 2, 4)




Slide adapted from V. Pavlovic

Alpha-Beta Example

minimax(C, 1, 4) is returned to


beta = -10, minimum seen so far



Slide adapted from V. Pavlovic

Alpha-Beta Example


C's beta (-10) < A's alpha (-5): Stop expanding C (alpha cut-off)



Slide adapted from V. Pavlovic

Alpha-Beta Example


minimax(D, 1, 4)



Slide adapted from V. Pavlovic

Alpha-Beta Example


minimax(D, 1, 4) is returned to




Slide adapted from V. Pavlovic

Alpha-Beta Example

$\text{minimax}(D, 1, 4)$ is returned to

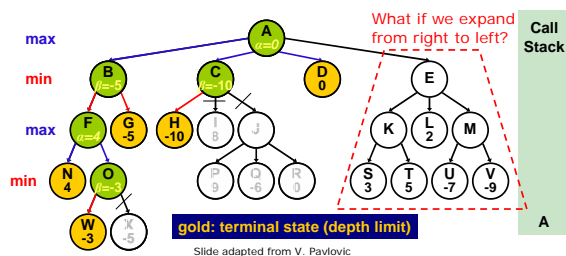


Alpha-Beta Example




Alpha-Beta Example

$\text{minimax}(D, 1, 4)$ is returned to



Alpha-Beta Example



Alpha-Beta pruning rule


Stop expanding

max node n if $\alpha(n) > \beta$ higher in the tree
 min node n if $\beta(n) < \alpha$ higher in the tree

Alpha-Beta pruning rule

Stop expanding


max node n if $\alpha(n) > \beta$ higher in the tree
 min node n if $\beta(n) < \alpha$ higher in the tree



Alpha-Beta pruning rule

Stop expanding

max node n if $\alpha(n) > \beta$ higher in the tree
 min node n if $\beta(n) < \alpha$ higher in the tree




Which nodes will not be expanded when expanding from left to right?

Alpha-Beta pruning rule

Stop expanding

max node n if $\alpha(n) > \beta$ higher in the tree
 min node n if $\beta(n) < \alpha$ higher in the tree




Which nodes will not be expanded when expanding from right to left?

Alpha-Beta pruning rule

Stop expanding

max node n if $\alpha(n) > \beta$ higher in the tree
 min node n if $\beta(n) < \alpha$ higher in the tree



Which nodes will not be expanded when expanding from right to left?

3. Cut the search short

- Use depth-limit and estimate utility for non-terminal nodes (evaluation function)
 - Static board evaluation (SBE)
 - Must be easy to compute

Example, chess:

$$SBE = \alpha \text{ "Material Balance"} + \beta \text{ "Center Control"} + \gamma \dots$$

Material balance = value of white pieces – value of black pieces, where pawn = +1, knight & bishop = +3, rook = +5, queen = +9, king = ?

The parameters ($\alpha, \beta, \gamma, \dots$) can be learned (adjusted) from experience.

http://en.wikipedia.org/wiki/Computer_chess

Leaf evaluation

For most chess positions, computers cannot look ahead to all final possible positions. Instead, they must look ahead a few plies and then evaluate the final board position. The algorithm that evaluates final board positions is termed the "evaluation function", and these algorithms are often vastly different between different chess programs.

Nearly all evaluation functions evaluate positions in units and at the least consider material value. Thus, they will count up the amount of material on the board for each side (where a **pawn** is worth exactly 1 point, a **knight** is worth 3 points, a **bishop** is worth 3 points, a **rook** is worth 5 points and a **queen** is worth 9 points). The **king** is impossible to value since its loss causes the loss of the game. For the purposes of programming chess computers, however, it is often assigned a value of appr. 200 points.

Evaluation functions take many other factors into account, however, such as pawn structure, the fact that doubled bishops are usually worth more, centralized pieces are worth more, and so on. The protection of kings is usually considered, as well as the phase of the game (opening, middle or endgame).

4. Book moves

- Build a database (look-up table) of endgames, openings, etc.
- Use this instead of minimax when possible.

http://en.wikipedia.org/wiki/Computer_chess

Using endgame databases

...


Nalimov Endgame Tablebases, which use state-of-the-art compression techniques, require 7.05 GB of hard disk space for all five-piece endings. It is estimated that to cover all the six-piece endings will require at least 1 [terabyte](#). Seven-piece tablebases are currently a long way off.

While Nalimov Endgame Tablebases handle [en passant](#) positions, they assume that castling is not possible. This is a minor flaw that is probably of interest only to endgame specialists.

More importantly, they do not take into account the fifty move rule. Nalimov Endgame Tablebases only store the shortest possible mate (by ply) for each position. However in certain rare positions, the stronger side cannot force win before running into the [fifty move rule](#). A chess program that searches the database and obtains a value of mate in 85 will not be able to know in advance if such a position is actually a draw according to the fifty move rule, or if it is a win, because there will be a piece exchange or pawn move along the way. Various solutions including the addition of a "distance to zero" (DTZ) counter have been proposed to handle this problem, but none have been implemented yet.

Games with chance

- Dice games, card games, ...
- Extend the minimax tree with chance layers.



Animation adapted from V. Pavlovic