## Game Playing

Why do AI researchers study game playing?

1. It's a good reasoning problem, formal and nontrivial.
2. Direct comparison with humans and other computer programs is easy.

## What Kinds of Games?

Mainly games of strategy with the following characteristics:

1. Sequence of moves to play
2. Rules that specify possible moves
3. Rules that specify a payment for each move
4. Objective is to maximize your payment

## Games vs. Search Problems

- Unpredictable opponent $\rightarrow$ specifying a move for every possible opponent reply
- Time limits $\rightarrow$ unlikely to find goal, must approximate


## Two-Player Game



# Game Tree (2-player, Deterministic, Turns) 



## Mini-Max Terminology

- utility function: the function applied to leaf nodes
- backed-up value
- of a max-position: the value of its largest successor
- of a min-position: the value of its smallest successor
- minimax procedure: search down several levels; at the bottom level apply the utility function, back-up values all the way up to the root node, and that node selects the move.


## Minimax

- Perfect play for deterministic games
- Idea: choose move to position with highest minimax value

> = best achievable payoff against best play

MAX

MIN


## Minimax Strategy

- Why do we take the min value every other level of the tree?
- These nodes represent the opponent's choice of move.
- The computer assumes that the human will choose that move that is of least value to the computer.


## Minimax algorithm

```
function Minimax-Decision(state) returns an action
v\leftarrow 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 Utllity(state)
    v\leftarrow-\infty
    for a,s in SUCCESSORS(state) do
        v \leftarrow \operatorname { M a x } ( v , \operatorname { M i n } - \operatorname { V a l u e } ( s ) )
    return v
```

function Min-Value(state) returns a utility value
if Terminal-Test(state) then return Utllity(state)
$v \leftarrow \infty$
for $a, s$ in Successors(state) do
$v \leftarrow \operatorname{Min}(v, \operatorname{Max}-\operatorname{Value}(s))$
return $v$

## Tic Tac Toe



- Let p be a position in the game
- Define the utility function $f(p)$ by
$-f(p)=$
- largest positive number if $p$ is a win for computer
- smallest negative number if $p$ is a win for opponent
- RCDC - RCDO
- where RCDC is number of rows, columns and diagonals in which computer could still win
- and RCDO is number of rows, columns and diagonals in which opponent could still win.


## Sample Evaluations

- $\mathrm{X}=$ Computer; $\mathrm{O}=$ Opponent


$$
\begin{array}{l|l|l}
\mathrm{O} & \mathrm{X} \\
\hline \mathrm{X} & \mathrm{X} \\
\hline & &
\end{array}
$$

|  | $X$ | $O$ |
| :--- | :--- | :--- |
| rows  <br> cols  <br> diags  |  |  |


|  | $X$ | $O$ |
| :--- | :--- | :--- |
| rows  <br> cols  <br> diags  |  |  |

## Minimax is done depth-first



## Properties of Minimax

- Complete? Yes (if tree is finite)
- Optimal? Yes (against an optimal opponent)
- Time complexity? $\mathrm{O}\left(\mathrm{b}^{m}\right)$
- Space complexity? O(bm) (depth-first exploration)
- For chess, $b \approx 35, m \approx 100$ for "reasonable" games $\rightarrow$ exact solution completely infeasible Need to speed it up.


## Alpha-Beta Procedure

- The alpha-beta procedure can speed up a depth-first minimax search.
- Alpha: a lower bound on the value that a max node may ultimately be assigned

$$
v \geq \alpha
$$

- Beta: an upper bound on the value that a minimizing node may ultimately be assigned

$$
v \leq \beta
$$

## $\alpha-\beta$ pruning example

MAX


## $\alpha-\beta$ pruning example



## $\alpha-\beta$ pruning example



## $\alpha-\beta$ pruning example



## $\alpha-\beta$ pruning example



## Alpha Cutoff



What happens here? Is there an alpha cutoff?

## Beta Cutoff



## Alpha-Beta Pruning



## Alpha-Beta Pruning



## Properties of $\alpha-\beta$

- Pruning does not affect final result. This means that it gets the exact same result as does full minimax.
- Good move ordering improves effectiveness of pruning
- With "perfect ordering," time complexity $=\mathrm{O}\left(\mathrm{b}^{m / 2}\right)$
$\rightarrow$ doubles depth of search
- A simple example of the value of reasoning about which computations are relevant (a form of metareasoning)


## The $\alpha-\beta$ algorithm

function Alpha-Beta-Search(state) returns an action inputs: state, current state in game
$v \leftarrow \operatorname{Max}-\operatorname{Value}($ state $,-\infty,+\infty)$
return the action in SUCCESSORS(state) with value $v$
function Max-Value(state, $\alpha, \beta$ ) returns a utility value
inputs: state, current state in game
$\alpha$, the value of the best alternative for MAX along the path to state
$\beta$, the value of the best alternative for min along the path to state
if Terminal-Test(state) then return Utility (state)
$v \leftarrow-\infty$
for $a, s$ in Successors(state) do $v \leftarrow \operatorname{Max}(v, \operatorname{Min}-\operatorname{Value}(s, \alpha, \beta))$ if $v \geq \beta$ then return $v$ cutoff $\alpha \leftarrow \operatorname{MAX}(\alpha, v)$
return $v$


## The $\alpha-\beta$ algorithm

```
function Min-VALuE(state, \alpha, \beta) returns a utility value
    inputs: state, current state in game
    \alpha, the value of the best alternative for MAX along the path to state
    \beta, the value of the best alternative for mIN along the path to state
    if Terminal-Test(state) then return Utility(state)
    v\leftarrow+\infty
    for a,s in SUCCESSORS(state) do
        v\leftarrow\operatorname{Min}(v,\operatorname{Max}-\operatorname{Value}(s,\alpha,\beta))
        if v\leq\alpha then return v Cutoff
        \beta\leftarrow\operatorname{Min}(\beta,v)
    return v
```


## When do we get alpha cutoffs?



## Shallow Search Techniques

1. limited search for a few levels
2. reorder the level-1 sucessors
3. proceed with $\alpha-\beta$ minimax search

## Additional Refinements

- Waiting for Quiescence: continue the search until no drastic change occurs from one level to the next.
- Secondary Search: after choosing a move, search a few more levels beneath it to be sure it still looks good.
- Book Moves: for some parts of the game (especially initial and end moves), keep a catalog of best moves to make.


## Evaluation functions

- For chess/checkers, typically linear weighted sum of features

$$
\operatorname{Eval}(s)=w_{1} f_{1}(s)+w_{2} f_{2}(s)+\ldots+w_{n} f_{n}(s)
$$

- e.g., $\mathrm{w}_{1}=9$ with
$\mathrm{f}_{1}(\mathrm{~s})=$ (number of white queens) - (number of black queens), etc.


## Example: Samuel's CheckerPlaying Program

- It uses a linear evaluation function
$f(n)=a_{1} x_{1}(n)+a_{2} x_{2}(n)+\ldots+a_{m} x_{m}(n)$

For example: $f=6 K+4 M+U$

- K = King Advantage
- M = Man Advantage
- U = Undenied Mobility Advantage (number of moves that Max has that Min can't jump after)


## Samuel's Checker Player

- In learning mode
- Computer acts as 2 players: $A$ and $B$
- A adjusts its coefficients after every move
- B uses the static utility function
- If $A$ wins, its function is given to $B$


## Samuel's Checker Player

- How does A change its function?


## 1. Coefficent replacement

$\triangle$ (node ) = backed-up value(node) - initial value(node)
if $\Delta>0$ then terms that contributed positively are given more weight and terms that contributed negatively get less weight
if $\Delta<0$ then terms that contributed negatively are given more weight and terms that contributed positively get less weight

## Samuel's Checker Player

- How does A change its function?

2. Term Replacement 38 terms altogether 16 used in the utility function at any one time

Terms that consistently correlate low with the function value are removed and added to the end of the term queue.

They are replaced by terms from the front of the term queue.

## Kalah



To move, pick up all the stones in one of your holes, and put one stone in each hole, starting at the next one, including your Kalah and skipping the opponent's Kalah.

## Kalah

- If the last stone lands in your Kalah, you get another turn.
- If the last stone lands in your empty hole, take all the stones from your opponent's hole directly across from it and put them in your Kalah AND put that last stone in your Kalah too.
- If all of your holes become empty, the opponent keeps the rest of the stones.
- The winner is the player who has the most stones in his Kalah at the end of the game.


## 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?

$$
b^{m}=10^{6}, b=35 \rightarrow m=4
$$

4-ply lookahead is a hopeless chess player!

- 4-ply $\approx$ human novice
- 8-ply $\approx$ typical PC, human master
- 12-ply $\approx$ Deep Blue, Kasparov


## Deterministic Games in Practice

- Checkers: Chinook ended 40-year-reign of human world champion Marion Tinsley in 1994. Used a precomputed endgame database defining perfect play for all positions involving 8 or fewer pieces on the board, a total of 444 billion positions.
- Chess: Deep Blue defeated human world champion Garry Kasparov in a six-game match in 1997. Deep Blue searches 200 million positions per second, uses very sophisticated evaluation, and undisclosed methods for extending some lines of search up to 40 ply.
- Othello: human champions refuse to compete against computers, who are too good.
- Go: The game of Go has been a fertile subject of artificial intelligence research for decades, culminating in 2017 with
AlphaGo Master winning three of three games against Ke Jie, who at the time continuously held the world No. 1 ranking for two years.


## Games of Chance

- What about games that involve chance, such as
- rolling dice
- picking a card
- Use three kinds of nodes:
- max nodes
- min nodes
- chance nodes



## Games of Chance


chance node with max children

$$
\operatorname{expectimax}(\mathrm{c})=\sum_{\mathrm{i}} \mathrm{P}\left(\mathrm{~d}_{\mathrm{i}}\right) \max _{\mathrm{s} \text { in } \mathrm{S}\left(\mathrm{c}, \mathrm{~d}_{\mathrm{i}}\right)}^{\operatorname{macked}-u p-v a l u e(\mathrm{~s}))}
$$

expectimin $\left(\mathrm{c}^{\prime}\right)=\sum \mathrm{P}\left(\mathrm{d}_{\mathrm{i}}\right) \min ($ backed-up-value(s))
s in $\mathrm{S}\left(\mathrm{c}, \mathrm{d}_{\mathrm{i}}\right)$

## Example Tree with Chance



## Complexity

- Instead of $O\left(b^{m}\right)$, it is $O\left(b^{m} n^{m}\right)$ where $n$ is the number of chance outcomes.
- Since the complexity is higher (both time and space), we cannot search as deeply.
- Pruning algorithms may be applied.


## Summary

- Games are fun to work on!
- They illustrate several important points about AI.
- Perfection is unattainable $\rightarrow$ must approximate.
- Game playing programs have shown the world what Al can do.

