


CSE 573: Artificial Intelligence

Adversarial Search

Dan Weld



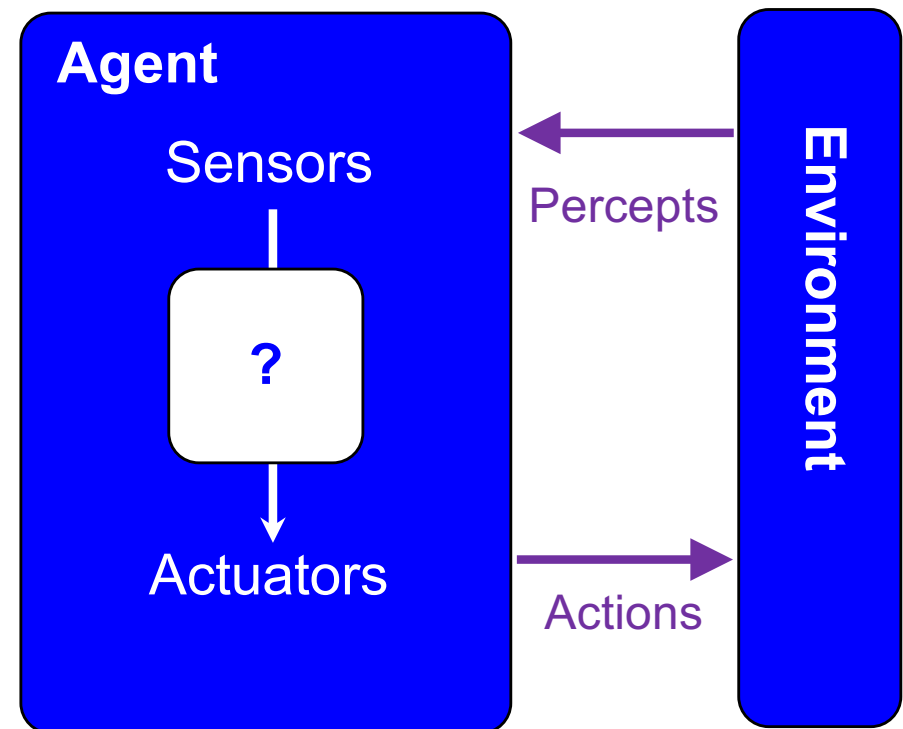
Based on slides from

Dan Klein, Stuart Russell, Pieter Abbeel, Andrew Moore and Luke Zettlemoyer

(best illustrations from ai.berkeley.edu)

Types of Environments

- Fully observable *vs.* partially observable
- Single agent *vs.* *multi-agent*
- Deterministic *vs.* *stochastic*
- Episodic *vs.* sequential
- Discrete *vs.* continuous





Types of Games

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

Number of Players? 1, 2, ...?

Zero-Sum Games



■ Zero-Sum Games

- Agents have opposite utilities (values on outcomes)
- Lets us think of a single value that one maximizes and the other minimizes
- Adversarial, pure competition

■ General Games

- Agents have independent utilities (values on outcomes)
- Cooperation, indifference, competition, & more are possible
- More later on non-zero-sum games


Minimax Values

States Under Agent's Control:

$$V(s) = \max_{s' \in \text{successors}(s)} V(s')$$

States Under Opponent's Control:

$$V(s') = \min_{s \in \text{successors}(s')} V(s)$$



Terminal States:

$$V(s) = \text{known}$$

Minimax Implementation

Need **Base case** for recursion

```
def max-value(state):  
    if leaf?(state), return U(state)  
    initialize v = -∞  
    for each c in children(state)  
        v = max(v, min-value(c))  
    return v
```

```
def min-value(state):  
    if leaf?(state), return U(state)  
    initialize v = +∞  
    for each c in children(state)  
        v = min(v, max-value(c))  
    return v
```


$$V(s) = \max_{s' \in \text{successors}(s)} V(s')$$

$$V(s') = \min_{s \in \text{successors}(s')} V(s)$$

Alpha-Beta Quiz

Max:

Min:




Search depth-first
Left to right
Order is important

Do all nodes matter?

Stochastic Single-Player


- What if we don't know what the result of an action will be? E.g.,
 - In solitaire, shuffle is unknown
 - In minesweeper, mine locations
- Can do **expectimax search**
 - Chance nodes, like actions except the environment controls the action chosen
 - Max nodes as before
 - Chance nodes take average (expectation) of value of children



ExpectiMax Search

In ExpectiMax search, we have a probabilistic model of how the opponent (or environment) will behave in any state

- Model could be a simple uniform distribution (roll a die)... or more complex
- We have a node for every outcome out of our control: opponent or environment



For now, assume \forall states we magically have a distribution to assign probabilities to enemy-actions / environment outcomes

Expectimax Pseudocode

```
def value(s)
```

```
  if s is a max node return maxValue(s)
```

```
  if s is an exp node return expValue(s)
```

```
  if s is a terminal node return evaluation(s)
```

```
def maxValue(s)
```

```
  values = [value(s') for s' in successors(s)]
```


```
  return max(values)
```

```
def expValue(s)
```

```
  values = [value(s') for s' in successors(s)]
```

```
  weights = [probability(s, s') for s' in successors(s)]
```

```
  return expectation(values, weights)
```




ExpectiMax for Pacman

- Note: that we've gotten away from thinking that the ghosts are trying to minimize pacman's score
- Instead, they are now a part of the environment
- Pacman has a belief (distribution) over how they will act
- **Quiz:** Can we see minimax as a special case of expectimax?
- **Quiz:** what would pacman's computation look like if we assumed that the ghosts were doing 1-ply minimax and taking the result 80% of the time, otherwise moving randomly?

Expectimax for Pacman


Results from playing 5 games

| | Minimizing Ghost | Random Ghost |
|-------------------|--------------------------------|-------------------------------|
| Minimax Pacman | Won 5/5 Avg. Score: 493 | Won 5/5 Avg. Score: 483 |
| Expectimax Pacman | Won 1/5 Avg. Score: -303 | Won 5/5 Avg. Score: 503 |



Pacman does depth 4 search with an eval function that avoids trouble
Minimizing ghost does depth 2 search with an eval function that seeks Pacman


ExpectiMax Pruning?



- Not easy
 - exact: need bounds on possible values
 - approximate: sample high-probability branches


ExpectiMax Evaluation

- Evaluation functions quickly return an estimate for a node's true value (which value, expectimax or minimax?)
- For minimax, evaluation function scale doesn't matter
 - We just want better states to have higher evaluations (get the ordering right)
 - We call this **insensitivity to monotonic transformations**
- For expectimax, we need *magnitudes* to be meaningful



Mixed Layer Types

- E.g. Backgammon
- Expecti-Mini-Max
 - Environment is an extra player that moves after each agent
 - Chance nodes take expectations, otherwise like minimax



if *state* is a MAX node then

return the highest EXPECTIMINIMAX-VALUE of SUCCESSORS(*state*)

if *state* is a MIN node then


return the lowest EXPECTIMINIMAX-VALUE of SUCCESSORS(*state*)

if *state* is a chance node then

return average of EXPECTIMINIMAX-VALUE of SUCCESSORS(*state*)

Stochastic Two-Player


- Dice rolls increase b : 21 possible rolls with 2 dice
 - Backgammon \approx 20 legal moves
 - Depth 4 = $20 \times (21 \times 20)^3 = 1.2 \times 10^9$
- As depth increases, probability of reaching a given node shrinks
 - So value of lookahead is diminished
 - So limiting depth is less damaging
 - But pruning is less possible...
- TDGammon used depth-2 search + very good eval function
 - Learned via NN & reinforcement learning
 - **World-champion level play (1992, Gerald Tesauro)**



Multi-player Non-Zero-Sum Games

Similar to minimax:

- Utilities are now tuples
- Each player maximizes their own entry at each node
- Propagate (or back up) nodes from children
- Can give rise to cooperation and competition dynamically...



In this example... three agents