## AI: Representation and Problem Solving

## Adversarial Search



Instructors: Tuomas Sandholm and Nihar Shah
Slide credits: CMU AI, http://ai.berkeley.edu

## Outline

History / Overview
Zero-Sum Games (Minimax)
Evaluation Functions
Search Efficiency ( $\alpha-\beta$ Pruning)
Games of Chance (Expectimax)


## Game Playing State-of-the-Art

Checkers:

- 1950: First computer player
- 1959: Samuel's self-taught program
- 1994: First computer world champion: Chinook ended 40-year-reign of human champion Marion Tinsley using complete 8-piece endgame
- 2007: Checkers solved! Endgame database of 39 trillion states

Chess:

- 1945-1960: Zuse, Wiener, Shannon, Turing, Newell \& Simon, McCarthy
- 1960s onward: gradual improvement under "standard model"
- 1997: special-purpose chess machine Deep Blue defeats human champion Gary Kasparov in a six-game match. Deep Blue examined 200M positions per second and extended some lines of search up to 40 ply. Current programs running on a PC rate > 3200 (vs 2870 for Magnus Carlsen)
Go:
- 1968: Zobrist's program plays legal Go, barely (b>300!)
- 2005-2014: Monte Carlo tree search enables rapid advances: current programs beat strong amateurs, and professionals with a 3-4 stone handicap
- 2015: AlphaGo from DeepMind beats best player Lee Sedol


## Poker:

- 1921: Borel introduces poker as the game theory benchmark
- 1950s: 3-card-deck tiny variant (Kuhn poker) solved by Kuhn, Nash, etc.
- 1950s-1970s: rule-based Als; not strong
- 1990s: ML-based Als; not strong
- 2000s-present: Game-theory-based Als
- 2008: Superhuman play in 2-player limit Texas hold’em [Bowling et al.]
- 2015: Near-optimal play in 2-player limit Texas hold'em [Bowling et al.]
- 2017: Superhuman AI Libratus for 2-player no-limit Texas hold'em [Brown \& Sandholm]
- 2019: Superhuman AI Pluribus for 2-player no-limit Texas hold’em [Brown \& Sandholm]



## Types of Games

Many different types of game!

Axes:

- Deterministic or stochastic?
- Perfect information (fully observable)?
- One, two, or more players?
- Turn-taking or simultaneous?

- Zero sum?

Want algorithms for calculating a contingent plan (a.k.a. strategy or policy) which recommends a move for every possible eventuality

## Zero-Sum Games



- Two-Player Zero-Sum Games
- Agents haye opposite utilities
- Pure competition:
- One maximizes, the other minimizes

- General Games
- Agents have independent utilities
- Cooperation, indifference, competition, shifting alliances, and more are all possible


## "Standard" Games

Standard games are deterministic, observable, two-player, turn-taking, zero-sum
Game formulation:

- Initial state: $\mathrm{s}_{0}$
- Players: Player(s) indicates whose move it is
- Actions: Actions(s) for player on move
- Transition model: Result(s,a)
- Terminal test: Terminal-Test(s)
- Terminal values: Utility(s,p) for player $p$
- Or just Utility(s) for player making the decision at root


Adversarial Search


Single-Agent Trees


## Minimax

States
Actions
Values


## Minimax

States
Actions
Values


## Minimax Code

```
def max_value(state):
    if state.is_leaf:
        return state.value
    # TODO Also handle depth limit
    best_value = -10000000
    for action in state.actions:
        next_state = state.result(action)
        next_value = min_value(next_state)
        if next_value > best_value:
            best_value = next_value
    return best_value
def min_value(state):
```


## Poll 1 (+ worksheet Poll 2 and 3 for Q1a/b)

What is the minimax value at the root?
A) 2
B) 3
C) 6
D) 12
E) 14


## Poll 1

What is the minimax value at the root?


## Minimax Notation

```
def max_value(state):
    if state.is_leaf:
        return state.value
    # TODO Also handle depth limit
    best_value = -10000000
    for action in state.actions:
        next_state = state.result(action)
        next_value = min_value(next_state)
        if next_value > best_value:
            best_value = next_value
    return best_value
def min_value(state):
```



## Minimax Notation



$$
\begin{aligned}
V(s)= & \max _{a} V\left(s^{\prime}\right) \\
& \text { where } s^{\prime}=\operatorname{result}(s, a)
\end{aligned}
$$

$$
\begin{aligned}
& \hat{a}=\underset{a}{\operatorname{argmax}} V\left(s^{\prime}\right), \\
& \quad \text { where } s^{\prime}=\operatorname{result}(s, a)
\end{aligned}
$$

## Generic Game Tree Pseudocode

function minimax_decision( state )
return $\operatorname{argmax}_{\mathrm{a}}$ in state.actions value( state.result(a) )
function value( state)
if state.is_leaf
return state.value
if state.player is MAX
return $\max _{\mathrm{a}}$ in state.actions value( state.result(a) )
if state.player is MIN
return $\mathrm{min}_{\mathrm{a}}$ in state.actions value( state.result(a) )

## Generalized minimax (better name: backward induction)

What if the game is not zero-sum, or has multiple players?
Generalization of minimax:

- Terminals have utility tuples
- Node values are also utility tuples
- Each player maximizes its own component
- Can give rise to cooperation and competition dynamically...

omㅇ


## Minimax Efficiency

How efficient is minimax?

- Just like (exhaustive) DFS
- Time: O(b ${ }^{m}$ )
- Space: O(bm)

Example: For chess, $\mathrm{b} \approx 35, \mathrm{~m} \approx 100$

- Exact solution is completely infeasible
- Humans can't do this either, so how do we play chess?
- Bounded rationality - Herbert Simon



## Resource Limits



## Resource Limits

Problem: In realistic games, cannot search to leaves!
Solution 1: Bounded lookahead

- Search only to a preset depth limit or horizon
- Use an evaluation function for non-terminal positions

Guarantee of optimal play is gone
More plies make a BIG difference

Example:

- Suppose we have 100 seconds, can explore 10K nodes / sec
- So can check 1 M nodes per move
- For chess, b=~35 so reaches about depth 4 - not so good



## Depth Matters

Evaluation functions are always imperfect

Deeper search => better play (usually)

Or, deeper search gives same quality of
 play with a less accurate evaluation function

An important example of the tradeoff between complexity of features and complexity of computation


## Evaluation Functions



## Evaluation Functions

Evaluation functions score non-terminals in depth-limited search


Ideal function: returns the actual minimax value of the position In practice: typically weighted linear sum of features:

- $\operatorname{EVAL}(\mathrm{s})=w_{1} f_{1}(s)+w_{2} f_{2}(s)+\ldots .+w_{n} f_{n}(s)$
- E.g., $w_{1}=9, f_{1}(s)=$ (num white queens - num black queens), etc.


## Evaluation for Pacman



Game Tree Pruning


## Minimax Example



## Alpha-Beta Example

$\boldsymbol{\alpha}=$ best option so far from any
MAX node on this path


The order of generation matters: more pruning is possible if good moves come first

## Alpha-Beta Implementation

$\alpha$ : MAX's best option on path to root
$\beta$ : MIN's best option on path to root
def max-value(state, $\alpha, \beta$ ):
initialize $v=-\infty$
for each successor of state:
$v=\max (v$, value(successor, $\alpha, \beta)$ )
if $v \geq \beta$
return $v$
$\alpha=\max (\alpha, v)$
return v

```
def min-value(state, \alpha, \beta):
    initialize v = +\infty
    for each successor of state:
        v=min(v, value(successor, \alpha, \beta))
        if v
        return v
        \beta=min}(\beta,v
    return v
```


## On your own

Which branches are pruned?
(Left to right traversal)
(Select all that apply)


## Poll 4

Which branches are pruned?
(Left to right traversal)
A) e, I
B) g, I
C) $\mathrm{g}, \mathrm{k}, \mathrm{l}$
D) $g, n$


Poll 4


## Alpha-Beta Code



## Alpha-Beta Code



## Alpha-Beta Pruning Properties

Theorem: This pruning has no effect on minimax value computed for the root!

Good child ordering improves effectiveness of pruning

- Iterative deepening helps with this

With "perfect ordering":

- Time complexity drops to $O\left(b^{m / 2}\right)$
- Doubles solvable depth!
- 1M nodes/move => depth=8, respectable


This is a simple example of metareasoning (computing about what to compute)

## Modeling Assumptions

Know your opponent


## Modeling Assumptions

Know your opponent


## Modeling Assumptions

Dangerous Pessimism
Assuming the worst case when it's not likely


Dangerous Optimism
Assuming chance when the world is adversarial


## Chance outcomes in trees



Tictactoe, chess
Minimax


Tetris, investing
Expectimax


Backgammon, Monopoly Expectiminimax

## Probabilities



## Probabilities

A random variable represents an event whose outcome is unknown

A probability distribution is an assignment of weights to outcomes

## Example: Traffic on freeway

- Random variable: $\mathrm{T}=$ whether there's traffic
- Outcomes: T in \{none, light, heavy\}
- Distribution:

$$
\mathrm{P}(\mathrm{~T}=\text { none })=0.25, \mathrm{P}(\mathrm{~T}=\text { light })=0.50, \mathrm{P}(\mathrm{~T}=\text { heavy })=0.25
$$

Probabilities over all possible outcomes sum to one

0.25

0.50

## Expected Value

Expected value of a function of a random variable:
Average the values of each outcome, weighted by the probability of that outcome

Example: How long to get to the airport?


## Expectations



Max node notation

$$
\begin{aligned}
V(s)= & \max _{a} V\left(s^{\prime}\right) \\
& \text { where } s^{\prime}=\operatorname{result}(s, a)
\end{aligned}
$$



Chance node notation
$V(s)=$

## Expectations



Max node notation

$$
\begin{aligned}
V(s)= & \max _{a} V\left(s^{\prime}\right) \\
& \text { where } s^{\prime}=\operatorname{result}(s, a)
\end{aligned}
$$

Chance node notation
$V(s)=\sum_{s^{\prime}} P\left(s^{\prime}\right) V\left(s^{\prime}\right)$

## On your own...

Expectimax tree search: Which action do we choose?

A: Left
B: Center
C: Right
D: Eight


## On your own...

Expectimax tree search: Which action do we choose?

A: Left
B: Center
C: Right
D: Eight


## Expectimax Code

function value( state )
if state.is_leaf
return state.value
if state.player is MAX
return $\max _{\mathrm{a}}$ in state.actions value( state.result(a) )
if state.player is MIN return $\min _{\mathrm{a}}$ in state.actions value( state.result(a) )
if state.player is CHANCE
return $\operatorname{sum}_{s}$ in state.next_states $P(s)$ * value (s)

## Expectimax Pruning?



## Modeling Assumptions



|  | Minimax <br> Ghost | Random <br> Ghost |
| :---: | :---: | :---: |
| Minimax <br> Pacman |  |  |
| Expectimax <br> Pacman |  |  |

Results from playing 5 games

## Summary

Games require decisions when optimality is impossible

- Bounded-depth search and approximate evaluation functions

Games force efficient use of computation

- E.g., alpha-beta pruning

Game playing has produced important research ideas

- Reinforcement learning (checkers)
- Iterative deepening (chess)
- Monte Carlo tree search (Go)
- Solution methods for partial-information games, e.g., in economics (poker)

Lots to do!

- E.g., video games present greater challenges: $b=10^{500},|S|=10^{4000}, m=10,000$
- See Prof. Sandholm course CS 15-888 Computational Game Solving


## Preview: MDP/Reinforcement Learning Notation



$$
V(s)=\max _{a} \sum_{s^{\prime}} P\left(s^{\prime}\right) V\left(s^{\prime}\right)
$$

## Preview: MDP/Reinforcement Learning Notation



## Preview: MDP/Reinforcement Learning Notation



## Why Expectimax?

Pretty great model for an agent in the world
Choose the action that has the: highest expected value


## Bonus Question

Let's say you know that your opponent is actually running a depth 1 minimax, using the result $80 \%$ of the time, and moving randomly otherwise
Question: What tree search should you use?
A: Minimax
B: Expectimax
C: Something completely different

