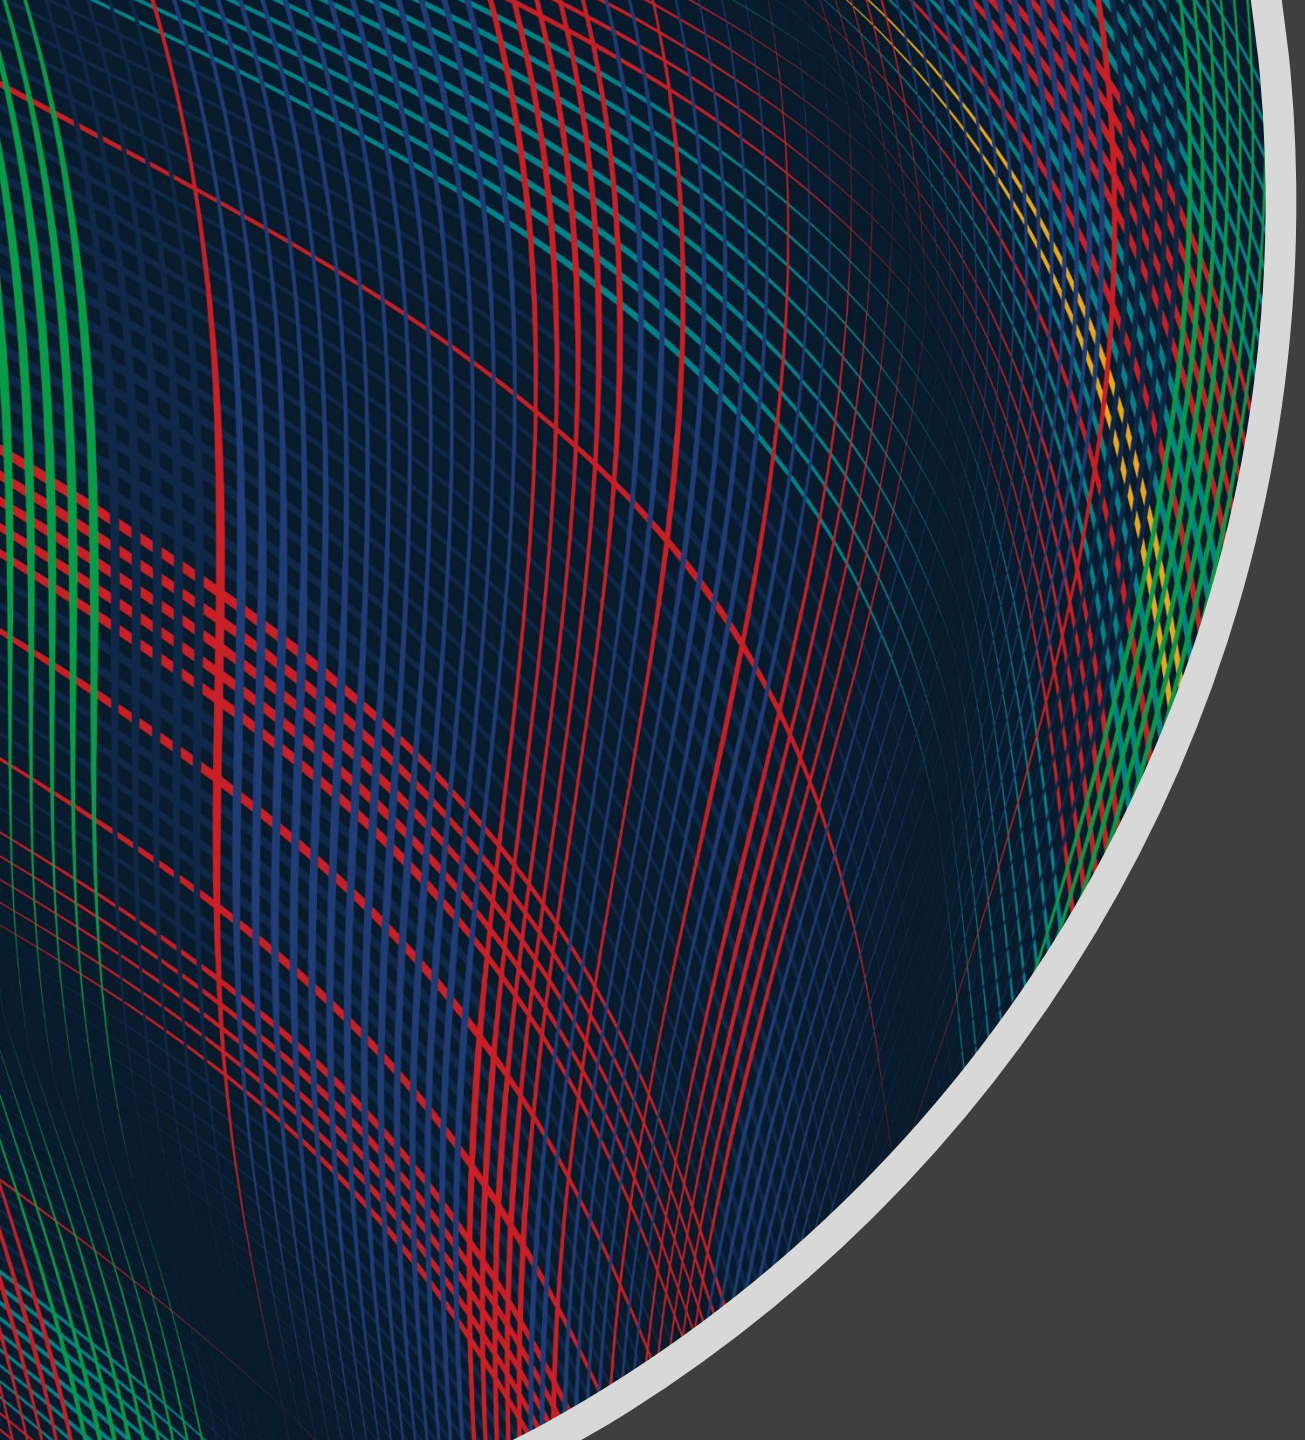

07-280 AI & ML I
Adversarial Search

Instructors:

Aviral Kumar & Nihar Shah

Minimax

Why the focus on games?

States

Actions

Values

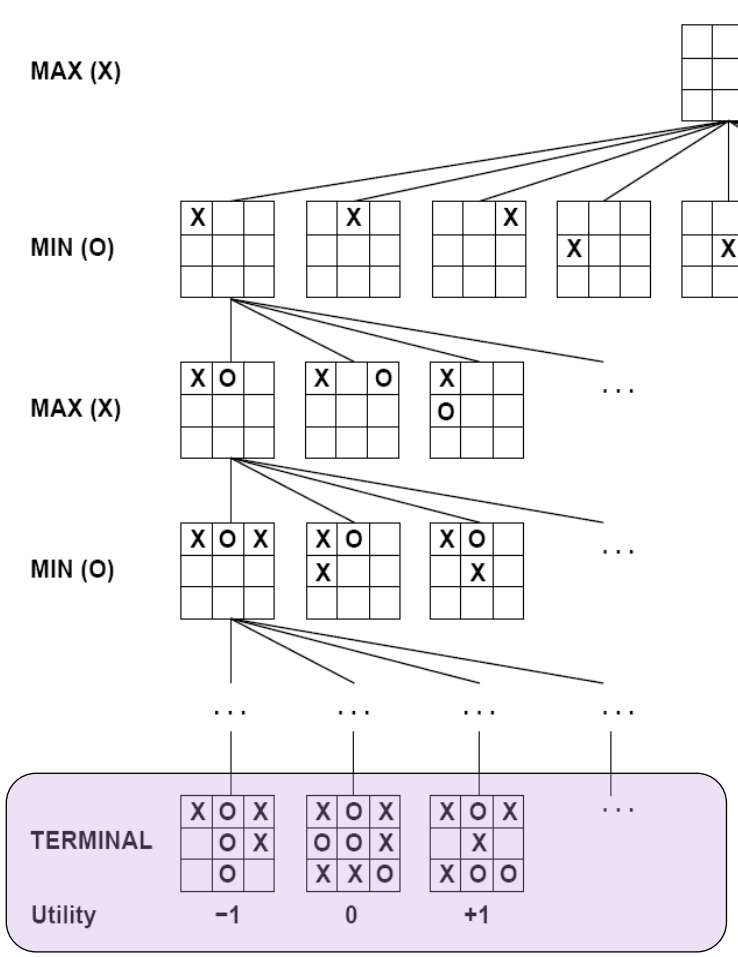

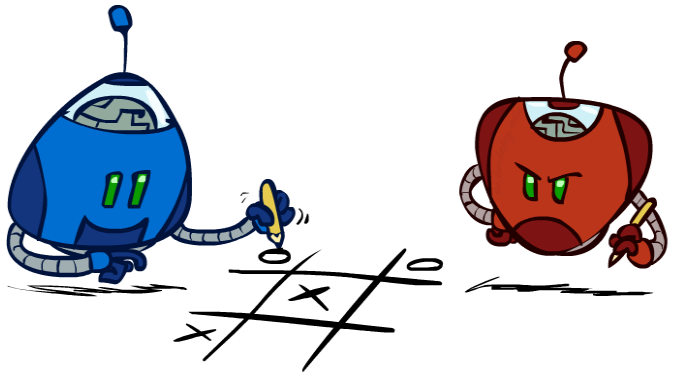

Image credit: Ketrina Yim, UCB CS 188

Intelligence and Uncertainty

Uncertainty can have lots of sources, including anything we attribute to random chance

- Hidden information
  - Cards in another player's hand
- Noise
  - Sensor noise
- Way to complicated to model
  - Leaves blowing in the wind
- Infinite number of possible configurations
- More possibilities than any computer can compute in a reasonable time
  - Tic-tac-toe → Checkers → Chess

Outline

Pre-reading

- History
- Standard/Zero-Sum Games
- Minimax
- Expected Value
- Expectimax

Evaluation Functions

Pruning

Modeling Assumptions

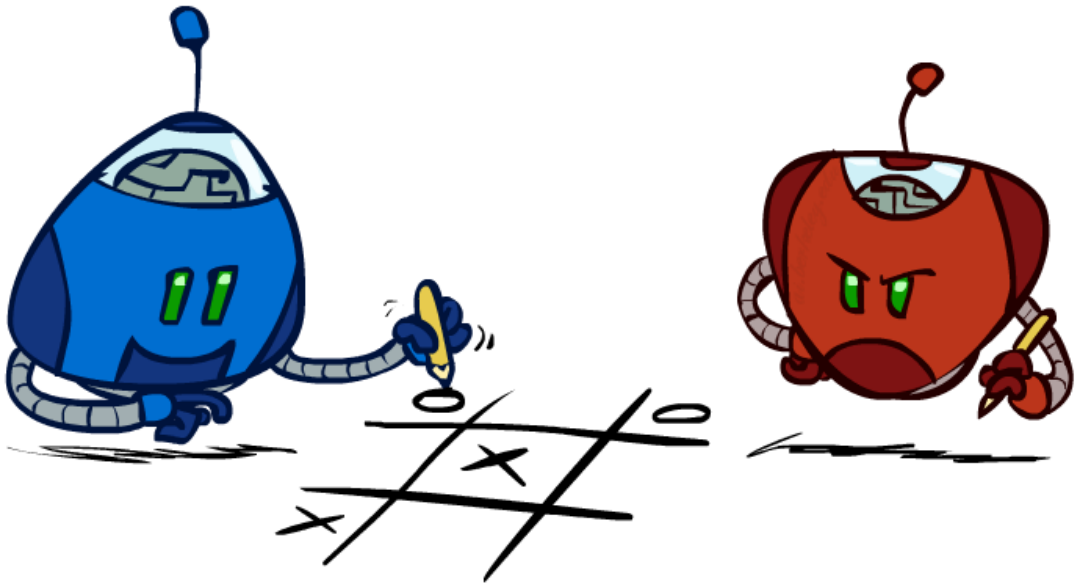

Image credit: Ketrina Yim, UCB CS 188 and well, Pat Virtue

Pre-reading

Zero-Sum Games

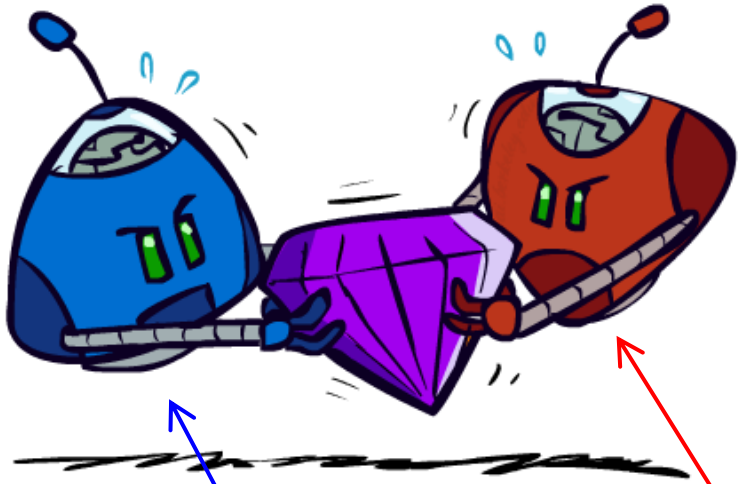

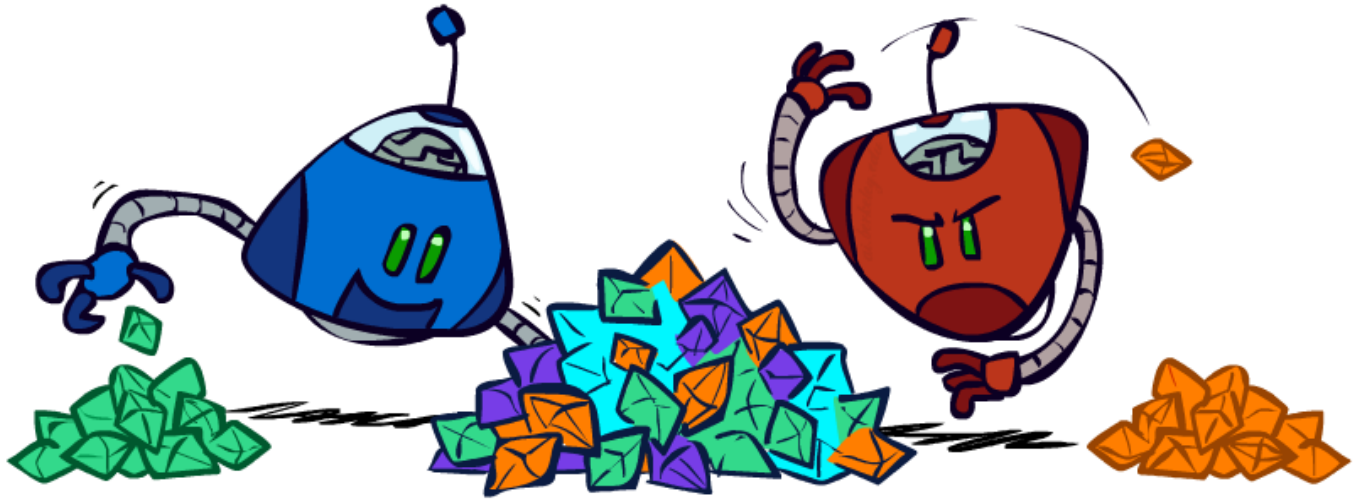

- Zero-Sum Games
  - Agents have ***opposite*** utilities
  - Pure competition:
    - One ***maximizes***, the other ***minimizes***

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

Image/slide credit: Ketrina Yim, UCB CS 188

“Standard” Games

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

Game formulation:

- Initial state: 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

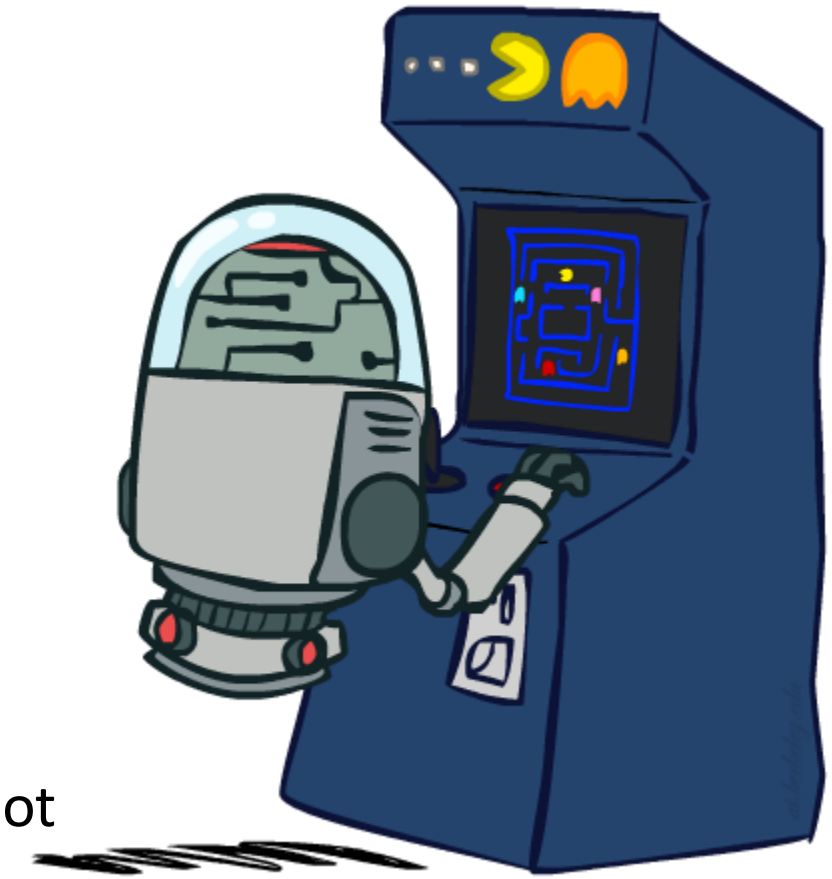

Image/slide credit: Ketrina Yim, UCB CS 188

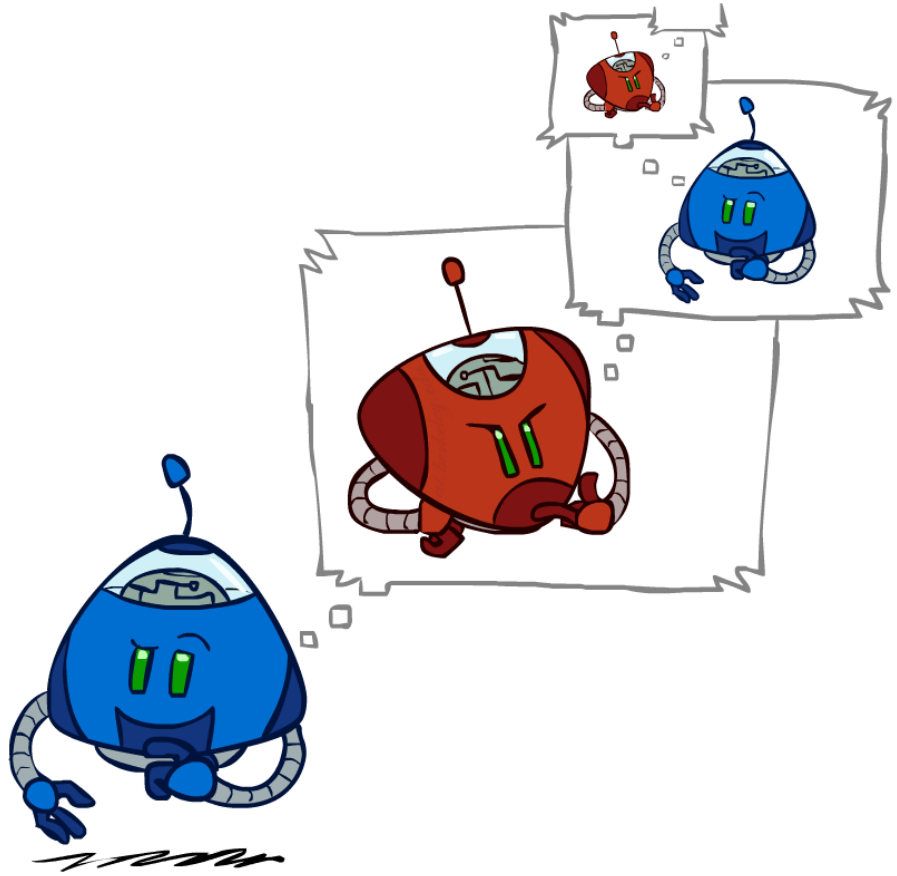

Pre-reading

Minimax Search

Image credit: Ketrina Yim, UCB CS 188

Single-Agent Trees

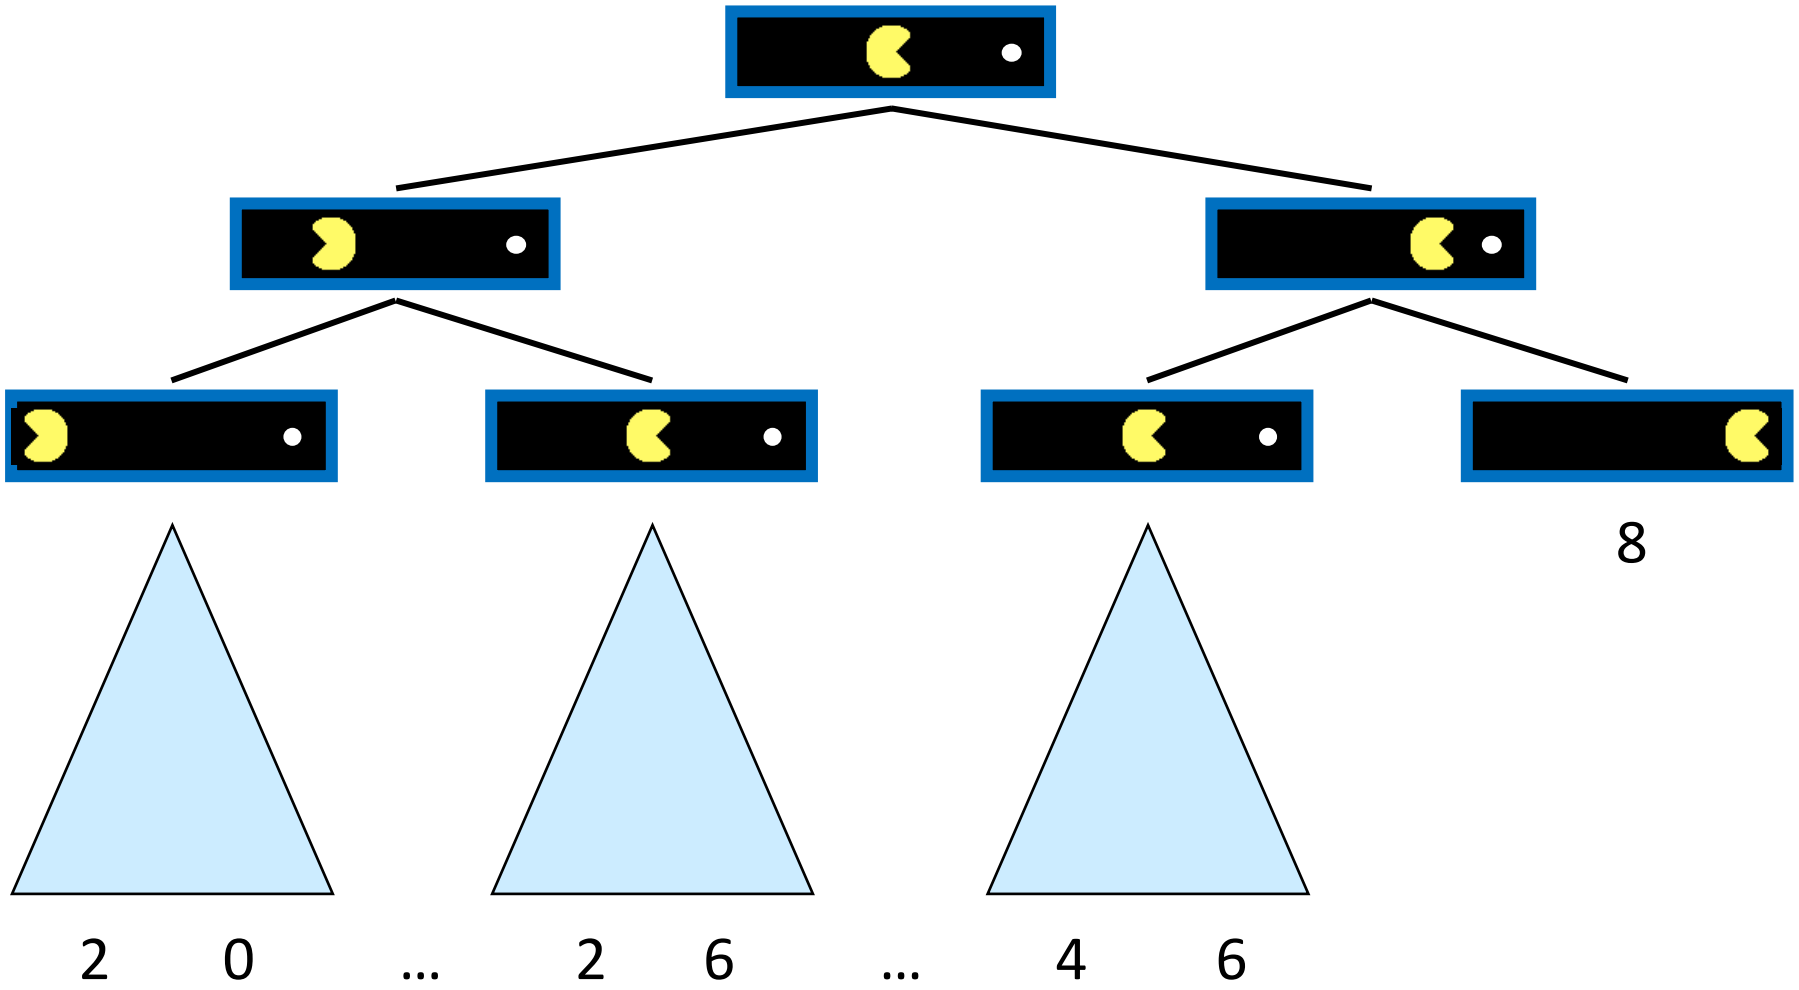

Slide credit: UCB CS 188

Minimax Example

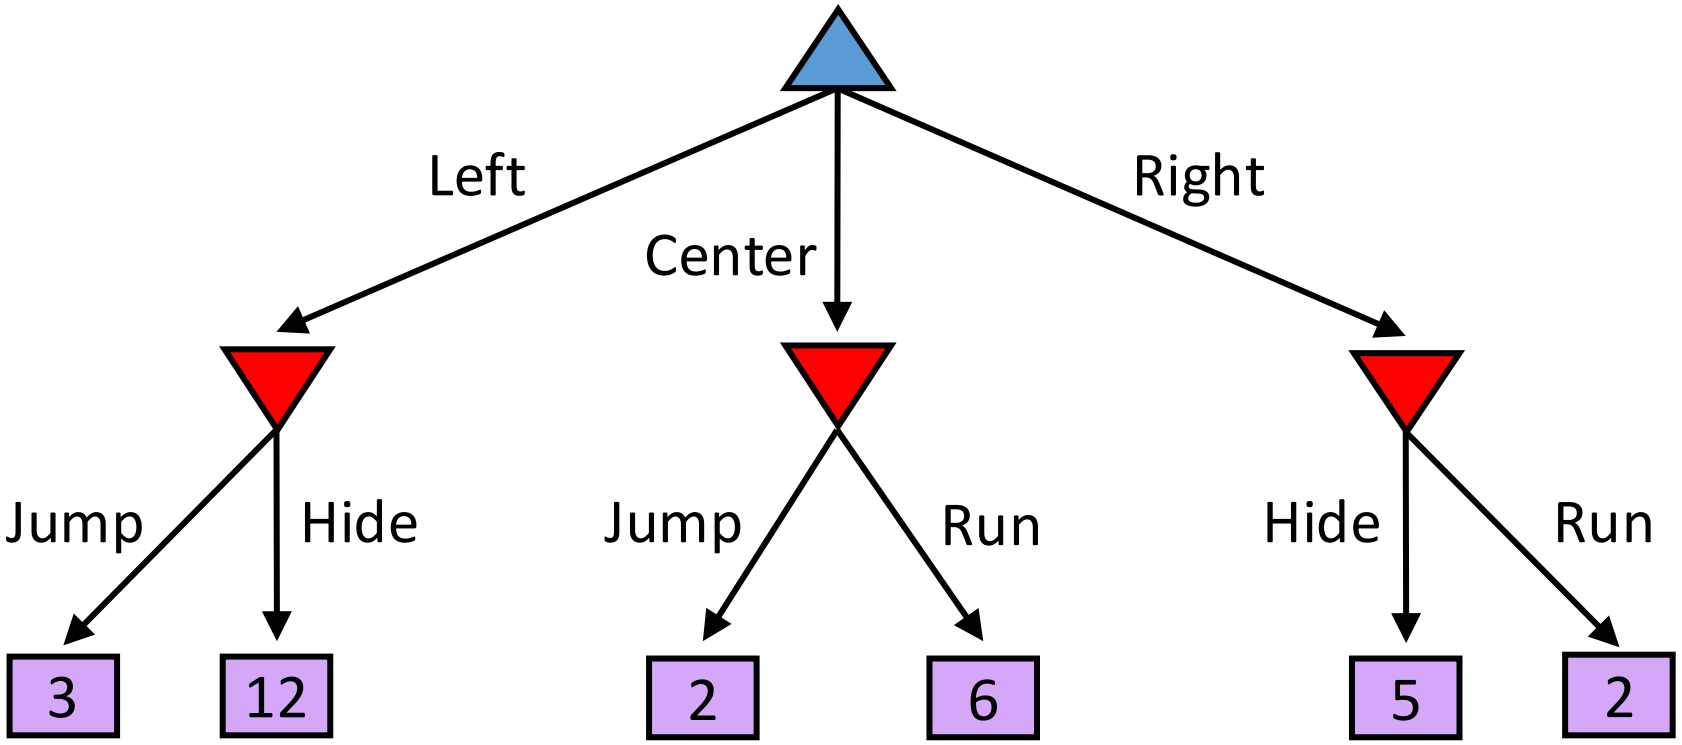

Generic Game Tree Pseudocode

```
function minimax_decision( state )
    return argmax 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 a in state.actions value( state.result(a) )

    if state.player is MIN
        return min a in state.actions value( state.result(a) )
```

Slide credit: UCB CS 188

Minimax Efficiency

How efficient is minimax?

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

Example: For chess, $b \approx 35$, $m \approx 100$

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

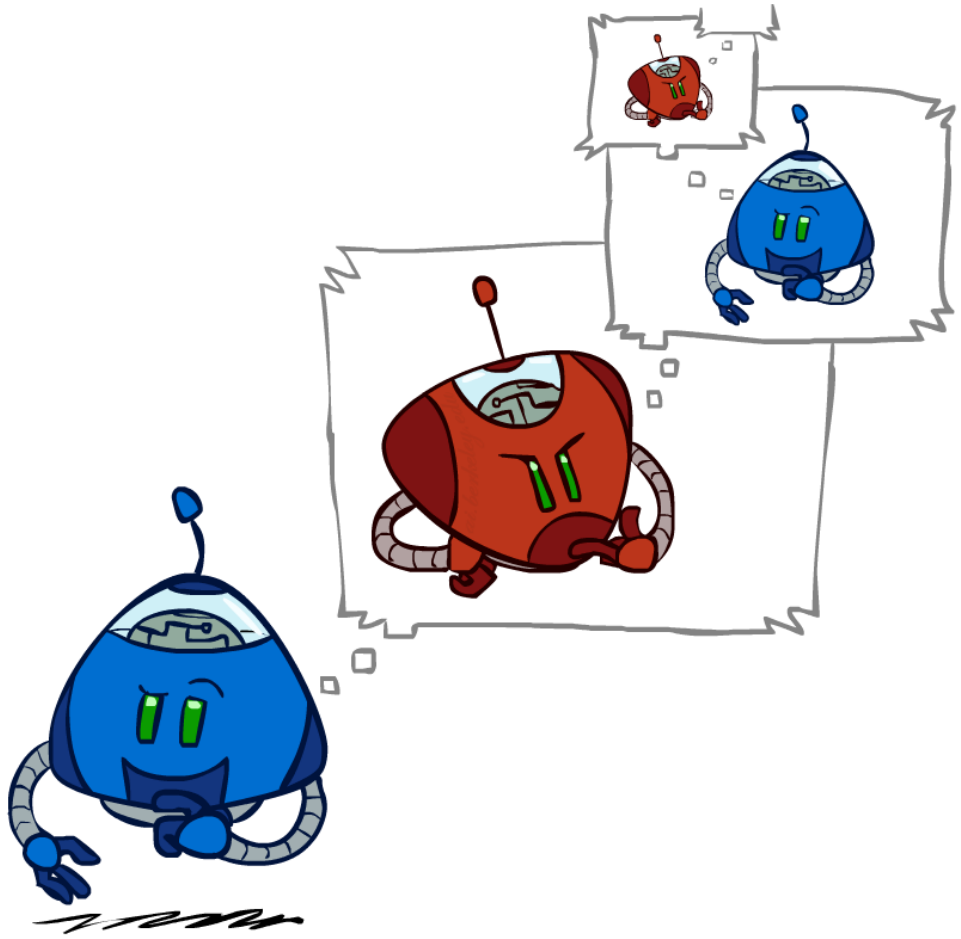

Image/slide credit: Ketrina Yim, UCB CS 188

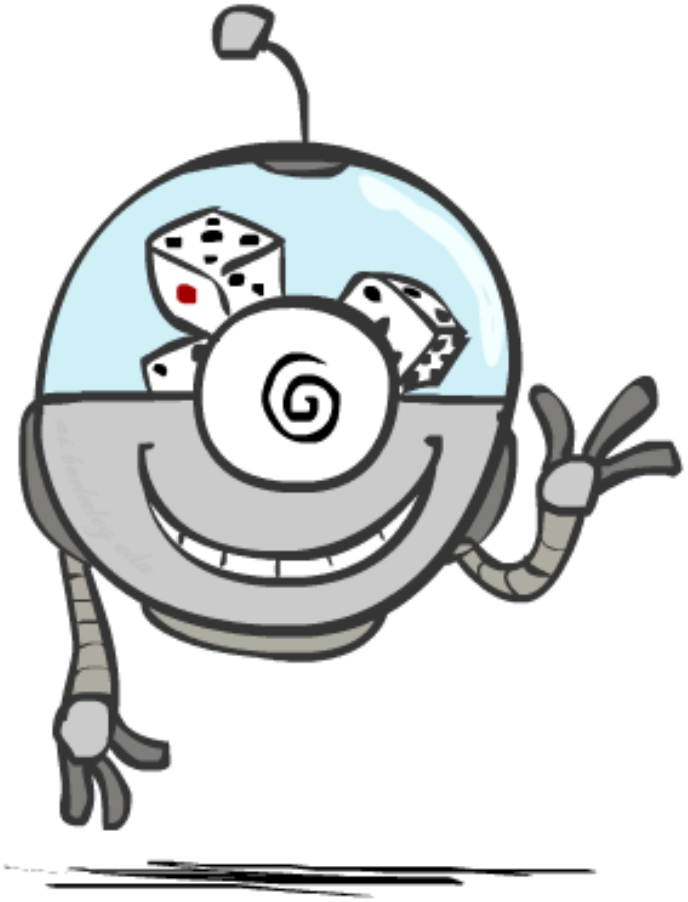

Pre-reading

Probability and Expected Value

Image credit: Ketrina Yim, UCB CS 188

Probabilities

A random variable represents an event whose outcome is unknown

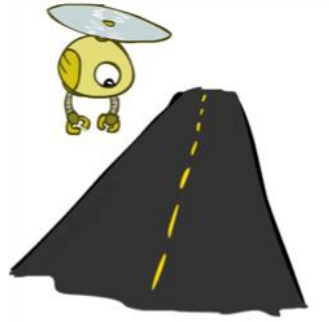

0.25

A probability distribution is an assignment of weights to outcomes

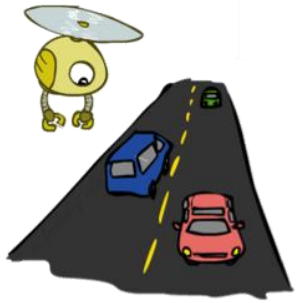

0.50

Example: Traffic on freeway

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

P(T=none) = 0.25, P(T=light) = 0.50, P(T=heavy) = 0.25

0.25

Probabilities over all possible outcomes sum to one

Image/slide credit: Ketrina Yim, UCB CS 188

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?

Time:	20 min	+	30 min	+	60 min
	x		x		x
Probability:	0.25		0.50		0.25

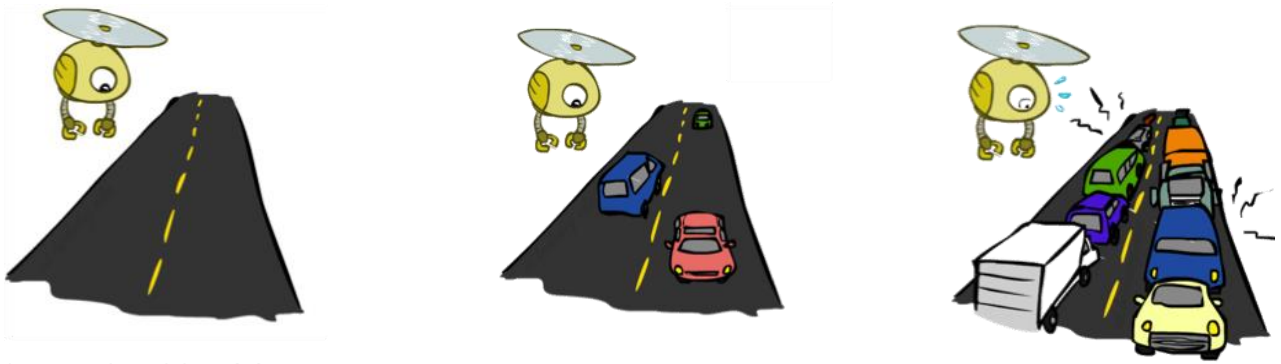

35 min

Image/slide credit: Ketrina Yim, UCB CS 188

Expected Value

Time:	20 min		30 min		60 min
	x	+	x	+	x
Probability:	0.25		0.50		0.25

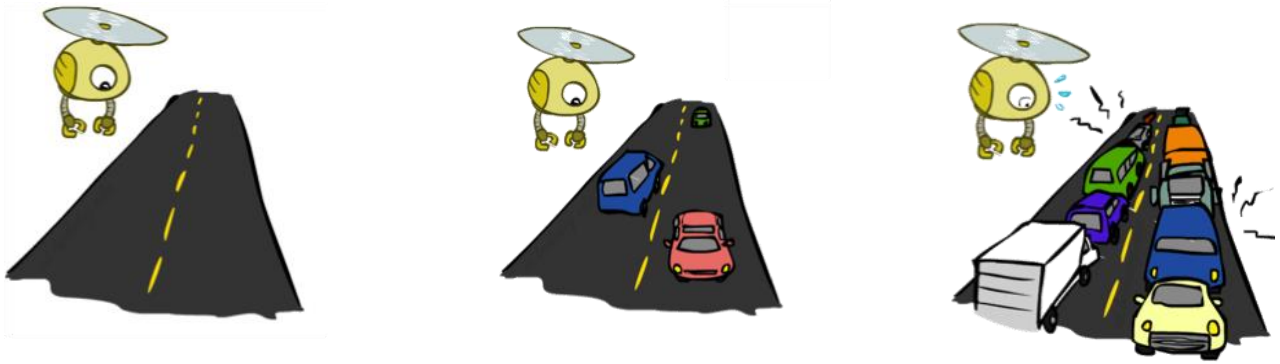

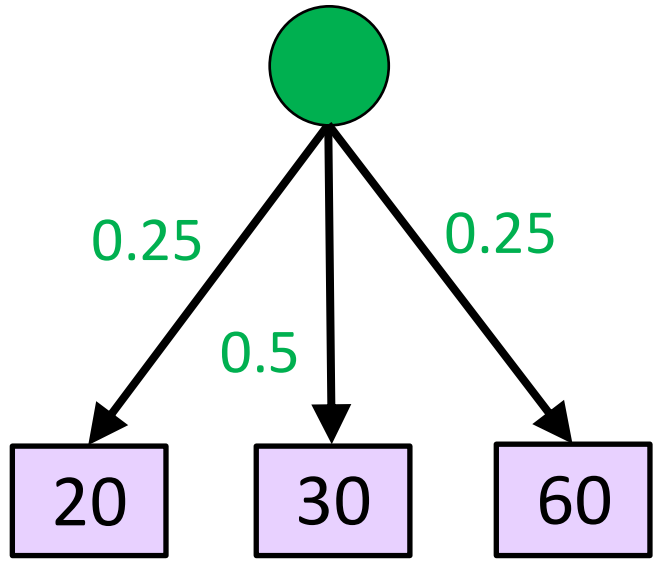

Max node notation

$$V(s) = \max_{a} V(s'),$$
$$\text{where } s' = result(s, a)$$

Chance node notation

$$V(s) = \sum_{s'} P(s')\, V(s')$$

Slide credit: UCB CS 188

Example

Expectimax tree search

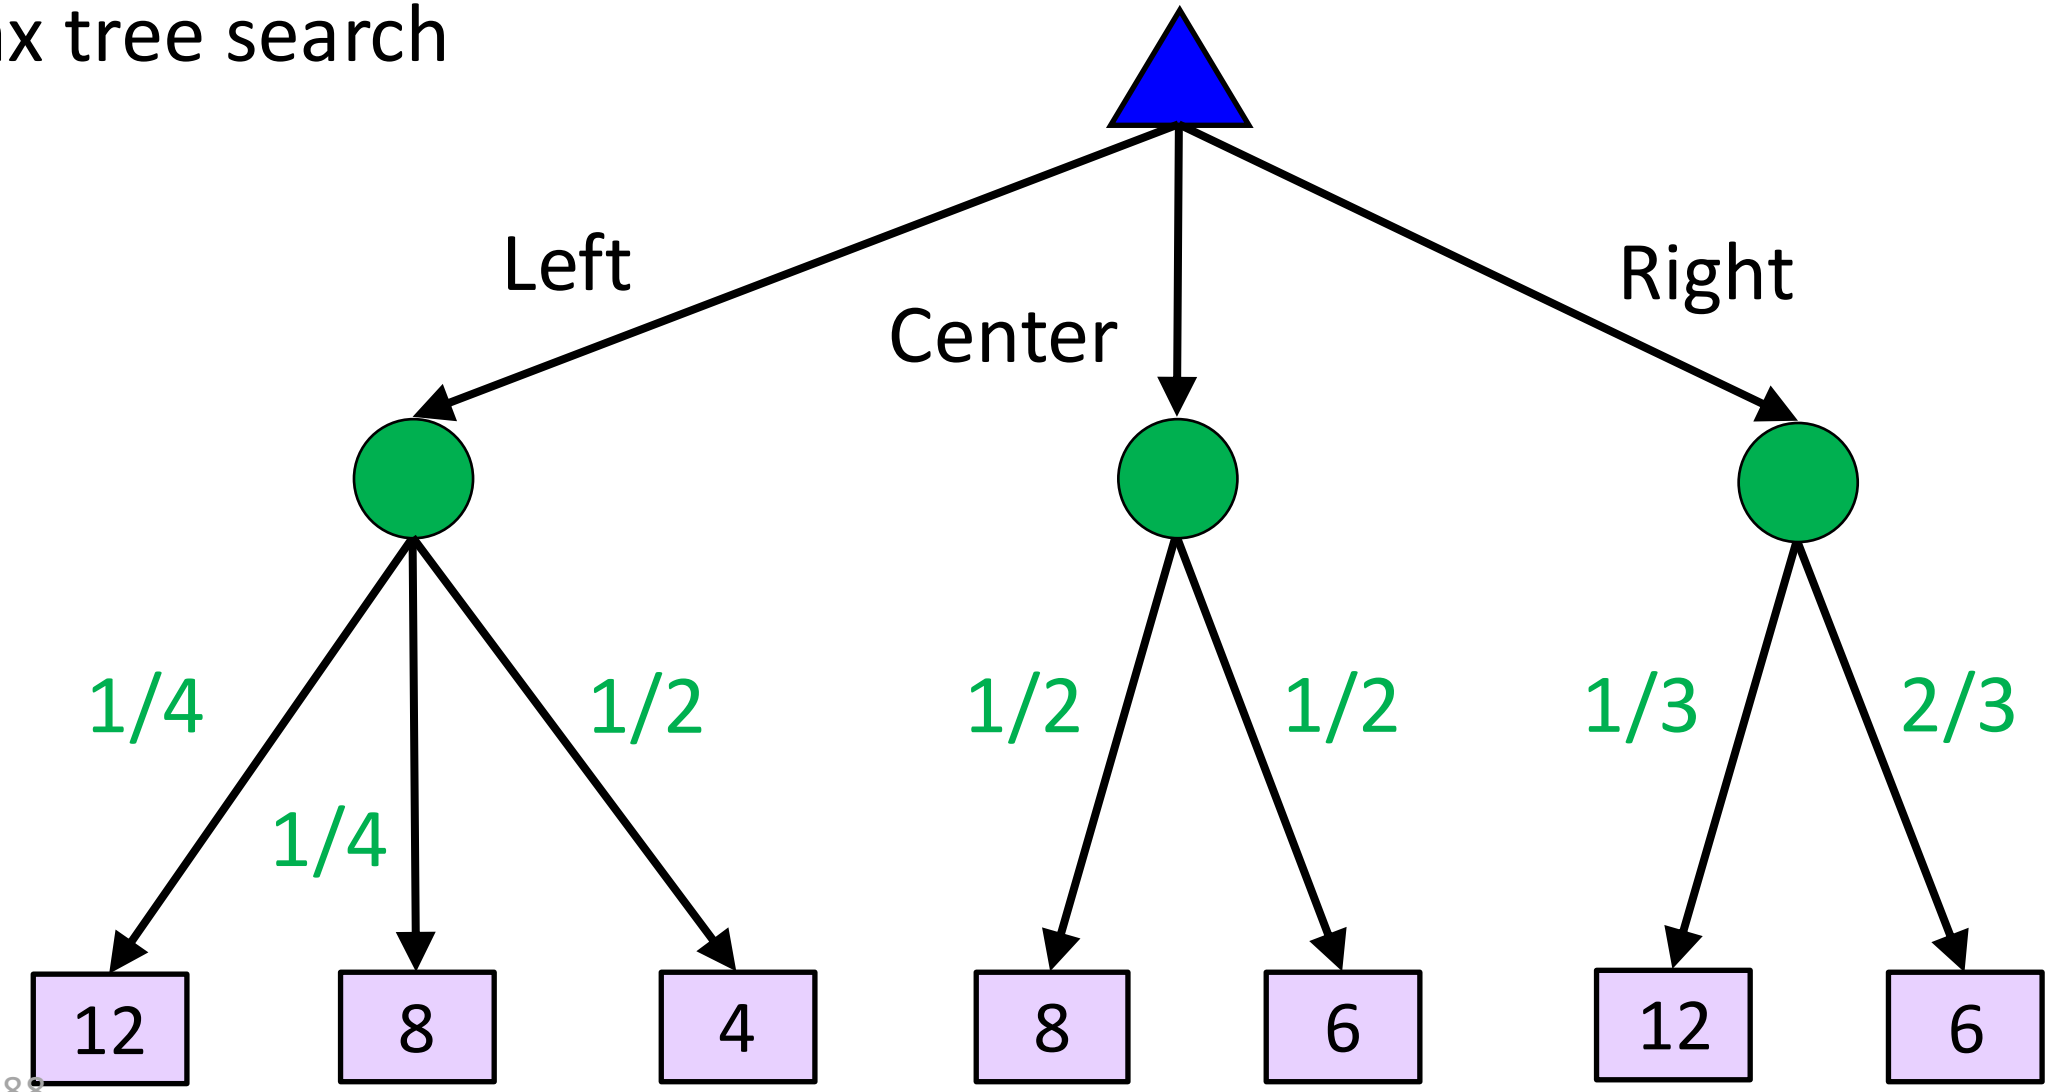

Updated: Generic Game Tree Pseudocode

```
function value( state )
    if state.is_leaf
        return state.value

    if state.player is MAX
        return max_{a in state.actions} value( state.result(a) )

    if state.player is MIN
        return min_{a in state.actions} value( state.result(a) )

    if state.player is CHANCE
        return sum_{s in state.next_states} P( s ) * value( s )
```

Slide credit: UCB CS 188

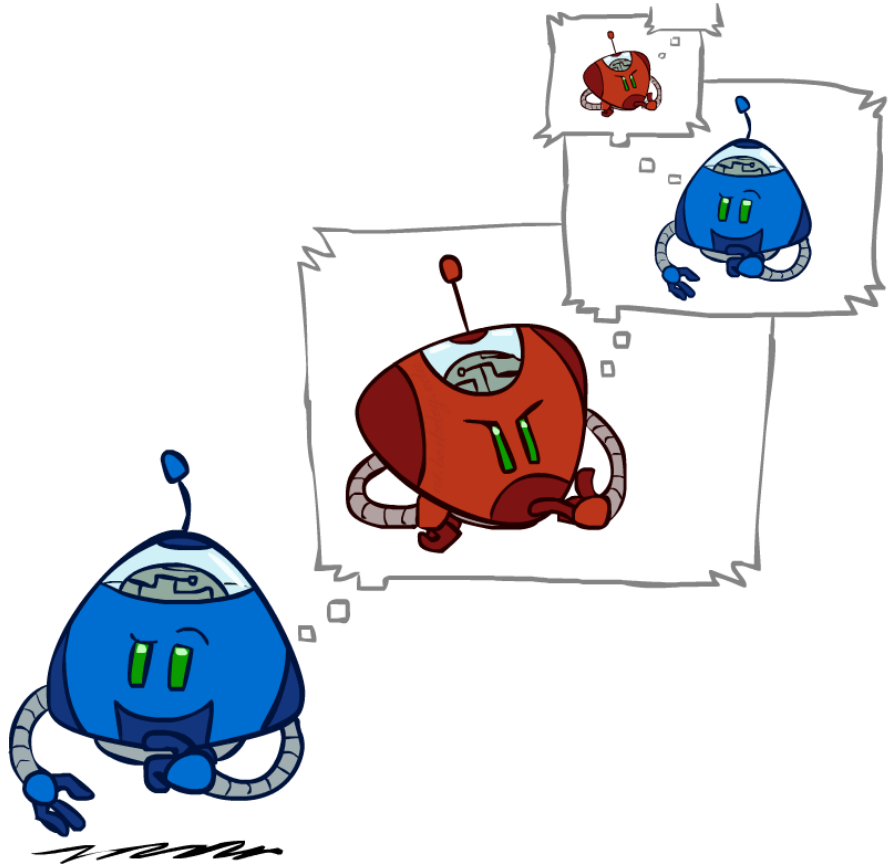

Adversarial Search

Image credit: Ketrina Yim, UCB CS 188

Minimax Code and Notation

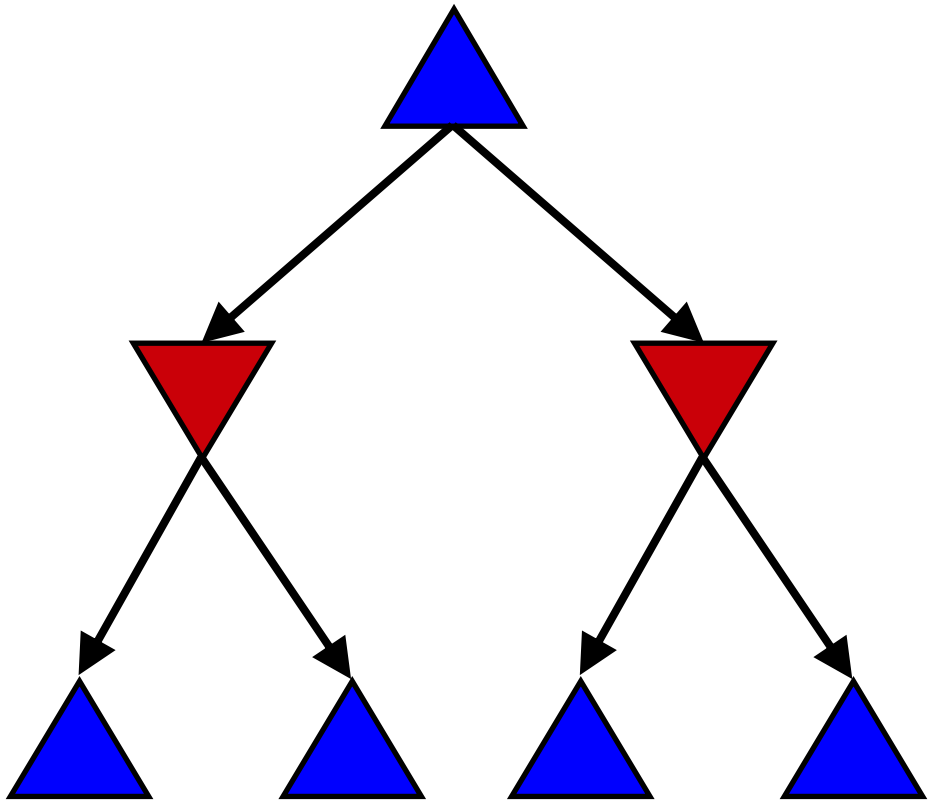

Minimax Notation

```
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):
```

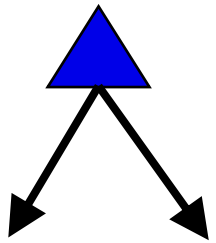

$$V(s) = \max_{a} V(s'),$$

where $s' = result(s, a)$

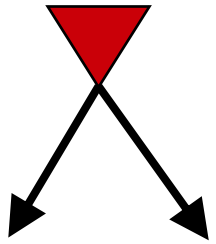

Minimax Notation

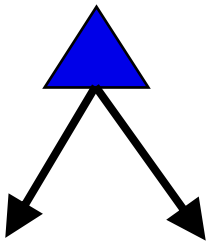

$$V(s) = \max_{a} V(s'),$$
$$\text{where } s' = result(s, a)$$

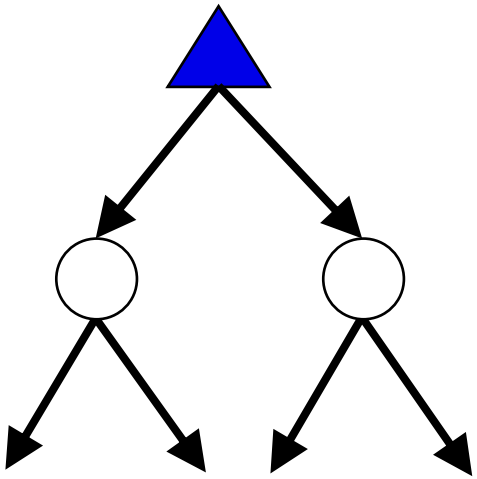

$$\hat{a} = \operatorname*{argmax}_{a} V(s'),$$
$$\text{where } s' = result(s, a)$$

Minimax Efficiency

How efficient is minimax?

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

Example: For chess, $b \approx 35$, $m \approx 100$

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

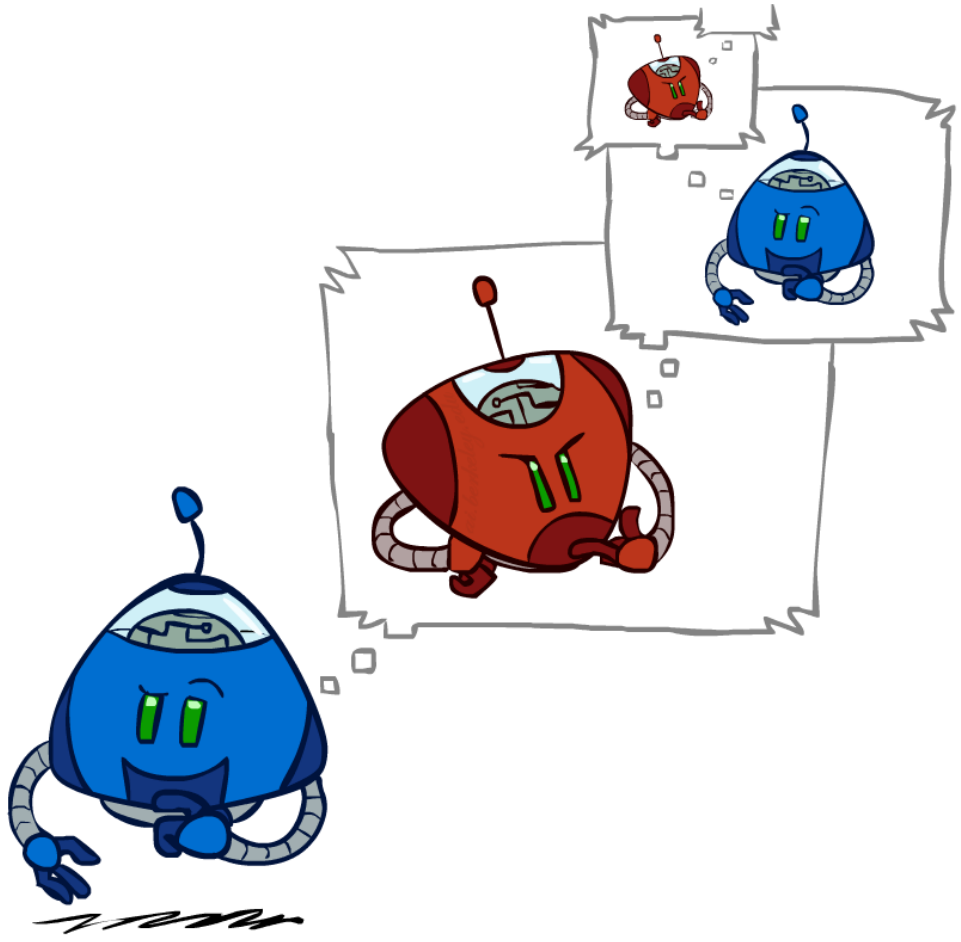

Slide credit: UCB CS 188

Resource Limits

Image/slide credit: Ketrina Yim, UCB CS 188

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 1M nodes per move
- For chess, b=~35 so reaches about depth 4 – not so good

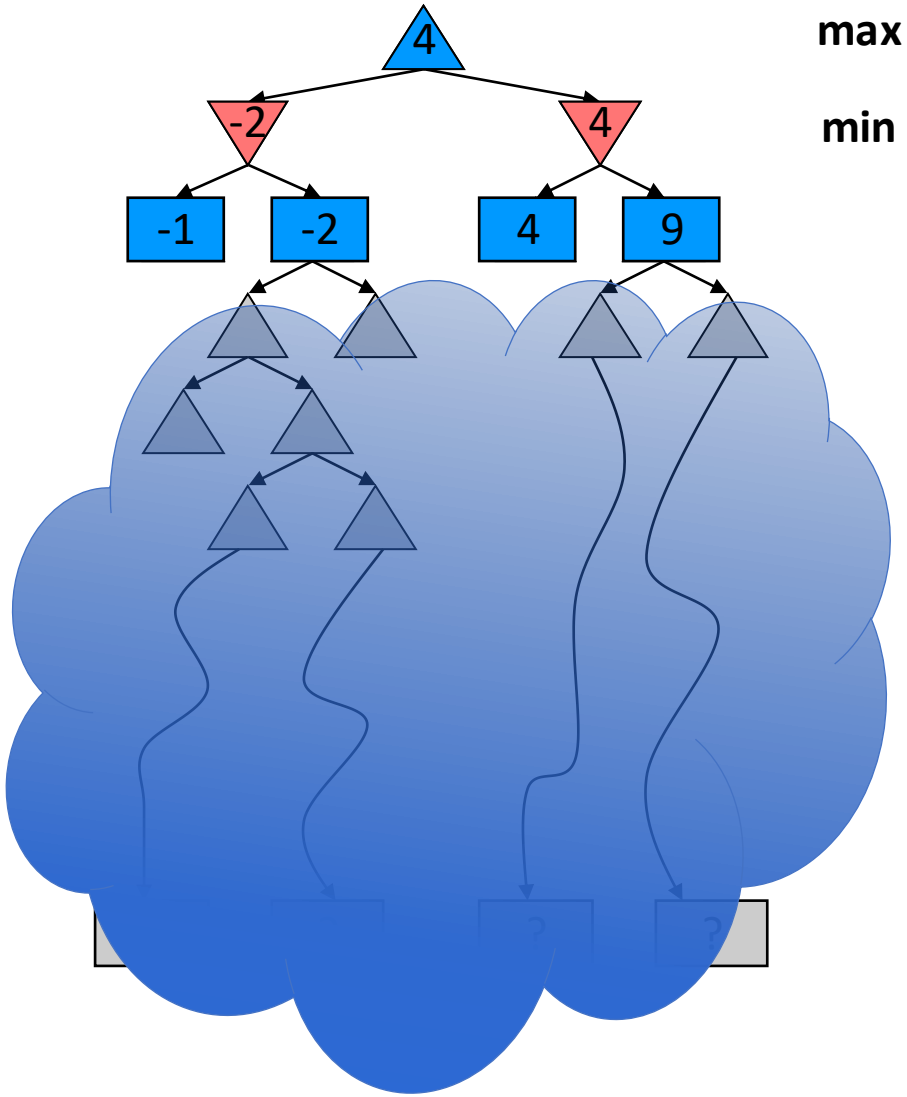

Slide credit: UCB CS 188

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

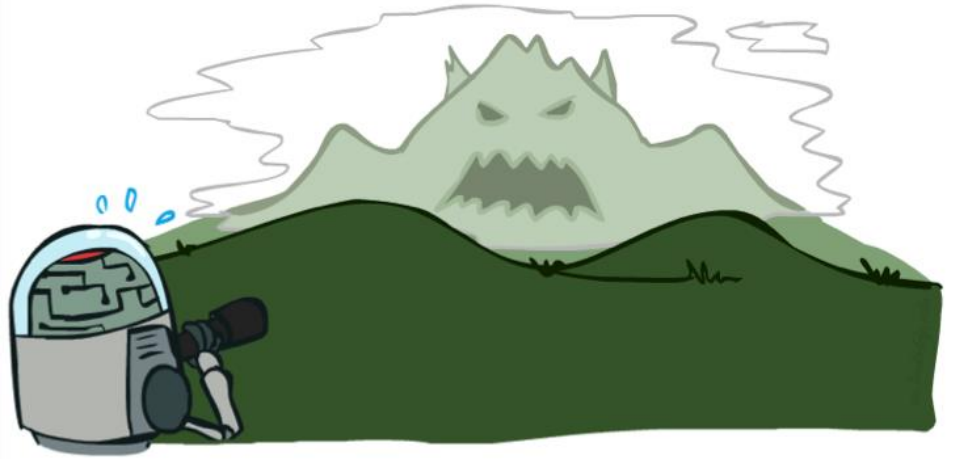

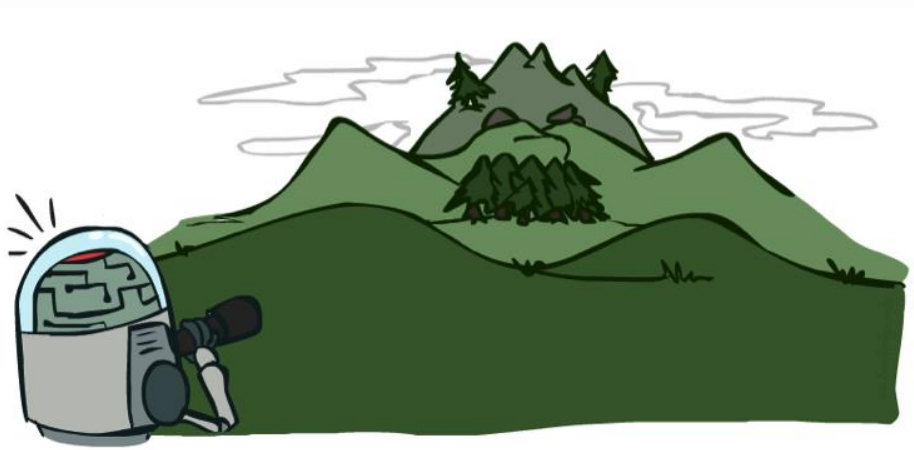

Image/slide credit: Ketrina Yim, UCB CS 188

[Demo: depth limited (L6D4, L6D5)]

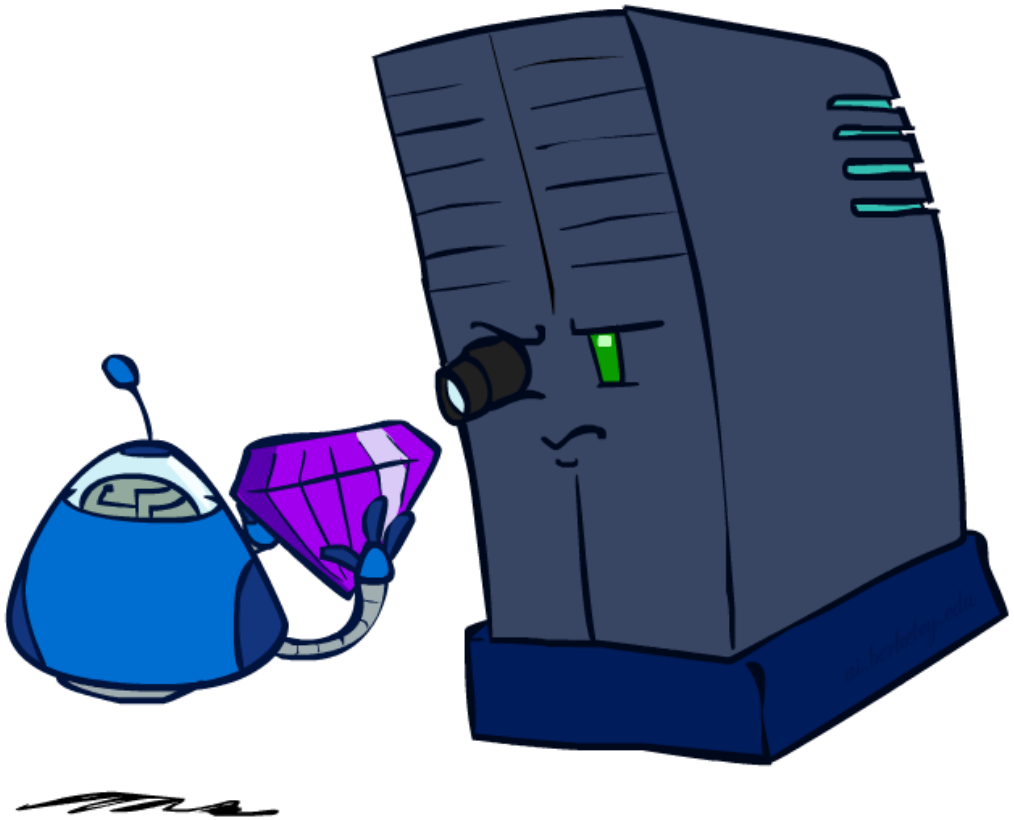

Evaluation Functions

Image credit: Ketrina Yim, UCB CS 188

Evaluation Functions

Evaluation functions score non-terminals in depth-limited search

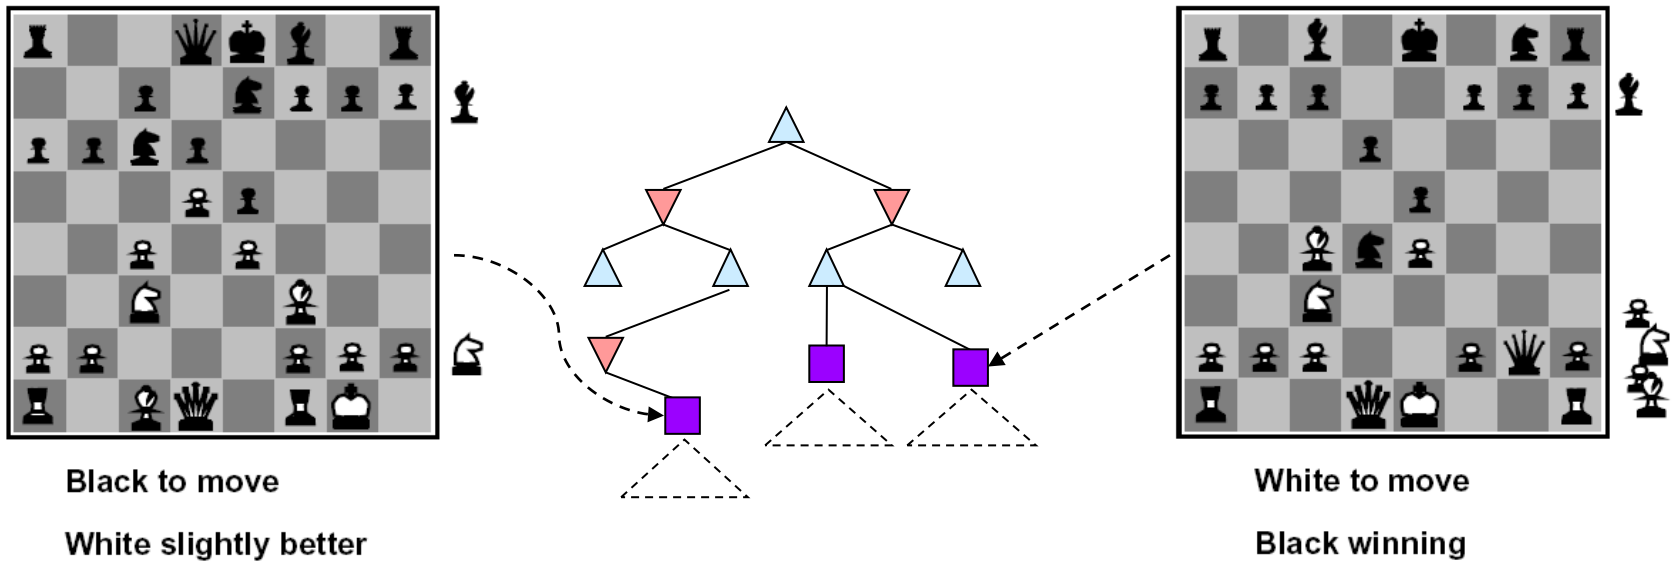

Ideal function: returns the actual minimax value of the position

In practice: typically weighted linear sum of features:

- EVAL(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.

Slide credit: UCB CS 188

Evaluation for Pacman

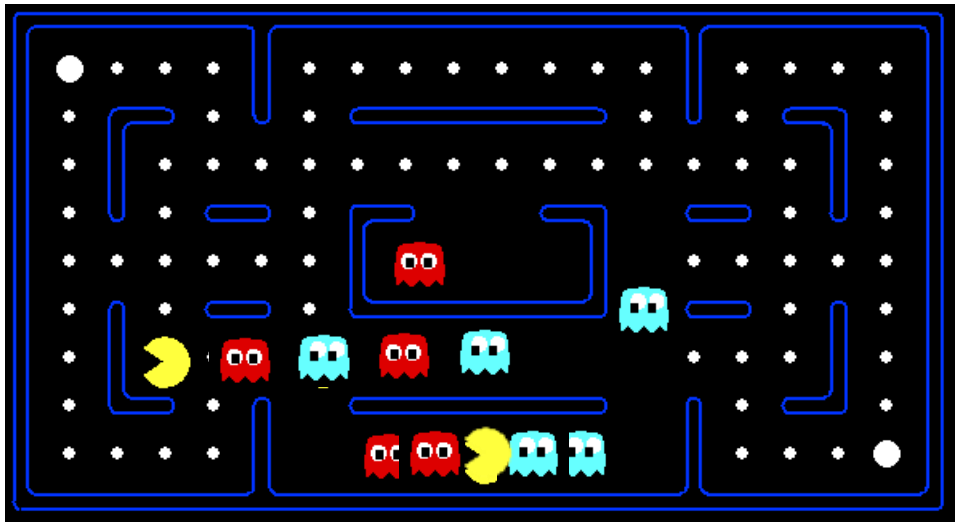

Slide credit: UCB CS 188

Generalized minimax

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…

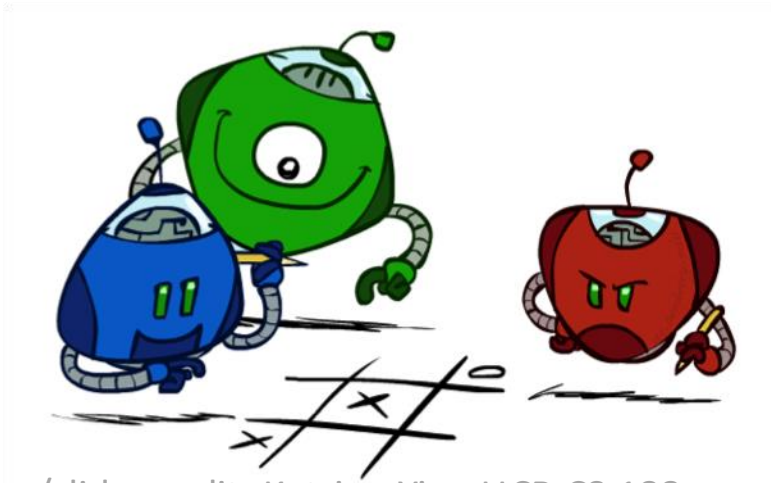

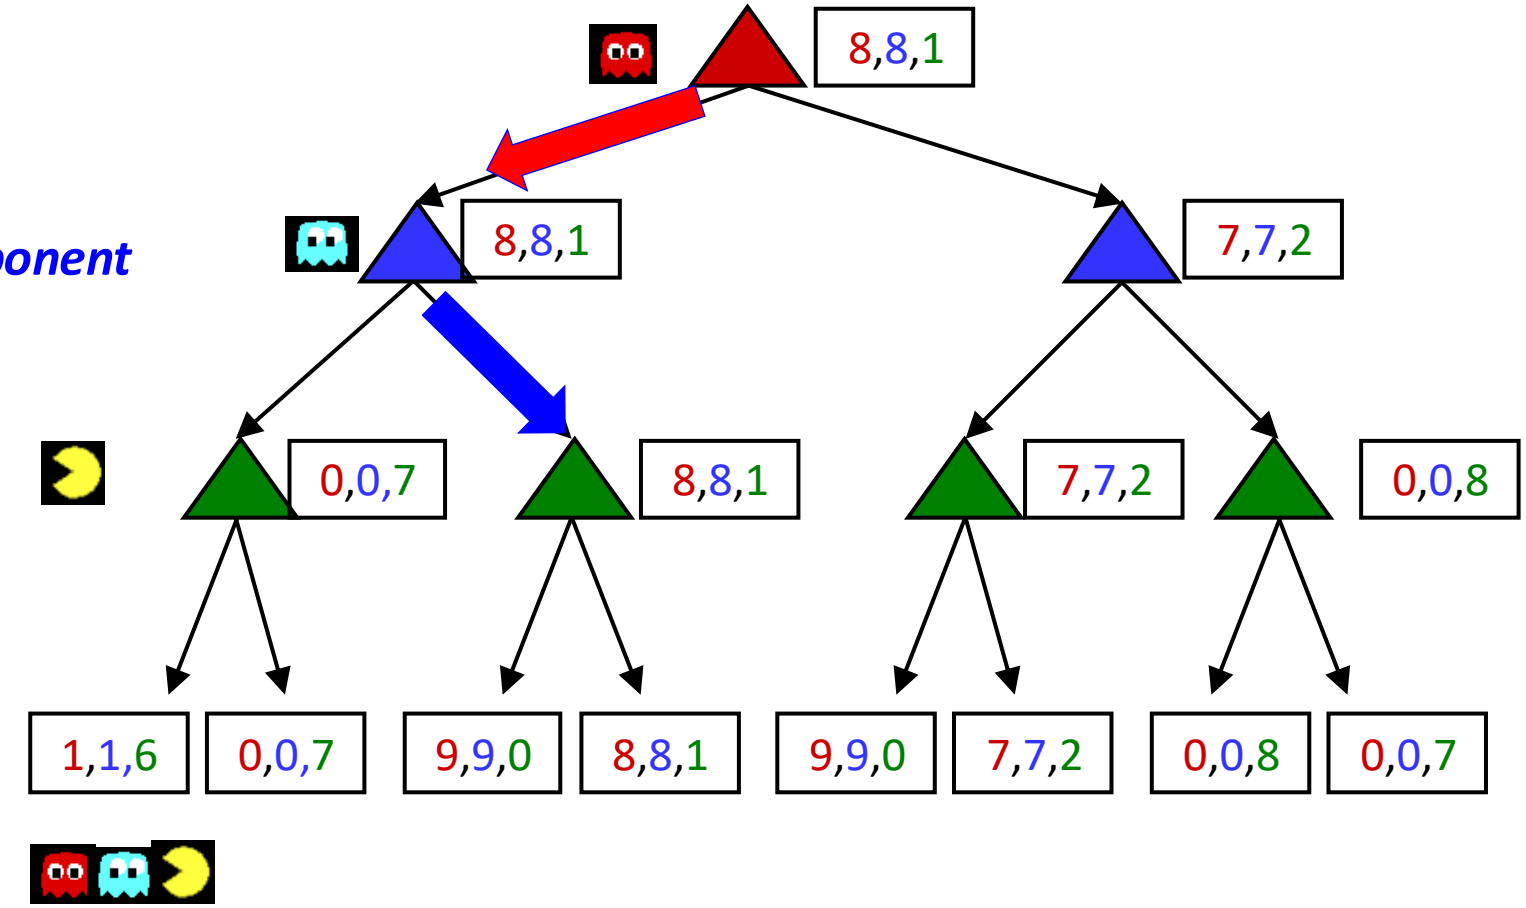

Image/slide credit: Ketrina Yim, UCB CS 188

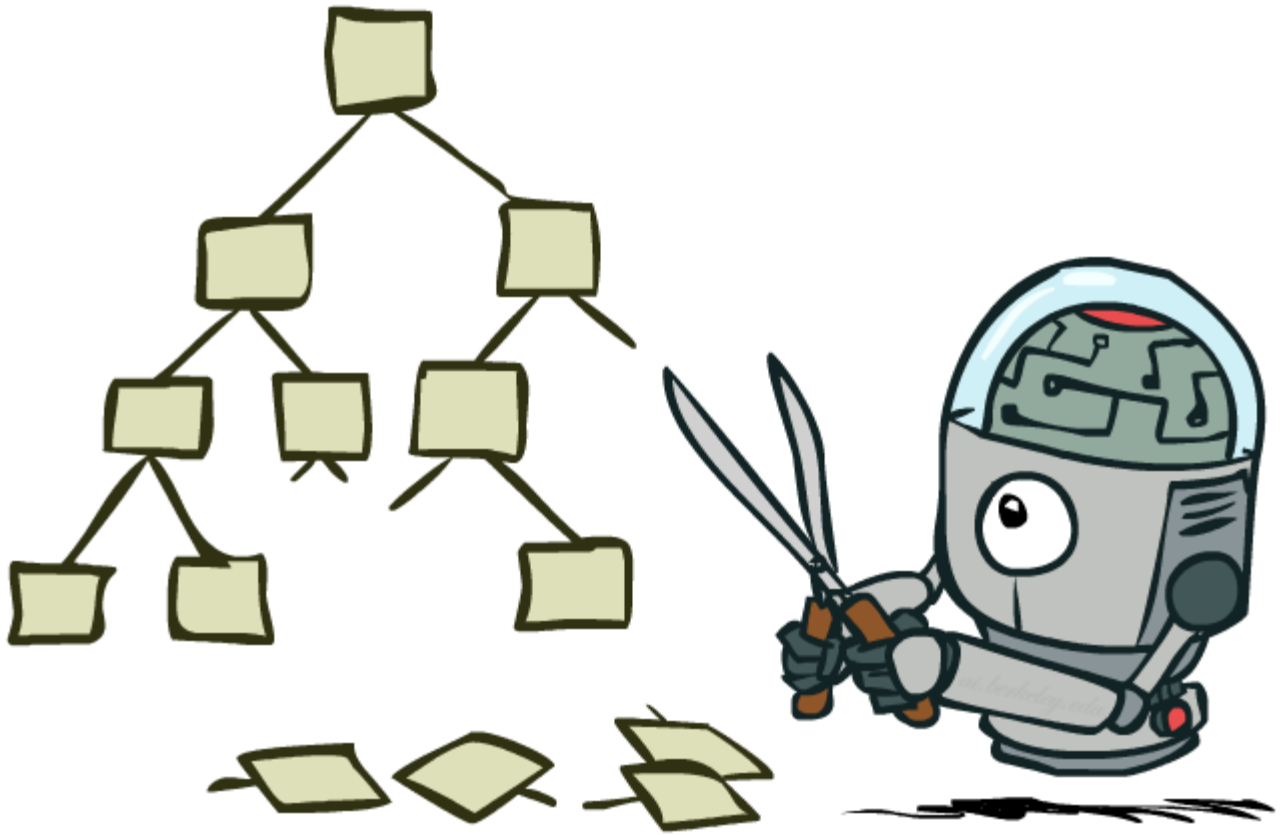

Game Tree Pruning

Image credit: Ketrina Yim, UCB CS 188

Alpha-Beta Example

α = best option so far from any MAX node on this path

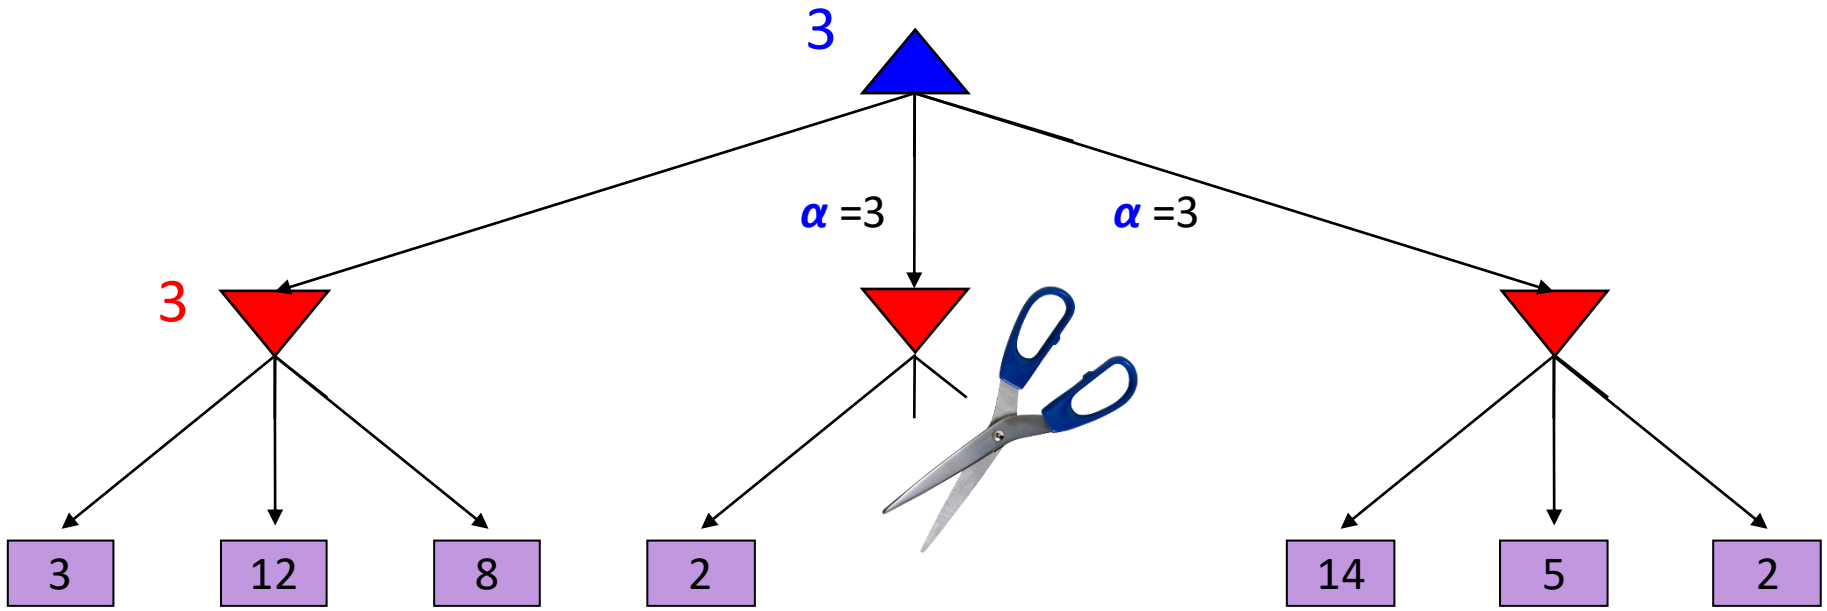

The order of generation matters: more pruning is possible if good moves come first

Slide credit: UCB CS 188

Poll 1

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

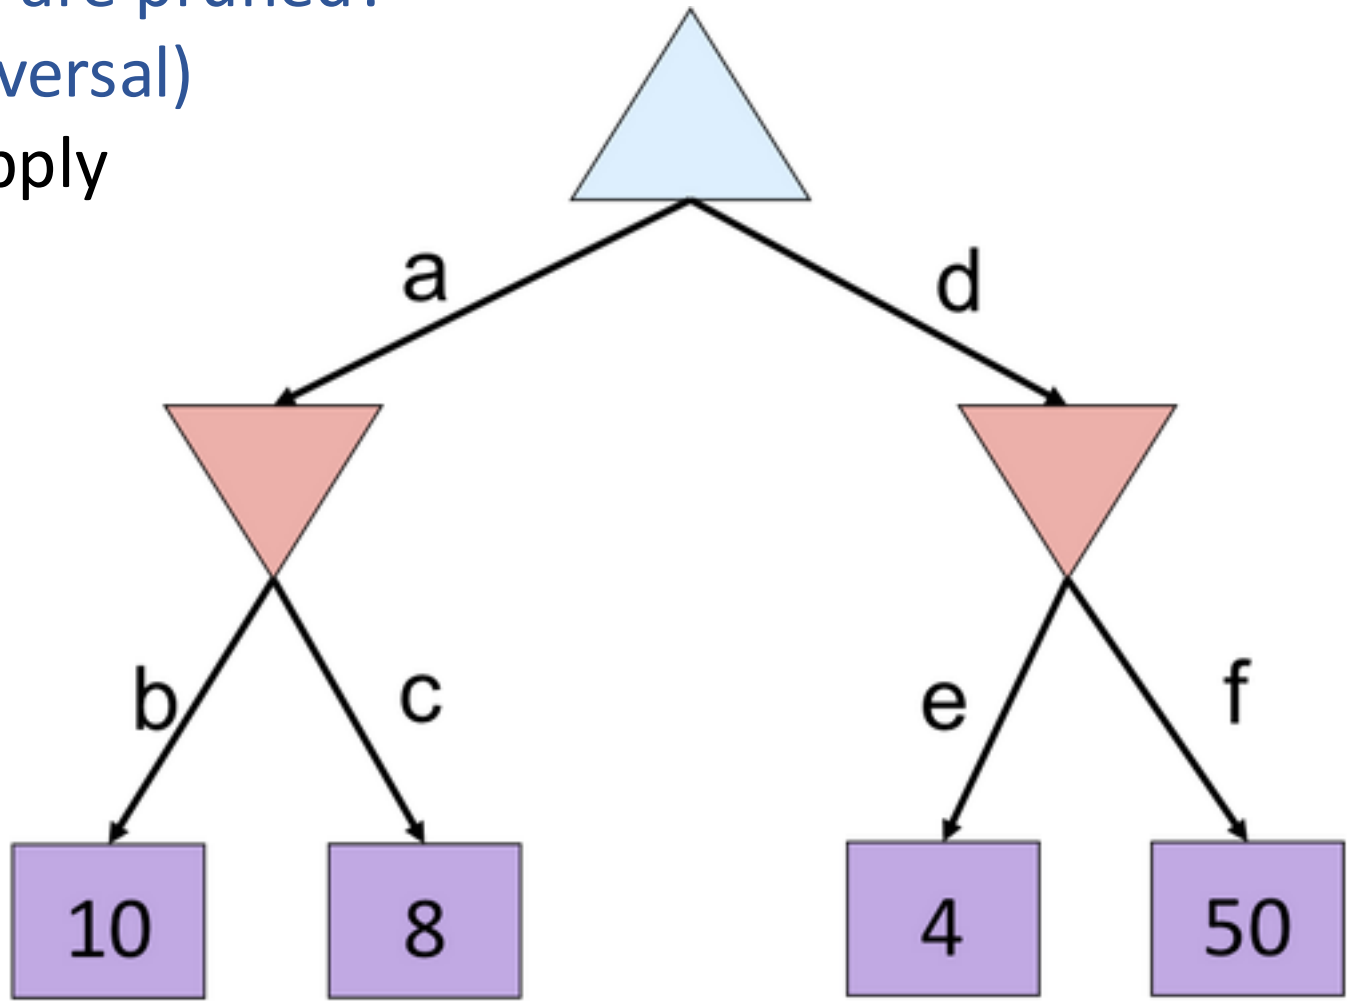

Slide credit: UCB CS 188

Example

Which branches are pruned?
(Left to right traversal)

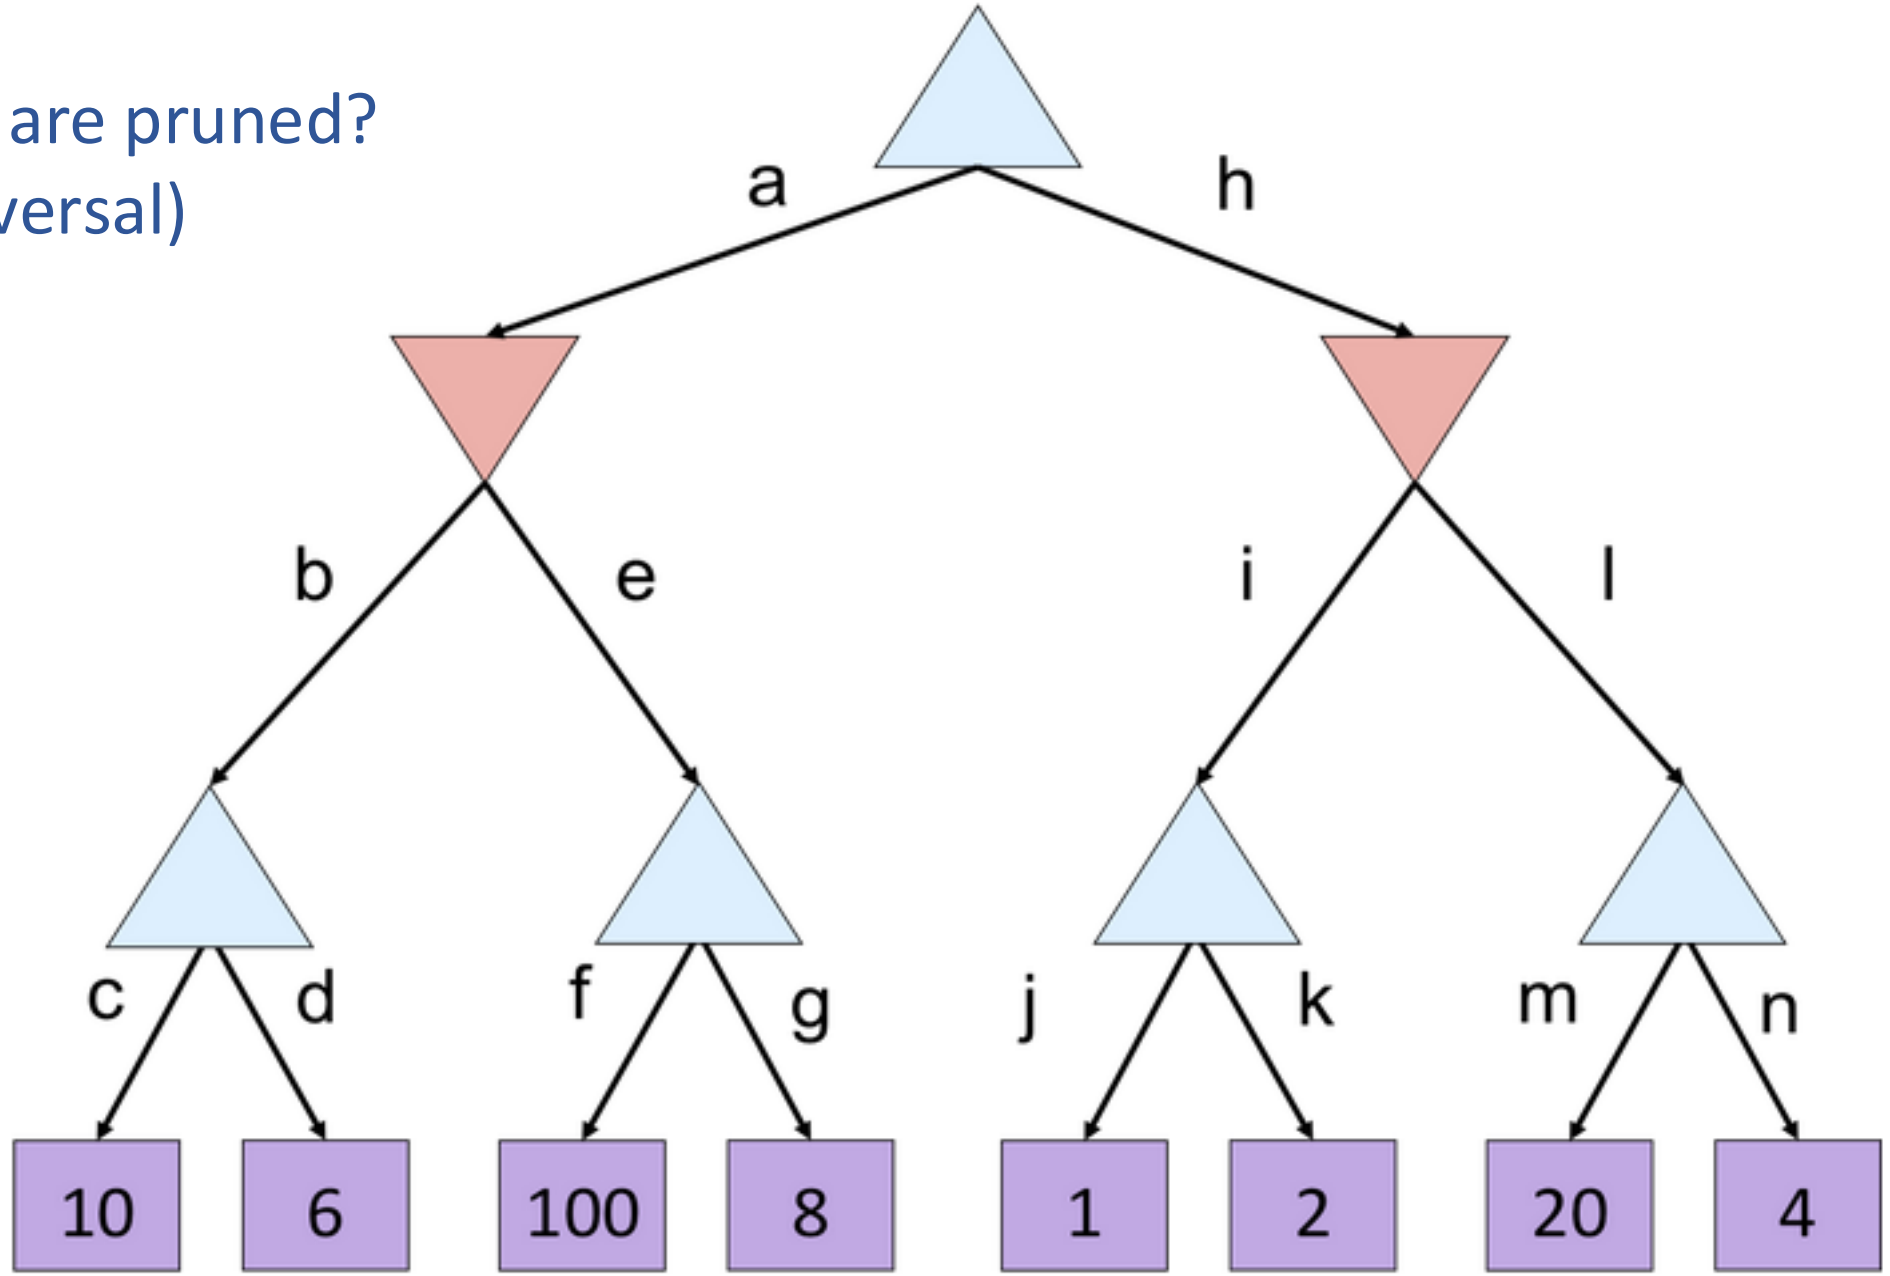

Slide credit: UCB CS 188

Alpha-Beta Implementation

α: MAX's best option on path to root
β: MIN's best option on path to root

```
def max-value(state, α, β):
    initialize v = -∞
    for each successor of state:
        v = max(v, value(successor, α, β))
        if v ≥ β
            return v
        α = max(α, v)
    return v
```

```
def min-value(state , α, β):
    initialize v = +∞
    for each successor of state:
        v = min(v, value(successor, α, β))
        if v ≤ α
            return v
        β = min(β, v)
    return v
```

Slide credit: UCB CS 188

Alpha-Beta version of the Example

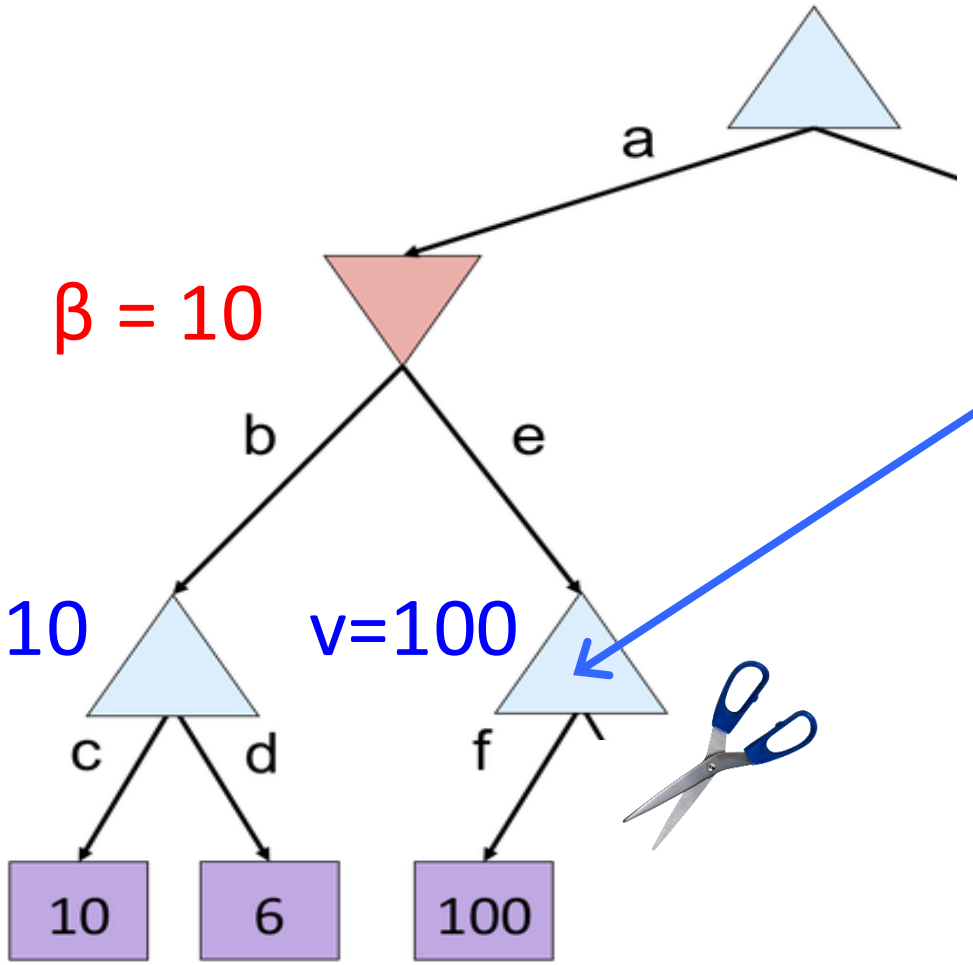

α: MAX's best option on path to root

β: MIN's best option on path to root

```
def max-value(state, α, β):
    initialize v = -∞
    for each successor of state:
        v = max(v, value(successor, α, β))
        if v ≥ β
            return v
        α = max(α, v)
    return v
```

Slide credit: UCB CS 188

Alpha-Beta version of the Example

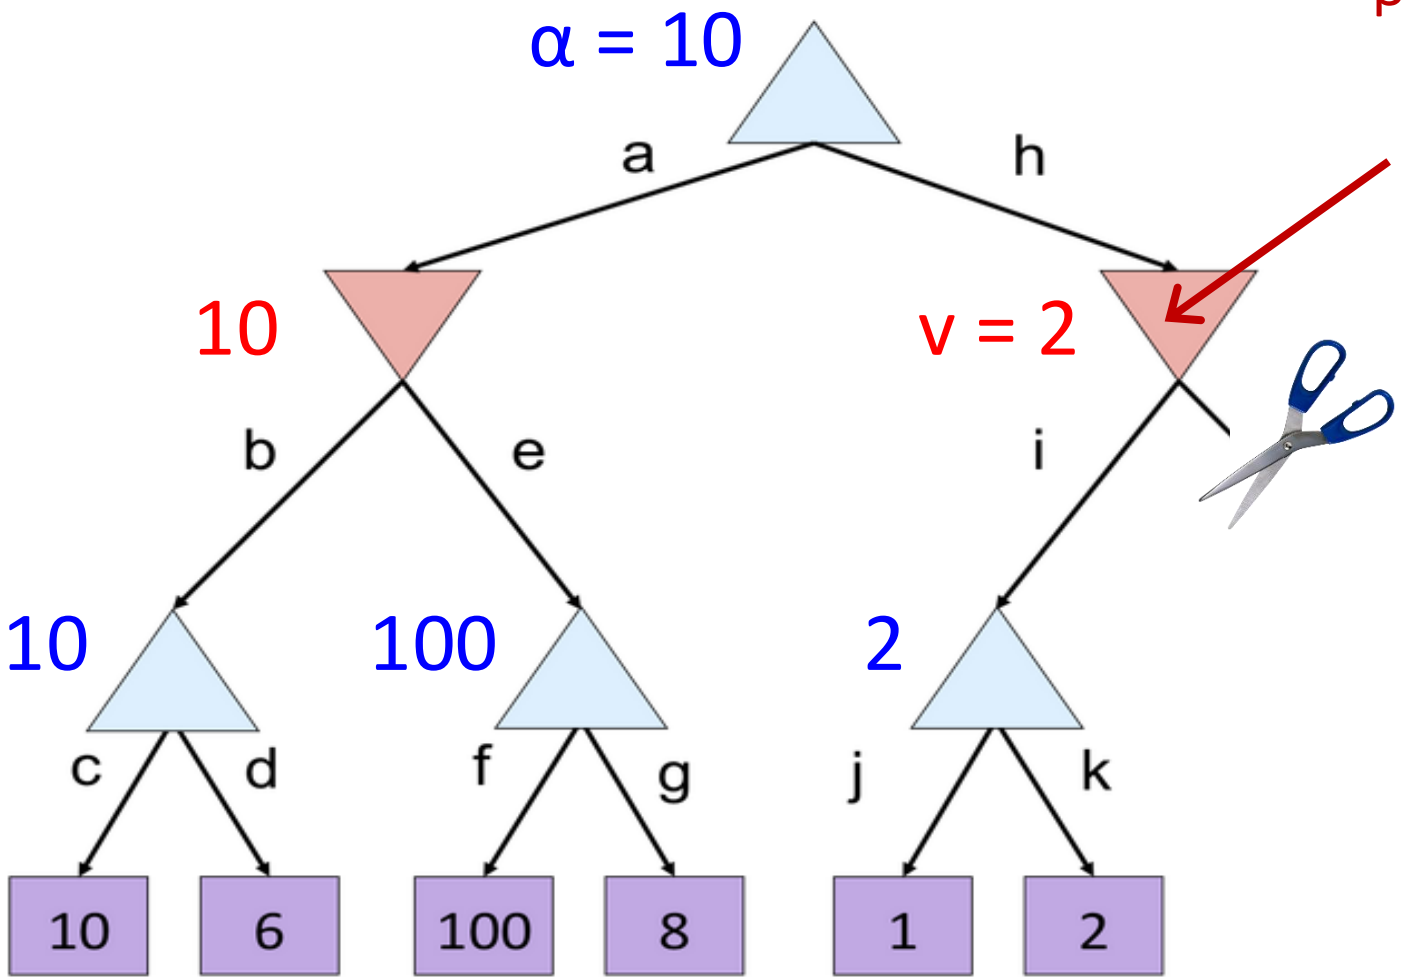

α: MAX's best option on path to root

β: MIN's best option on path to root

```
def min-value(state , α, β):
    initialize v = +∞
    for each successor of state:
        v = min(v, value(successor, α, β))
        if v ≤ α
            return v
        β = min(β, v)
    return v
```

Slide credit: UCB CS 188

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 $O(b^{m/2})$
- Doubles solvable depth!
- 1M nodes/move => depth=8, respectable

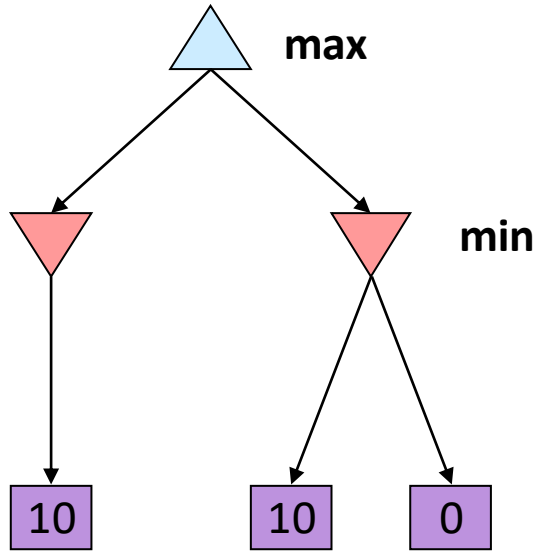

Slide credit: UCB CS 188

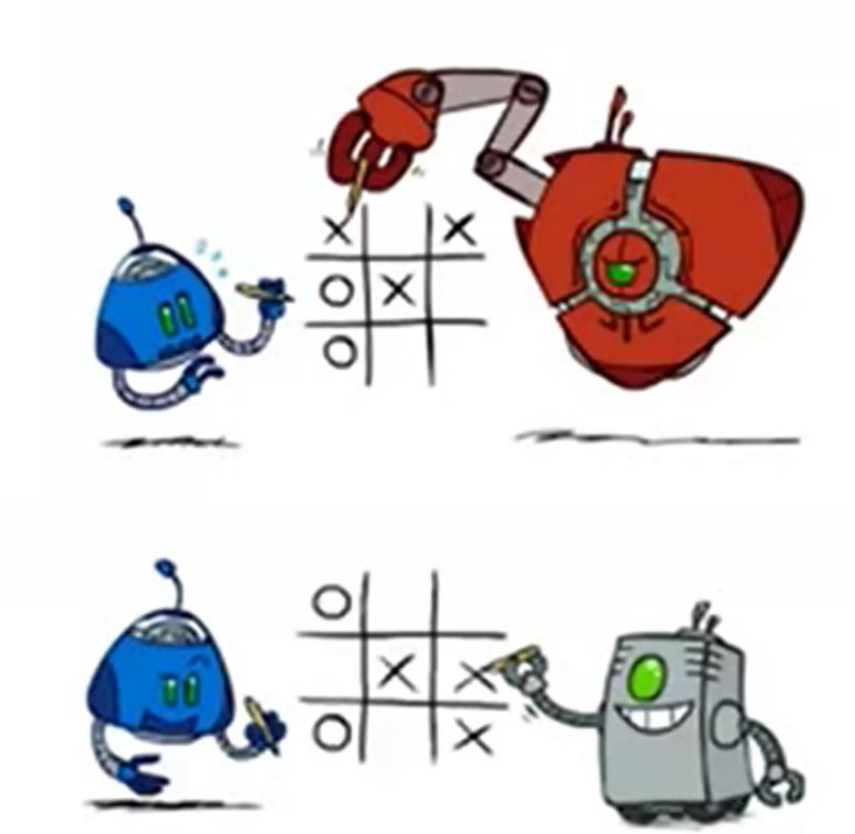

Modeling Assumptions

Image credit: Ketrina Yim, UCB CS 188

Modeling Assumptions

Know your opponent

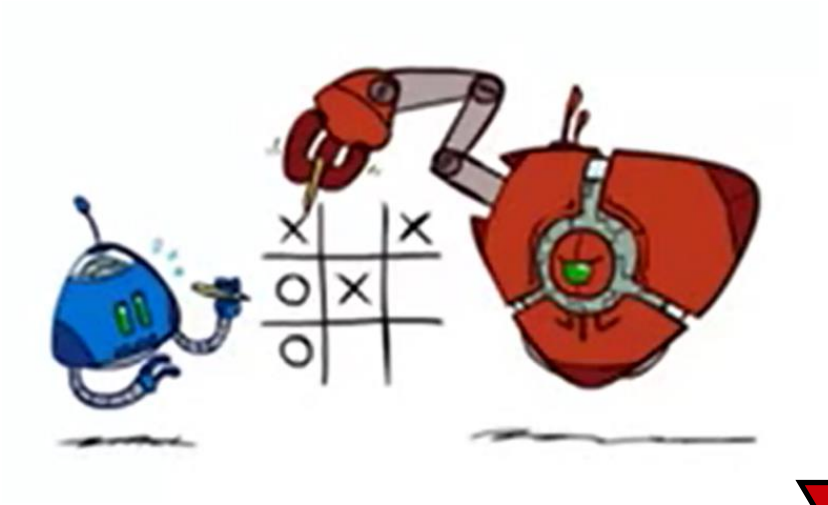

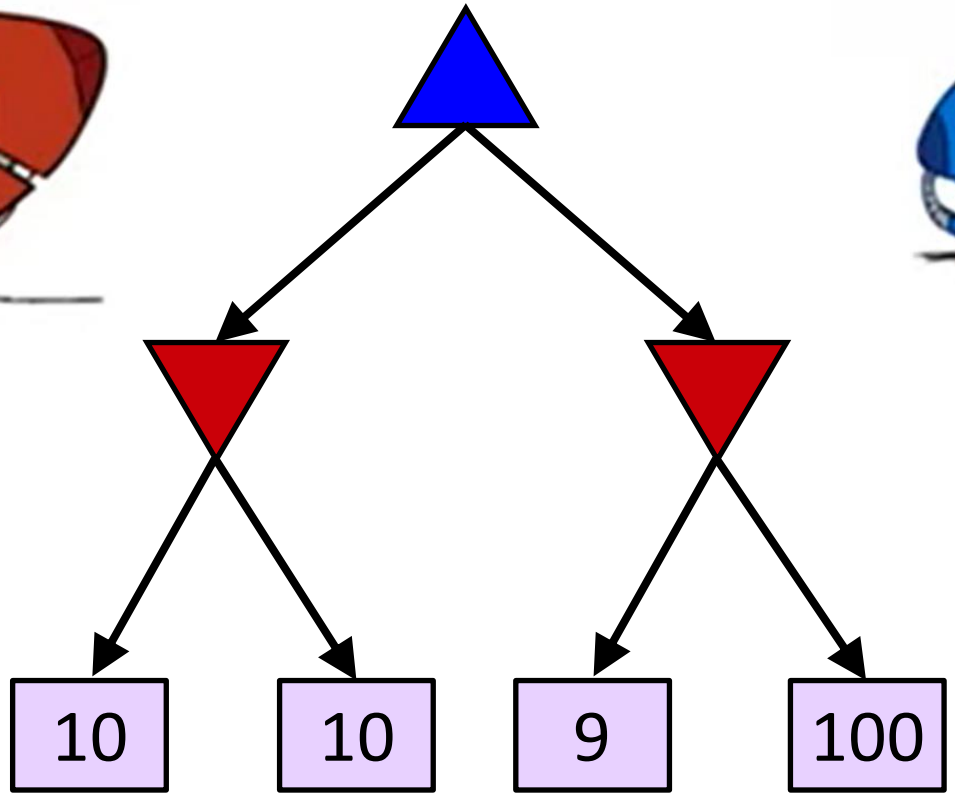

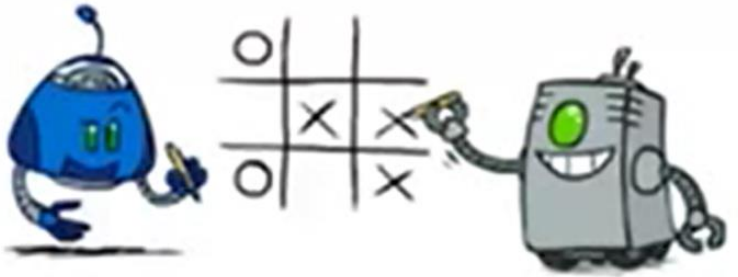

Image/slide credit: Ketrina Yim, UCB CS 188

Minimax Demo

Points

+500 win

-500 lose

-1 each move

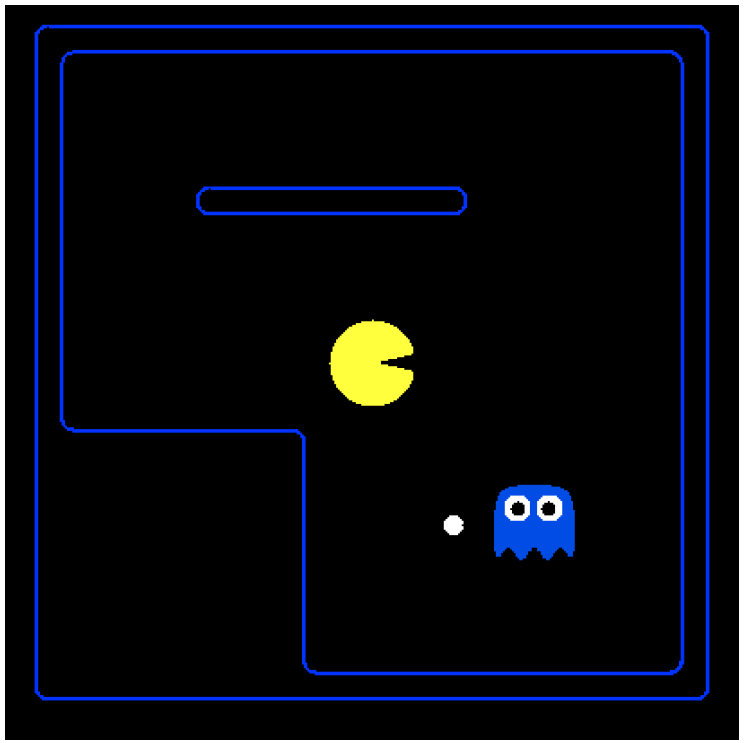

Fine print

- Pacman: uses depth 4 minimax
- Ghost: uses depth 2 minimax

Slide credit: UCB CS 188

[Demo: minimax, minimax L7D4)]

Assumptions vs. Reality

Which pair performs the worst (for Pacman)?

Which is the best?

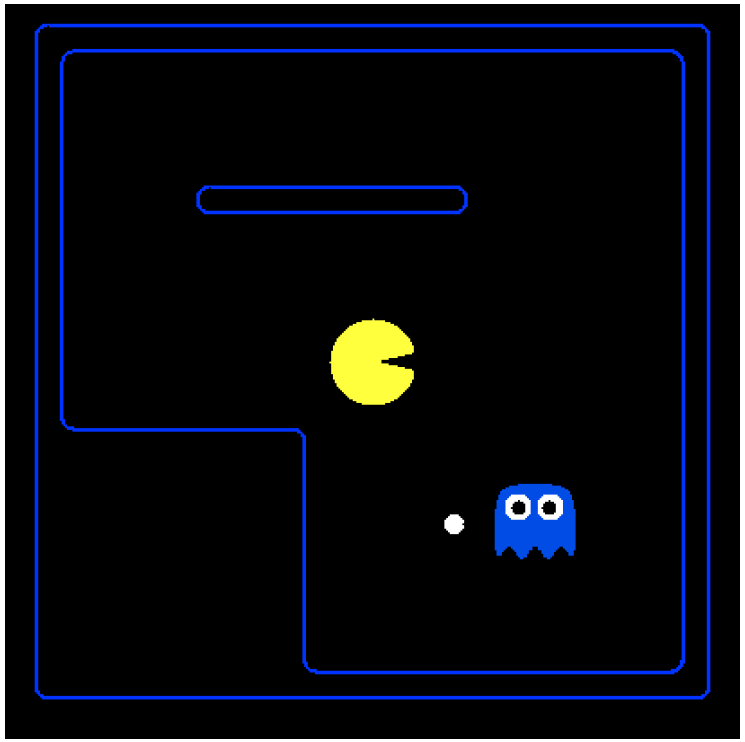

	Minimax Ghost	Random Ghost
Minimax Pacman		
Expectimax Pacman		

Results from playing 5 games

Slide credit: UCB CS 188

[Demo: random, minimax L7D6)]

Modeling Assumptions

Minimax autonomous vehicle?

Image: https://corporate.ford.com/innovation/autonomous-2021.html

Minimax Driver?

https://youtu.be/tzjgkMvTk5E?si=4twYU7UhYCkZqF73&t=89

Clip: How I Met Your Mother, CBS

Modeling Assumptions

Dangerous Pessimism

Assuming the worst case when it's not likely

Dangerous Optimism

Assuming chance when the world is adversarial

Image/slide credit: Ketrina Yim, UCB CS 188

Modeling Assumptions

Chance nodes: Expectimax

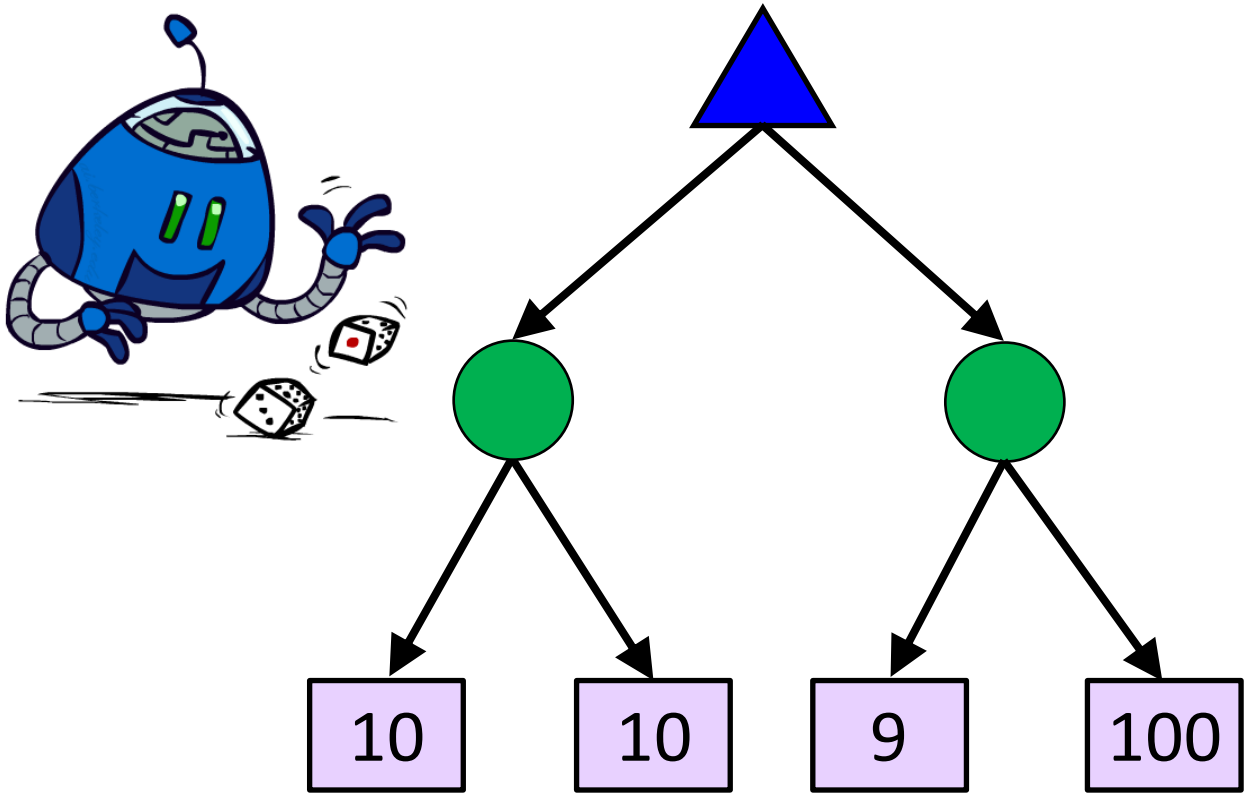

Image/slide credit: Ketrina Yim, UCB CS 188

Chance outcomes in trees

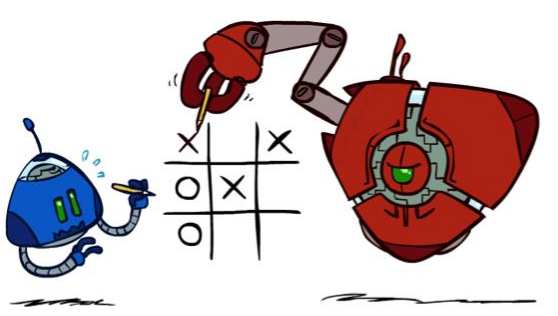

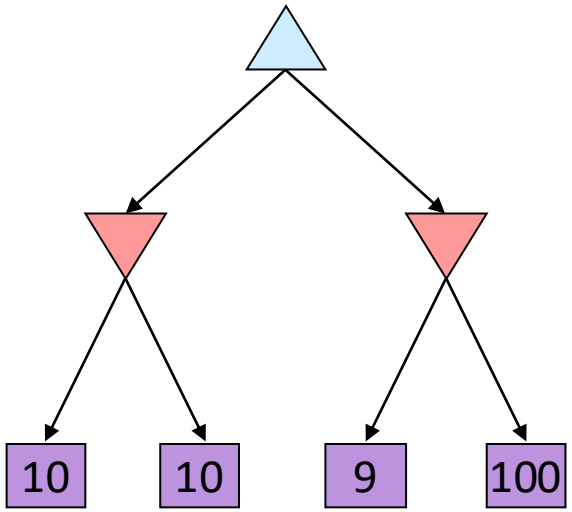

Tictactoe, chess

Minimax

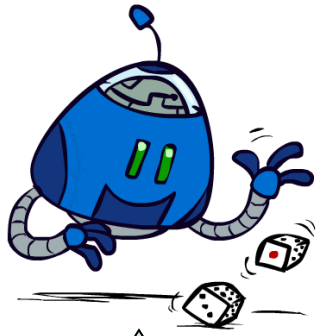

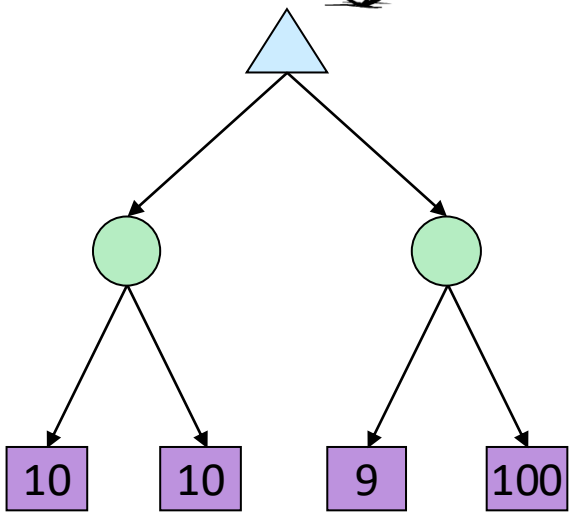

Tetris, investing

Expectimax

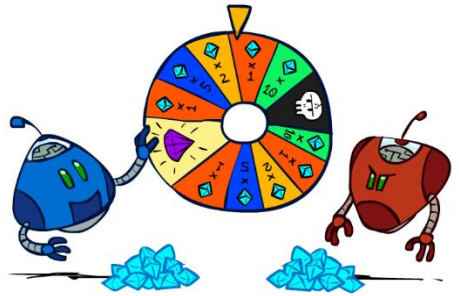

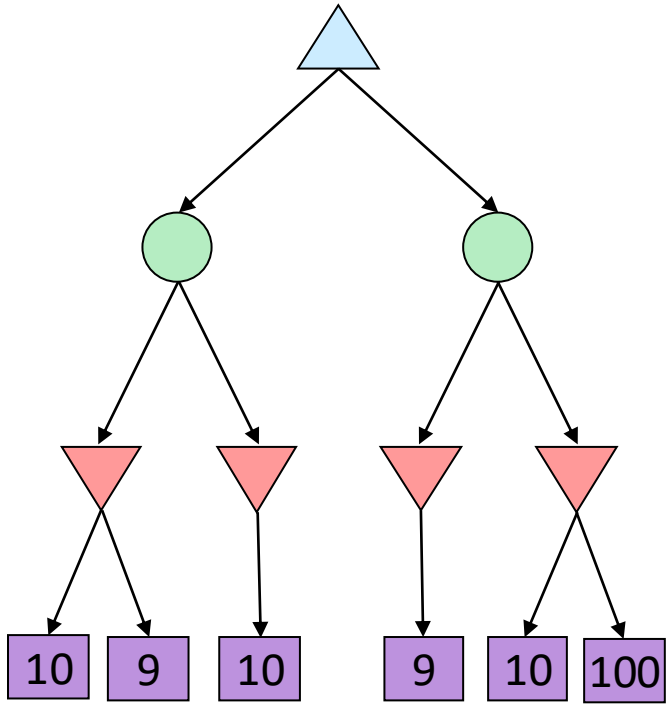

Backgammon, Monopoly

Expectiminimax

Image/slide credit: Ketrina Yim, UCB CS 188

Expectimax Pruning?

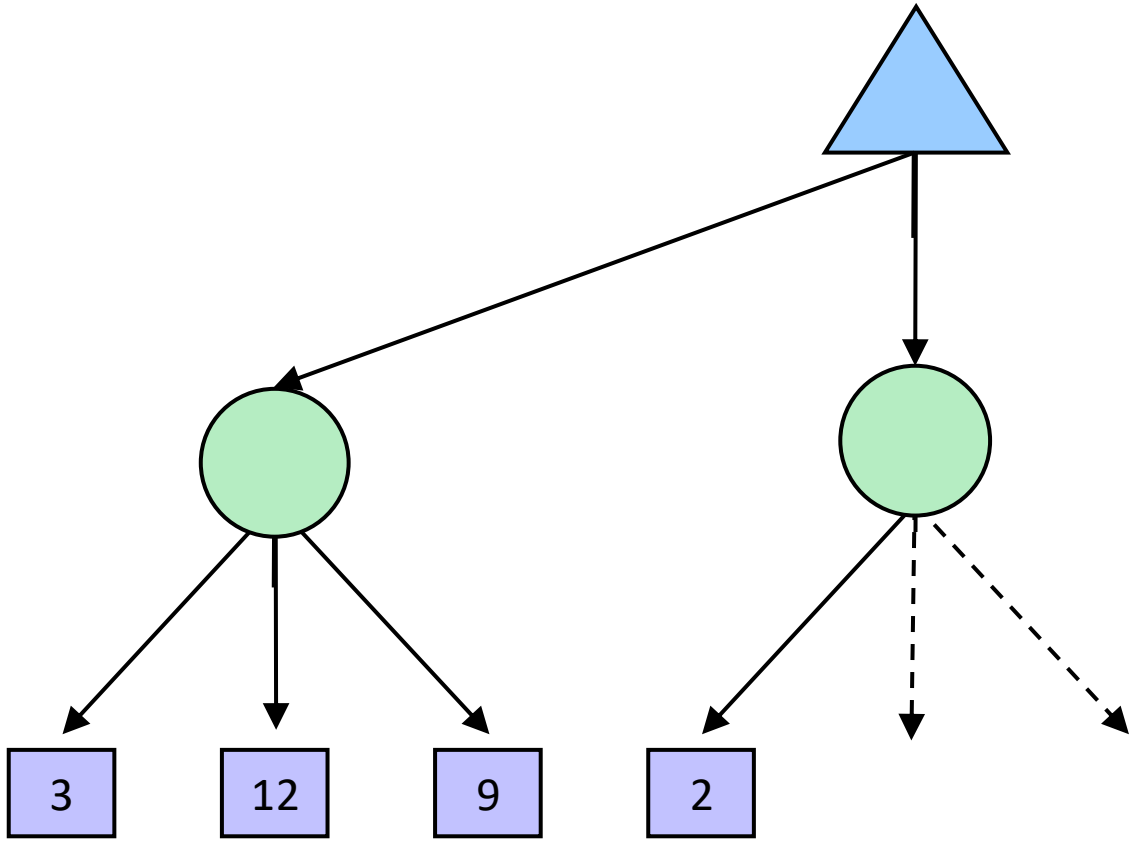

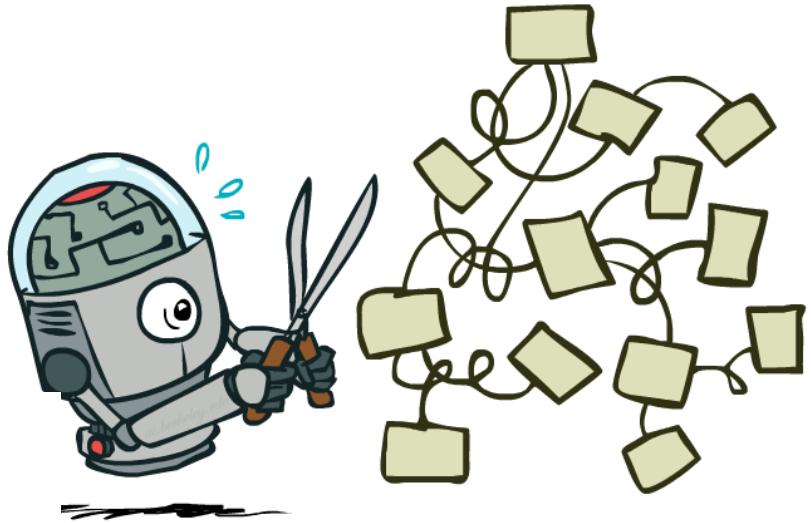

Image/slide credit: Ketrina Yim, UCB CS 188

Expectimax Code

```
function value( state )
    if state.is_leaf
        return state.value

    if state.player is MAX
        return max_{a in state.actions} value( state.result(a) )

    if state.player is MIN
        return min_{a in state.actions} value( state.result(a) )

    if state.player is CHANCE
        return sum_{s in state.next_states} P( s ) * value( s )
```

Slide credit: UCB CS 188

Preview: MDP/Reinforcement Learning Notation

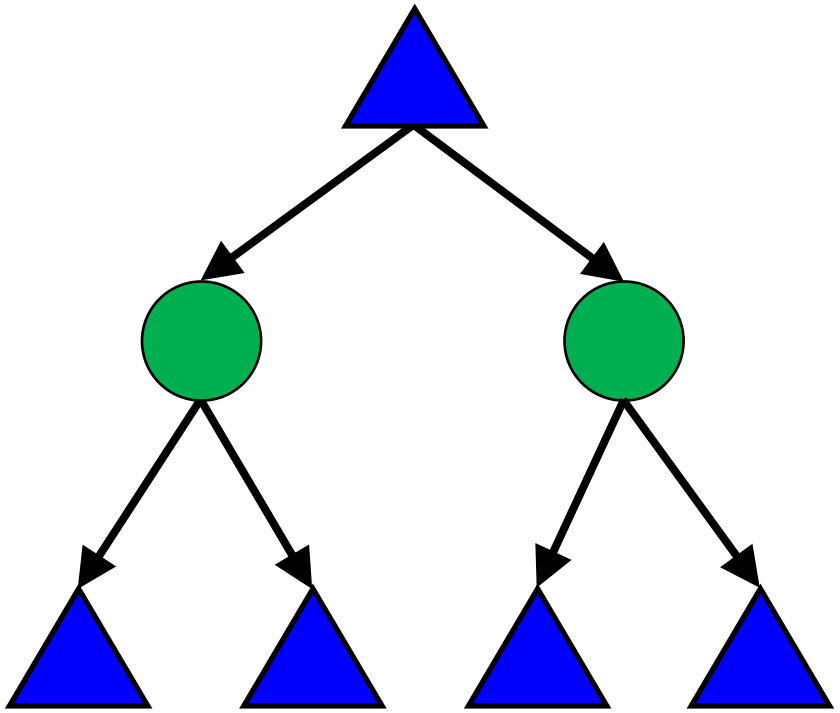

$$V(s) = \max_{a} \sum_{s'} P(s')\, V(s')$$

Preview: MDP/Reinforcement Learning Notation

Standard expectimax: $V(s) = \max_a \sum_{s'} P(s'|s,a)V(s')$

Bellman equations: $V(s) = \max_a \sum_{s'} P(s'|s,a)[R(s,a,s') + \gamma V(s')]$

Value iteration: $V_{k+1}(s) = \max_a \sum_{s'} P(s'|s,a)[R(s,a,s') + \gamma V_k(s')], \quad \forall\, s$

Q-iteration: $Q_{k+1}(s,a) = \sum_{s'} P(s'|s,a)[R(s,a,s') + \gamma \max_{a'} Q_k(s',a')], \quad \forall\, s,a$

Policy extraction: $\pi_V(s) = \operatorname{argmax}_a \sum_{s'} P(s'|s,a)[R(s,a,s') + \gamma V(s')], \quad \forall\, s$

Policy evaluation: $V_{k+1}^{\pi}(s) = \sum_{s'} P(s'|s,\pi(s))[R(s,\pi(s),s') + \gamma V_k^{\pi}(s')], \quad \forall\, s$

Policy improvement: $\pi_{new}(s) = \operatorname{argmax}_a \sum_{s'} P(s'|s,a)[R(s,a,s') + \gamma V^{\pi_{old}}(s')], \quad \forall\, s$

Preview: MDP/Reinforcement Learning Notation

Standard expectimax: $V(s) = \max_a \sum_{s'} P(s'|s,a)V(s')$

Bellman equations: $V(s) = \max_a \sum_{s'} P(s'|s,a)[R(s,a,s') + \gamma V(s')]$

Value iteration: $V_{k+1}(s) = \max_a \sum_{s'} P(s'|s,a)[R(s,a,s') + \gamma V_k(s')], \quad \forall\, s$

Q-iteration: $Q_{k+1}(s,a) = \sum_{s'} P(s'|s,a)[R(s,a,s') + \gamma \max_{a'} Q_k(s',a')], \quad \forall\, s,a$

Policy extraction: $\pi_V(s) = \operatorname{argmax}_a \sum_{s'} P(s'|s,a)[R(s,a,s') + \gamma V(s')], \quad \forall\, s$

Policy evaluation: $V^{\pi}_{k+1}(s) = \sum_{s'} P(s'|s,\pi(s))[R(s,\pi(s),s') + \gamma V^{\pi}_k(s')], \quad \forall\, s$

Policy improvement: $\pi_{new}(s) = \operatorname{argmax}_a \sum_{s'} P(s'|s,a)[R(s,a,s') + \gamma V^{\pi_{old}}(s')], \quad \forall\, s$

Why Expectimax?

Pretty great model for an agent in the world

Choose the action that has the: highest expected value

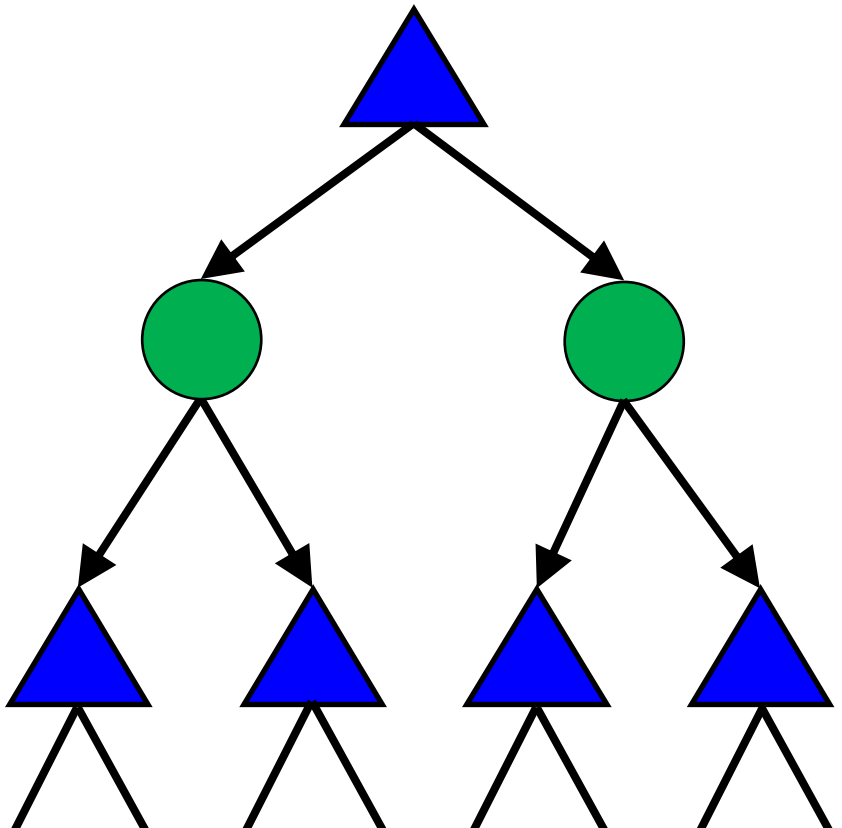

Slide credit: UCB CS 188

Summary

Games require decisions when optimality is impossible

- Bounded-depth search and approximate evaluation functions

Games force efficient use of computation

- 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 in economics (poker)

Video games present much greater challenges – lots to do!

- $b = 10^{500}$, $|S| = 10^{4000}$, $m = 10{,}000$

Slide credit: UCB CS 188