

# CS325 Artificial Intelligence

## Ch. 5, Games!

Cengiz Günay, Emory Univ.



vs.



Spring 2013

A lot of work is done on it. Why?

A lot of work is done on it. Why?

- Fun, provide entertainment
- Also, simpler than life: toy problems



A lot of work is done on it. Why?

- Fun, provide entertainment
- Also, simpler than life: toy problems

Types of game AIs:



A lot of work is done on it. Why?

- Fun, provide entertainment
- Also, simpler than life: toy problems

Types of game AIs:

- Adversaries



zerg rush

A lot of work is done on it. Why?

- Fun, provide entertainment
- Also, simpler than life: toy problems

Types of game AIs:

- Adversaries
- Simulated reality (non-playable characters, world reaction to player).



zerg rush



A lot of work is done on it. Why?

- Fun, provide entertainment
- Also, simpler than life: toy problems

Types of game AIs:

- Adversaries
- Simulated reality (non-playable characters, world reaction to player).
- Game theory (next class)



zerg rush



## Exit survey: Hidden Markov Models

- In the mining robot example, when is the uncertainty of the robot's trajectories reduced?
- How is Particle Filtering like and unlike a water filter?

## Entry survey: Adversarial Games (0.25 points of final grade)

- What algorithm would be useful in games? Give examples with two different algorithms you learned in class.
- How would you help an agent solve a problem against an adversary? Think of a game like chess or checkers for starters.



Can previous algorithms help in games?

- Single-state agent:

Can previous algorithms help in games?

- Single-state agent: vacuum world, solving a maze

Can previous algorithms help in games?

- Single-state agent: vacuum world, solving a maze
- Bayes Nets:

## Previously on AI for Games...

Can previous algorithms help in games?

- Single-state agent: vacuum world, solving a maze
- Bayes Nets: card games

Can previous algorithms help in games?

- Single-state agent: vacuum world, solving a maze
- Bayes Nets: card games
- Machine Learning:

Can previous algorithms help in games?

- Single-state agent: vacuum world, solving a maze
- Bayes Nets: card games
- Machine Learning: guessing games, learning user moves

Can previous algorithms help in games?

- Single-state agent: vacuum world, solving a maze
- Bayes Nets: card games
- Machine Learning: guessing games, learning user moves
- Logic, planning:

Can previous algorithms help in games?

- Single-state agent: vacuum world, solving a maze
- Bayes Nets: card games
- Machine Learning: guessing games, learning user moves
- Logic, planning: board game with complex rules (Machinarium)



Can previous algorithms help in games?

- Single-state agent: vacuum world, solving a maze
- Bayes Nets: card games
- Machine Learning: guessing games, learning user moves
- Logic, planning: board game with complex rules (Machinarium)
- MDPs, Reinforcement Learning:

Can previous algorithms help in games?

- Single-state agent: vacuum world, solving a maze
- Bayes Nets: card games
- Machine Learning: guessing games, learning user moves
- Logic, planning: board game with complex rules (Machinarium)
- MDPs, Reinforcement Learning: pathfinding, optimal strategy for zerg

Can previous algorithms help in games?

- Single-state agent: vacuum world, solving a maze
- Bayes Nets: card games
- Machine Learning: guessing games, learning user moves
- Logic, planning: board game with complex rules (Machinarium)
- MDPs, Reinforcement Learning: pathfinding, optimal strategy for zerg
- HMMs, Particle Filter:

Can previous algorithms help in games?

- Single-state agent: vacuum world, solving a maze
- Bayes Nets: card games
- Machine Learning: guessing games, learning user moves
- Logic, planning: board game with complex rules (Machinarium)
- MDPs, Reinforcement Learning: pathfinding, optimal strategy for zerg
- HMMs, Particle Filter: state estimation and future prediction, partially observable environment with traps (sonic?)

Can previous algorithms help in games?

- Single-state agent: vacuum world, solving a maze
- Bayes Nets: card games
- Machine Learning: guessing games, learning user moves
- Logic, planning: board game with complex rules (Machinarium)
- MDPs, Reinforcement Learning: pathfinding, optimal strategy for zerg
- HMMs, Particle Filter: state estimation and future prediction, partially observable environment with traps (sonic?)

**None for adversaries?**

# Properties of Games as Agent Environment

Stochastic	Part.-Observ.	Unknown	Adversarial	Game
				Chess, Checkers
				Robot Soccer
				Poker
				Hide-and-go-seek
				Starcraft
				Battle for Wesnoth
				Halo/CoD/MoH
				Solitaire
				Minesweeper
				Zuma

# Properties of Games as Agent Environment

Stochastic	Part.-Observ.	Unknown	Adversarial	Game
			X	Chess, Checkers Robot Soccer Poker Hide-and-go-seek Starcraft Battle for Wesnoth Halo/CoD/MoH Solitaire Minesweeper Zuma

# Properties of Games as Agent Environment

Stochastic	Part.-Observ.	Unknown	Adversarial	Game
			X	Chess, Checkers
X	X		X	Robot Soccer
				Poker
				Hide-and-go-seek
				Starcraft
				Battle for Wesnoth
				Halo/CoD/MoH
				Solitaire
				Minesweeper
				Zuma



# Properties of Games as Agent Environment

Stochastic	Part.-Observ.	Unknown	Adversarial	Game
			X	Chess, Checkers
X	X		X	Robot Soccer
	X		X	Poker
				Hide-and-go-seek
				Starcraft
				Battle for Wesnoth
				Halo/CoD/MoH
				Solitaire
				Minesweeper
				Zuma

# Properties of Games as Agent Environment

Stochastic	Part.-Observ.	Unknown	Adversarial	Game
			X	Chess, Checkers
X	X		X	Robot Soccer
	X		X	Poker
X	X	X	X	Hide-and-go-seek Starcraft Battle for Wesnoth Halo/CoD/MoH Solitaire Minesweeper Zuma

# Properties of Games as Agent Environment

Stochastic	Part.-Observ.	Unknown	Adversarial	Game
			X	Chess, Checkers
X	X		X	Robot Soccer
	X		X	Poker
X	X	X	X	Hide-and-go-seek
	X		X	Starcraft
				Battle for Wesnoth
				Halo/CoD/MoH
				Solitaire
				Minesweeper
				Zuma

# Properties of Games as Agent Environment

Stochastic	Part.-Observ.	Unknown	Adversarial	Game
			X	Chess, Checkers
X	X		X	Robot Soccer
	X		X	Poker
X	X	X	X	Hide-and-go-seek
	X		X	Starcraft
X	X		X	Battle for Wesnoth
				Halo/CoD/MoH
				Solitaire
				Minesweeper
				Zuma

# Properties of Games as Agent Environment

Stochastic	Part.-Observ.	Unknown	Adversarial	Game
			X	Chess, Checkers
X	X		X	Robot Soccer
	X		X	Poker
X	X	X	X	Hide-and-go-seek
	X		X	Starcraft
X	X		X	Battle for Wesnoth
X	X		X	Halo/CoD/MoH
				Solitaire
				Minesweeper
				Zuma

# Properties of Games as Agent Environment

Stochastic	Part.-Observ.	Unknown	Adversarial	Game
			X	Chess, Checkers
X	X		X	Robot Soccer
	X		X	Poker
X	X	X	X	Hide-and-go-seek
	X		X	Starcraft
X	X		X	Battle for Wesnoth
X	X		X	Halo/CoD/MoH
	X			Solitaire
	X			Minesweeper
	X			Zuma

# Single Player Games



Deterministic, single-state agent → Single-player game using **tree search**

- initial state
- player state
- possible actions
- results of actions
- utility values
- goal test

## Adversarial Games

- 1 Start by adapting single-state agent to games
- 2 Define adversary as someone who wants you to lose
- 3 And makes decisions based on the outcome of your moves



## Adversarial Games

- 1 Start by adapting single-state agent to games
- 2 Define adversary as someone who wants you to lose
- 3 And makes decisions based on the outcome of your moves

2-player games:

- Deterministic
- **Zero-sum**: Reward distributed between players
- **Minimax** algorithm:  
max & min players  
choose +/- utility, resp.

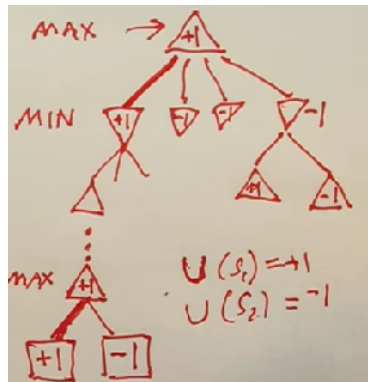
# Adversarial Games

## Adversarial Games

- 1 Start by adapting single-state agent to games
- 2 Define adversary as someone who wants you to lose
- 3 And makes decisions based on the outcome of your moves

2-player games:

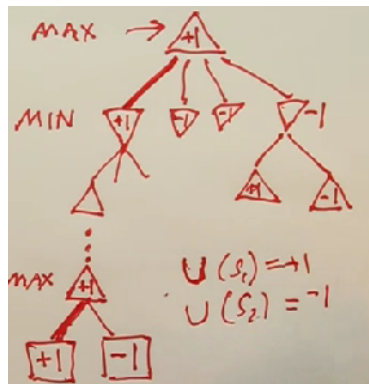
- Deterministic
- **Zero-sum**: Reward distributed between players
- **Minimax** algorithm:  
max & min players  
choose +/- utility, resp.



# 2-Player Value Function

defun value( $s$ ):

- if  $s$  is  $\square$ :  $U(s)$
- if  $s$  is  $\triangle$ :  $\max\text{Value}(s)$
- if  $s$  is  $\nabla$ :  $\min\text{Value}(s)$



# 2-Player Value Function

defun value( $s$ ):

- if  $s$  is  $\square$ :  $U(s)$
- if  $s$  is  $\triangle$ :  $\text{maxValue}(s)$
- if  $s$  is  $\nabla$ :  $\text{minValue}(s)$

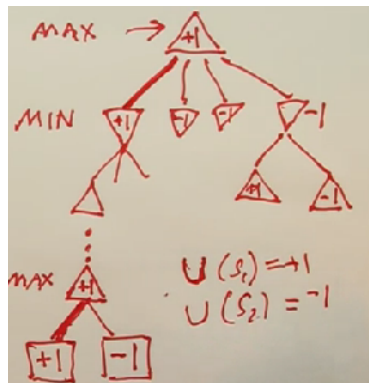
defun  $m = \text{maxValue}(s)$ :

$m = -\infty$

for  $(a, s')$  in successors( $s$ ):

$v = \text{value}(s')$

$m = \text{maxValue}(m, v)$



# 2-Player Value Function

defun value( $s$ ):

- if  $s$  is  $\square$ :  $U(s)$
- if  $s$  is  $\triangle$ :  $\text{maxValue}(s)$
- if  $s$  is  $\nabla$ :  $\text{minValue}(s)$


defun  $m = \text{maxValue}(s)$ :

$m = -\infty$

for  $(a, s')$  in successors( $s$ ):

$v = \text{value}(s')$

$m = \text{maxValue}(m, v)$

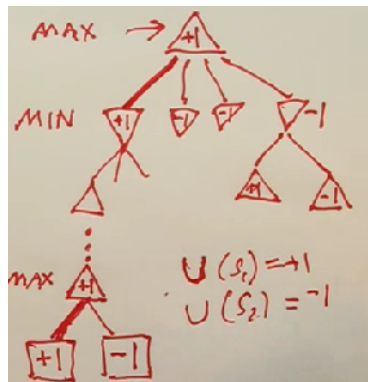


**Assumes opponent is perfect!**

# Time Complexity

For a tree with  
branching factor,  $b$ , and depth,  $m$ ?

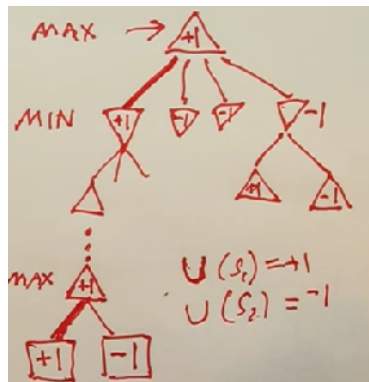
- 1  $O(bm)$
- 2  $O(b^m)$
- 3  $O(m^b)$



# Time Complexity

For a tree with  
branching factor,  $b$ , and depth,  $m$ ?

- 1  $O(bm)$
- 2  $O(b^m)$
- 3  $O(m^b)$



# Time Complexity

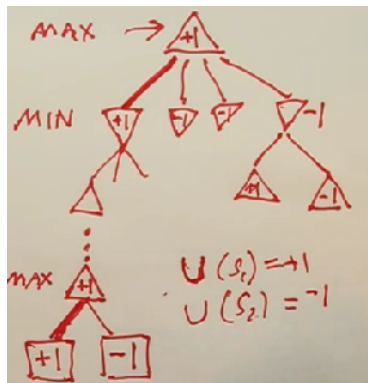
For a tree with  
branching factor,  $b$ , and depth,  $m$ ?

- 1  $O(bm)$
- 2  $O(b^m)$
- 3  $O(m^b)$

For chess:  $b \simeq 30$ ,  $m \simeq 40$   
How long would it take with:

- 1 billion processors  $\times$  1 billion/s evals?

- 1 seconds
- 2 minutes
- 3 hours
- 4 years
- 5 forever





# Time Complexity

For a tree with  
branching factor,  $b$ , and depth,  $m$ ?

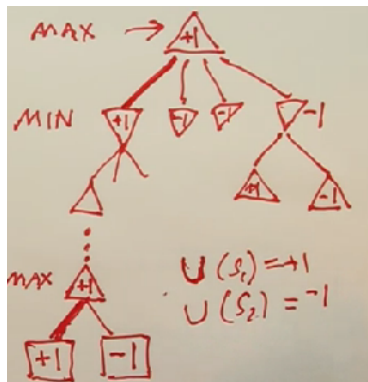
- 1  $O(bm)$
- 2  $O(b^m)$
- 3  $O(m^b)$

For chess:  $b \simeq 30$ ,  $m \simeq 40$

How long would it take with:

- 1 billion processors  $\times$  1 billion/s evals?


- 1 seconds
- 2 minutes
- 3 hours
- 4 years
- 5 **forever**



# Space Complexity

For a tree with  
branching factor,  $b$ , and depth,  $m$ ?

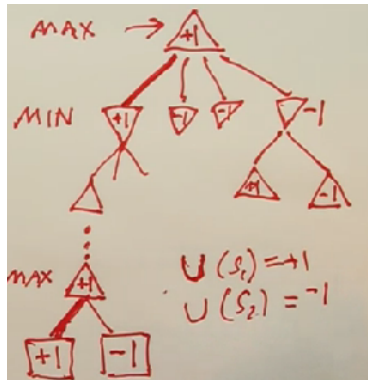
- 1  $O(bm)$
- 2  $O(b^m)$
- 3  $O(m^b)$



# Space Complexity

For a tree with  
branching factor,  $b$ , and depth,  $m$ ?

- 1  $O(bm)$
- 2  $O(b^m)$
- 3  $O(m^b)$

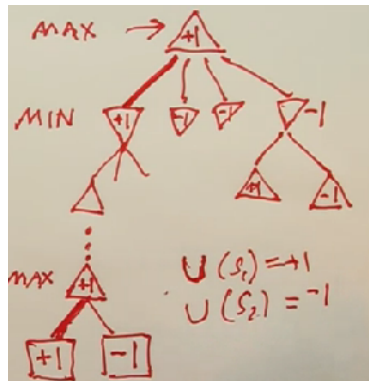


# Space Complexity

For a tree with  
branching factor,  $b$ , and depth,  $m$ ?

- 1  $O(bm)$
- 2  $O(b^m)$
- 3  $O(m^b)$

Do not need more than total number of nodes.



# Complexity Reduction?

## How to do it?

- 1 Reduce  $b$
- 2 Reduce  $m$
- 3 Tree  $\rightarrow$  graph

# Complexity Reduction?

## How to do it?

- 1 Reduce  $b$
- 2 Reduce  $m$
- 3 Tree  $\rightarrow$  graph

**All of the above!**

# Example

defun value( $s$ ):

- if  $s$  is  $\square$ :  $U(s)$
- if  $s$  is  $\Delta$ :  $\text{maxValue}(s)$
- if  $s$  is  $\nabla$ :  $\text{minValue}(s)$


defun  $m = \text{maxValue}(s)$ :

$m = -\infty$

for  $(a, s')$  in  $\text{successors}(s)$ :

$v = \text{value}(s')$

$m = \text{maxValue}(m, v)$



# Example

defun value( $s$ ):

- if  $s$  is  $\square$ :  $U(s)$
- if  $s$  is  $\Delta$ :  $\text{maxValue}(s)$
- if  $s$  is  $\nabla$ :  $\text{minValue}(s)$


defun  $m = \text{maxValue}(s)$ :

$m = -\infty$

for  $(a, s')$  in  $\text{successors}(s)$ :


$v = \text{value}(s')$

$m = \text{maxValue}(m, v)$







# Reducing Branching Factor, $b$



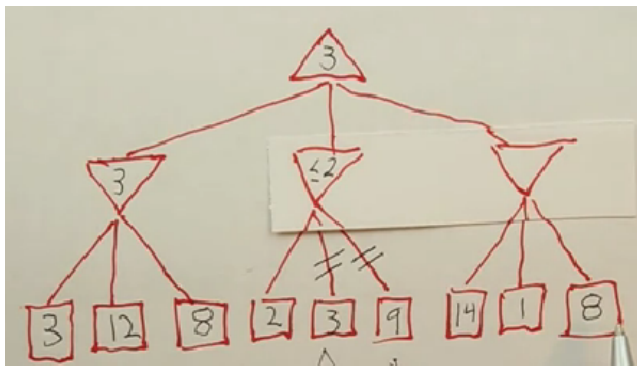
# Reducing Branching Factor, $b$



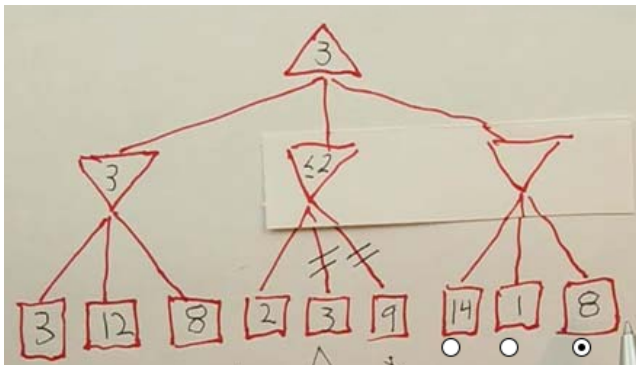
# Reducing Branching Factor, $b$



# Reducing Branching Factor, $b$



# Reducing Branching Factor, $b$



Which one to prune?

Select a cutoff:

- Limit  $m$  (e.g., plan 3 steps ahead in chess)
- Estimate terminal nodes' utility with evaluation function
  - like heuristics
- Learn from experience
- In chess, use board state, value of pieces, etc.
- For value of pieces:  $\text{eval}(s) = \sum_i w_i p_i$ 
  - can use machine learning for  $w_i$

# Formalize as Alpha-Beta Pruning

defun value( $s$ ):

  cutoff at depth  $m'$ : eval( $s$ )

  if  $s$  is  $\square$ :  $U(s)$

  if  $s$  is  $\Delta$ : maxValue( $s$ , *depth*,  $\alpha$ ,  $\beta$ )

  if  $s$  is  $\nabla$ : minValue( $s$ , *depth*,  $\alpha$ ,  $\beta$ )

where  $\alpha, \beta$  are overall max and min values,  
resp.

# Formalize as Alpha-Beta Pruning

defun value( $s$ ):

  cutoff at depth  $m'$ : eval( $s$ )

  if  $s$  is  $\square$ :  $U(s)$

  if  $s$  is  $\Delta$ : maxValue( $s$ ,  $depth$ ,  $\alpha$ ,  $\beta$ )

  if  $s$  is  $\nabla$ : minValue( $s$ ,  $depth$ ,  $\alpha$ ,  $\beta$ )

where  $\alpha, \beta$  are overall max and min values,  
resp.

defun  $v = \text{maxValue}(s, depth, \alpha, \beta)$ :

$v = -\infty$

  for  $(a, s')$  in successors( $s$ ):

$v = \max(v, \text{value}(s', depth + 1, \alpha, \beta))$

    if  $v > \beta$  return  $v$

$\alpha = \max(\alpha, v)$



# Formalize as Alpha-Beta Pruning

defun value( $s$ ):

  cutoff at depth  $m'$ : eval( $s$ )

  if  $s$  is  $\square$ :  $U(s)$

  if  $s$  is  $\Delta$ :  $\text{maxValue}(s, \text{depth}, \alpha, \beta)$

  if  $s$  is  $\nabla$ :  $\text{minValue}(s, \text{depth}, \alpha, \beta)$

where  $\alpha, \beta$  are overall max and min values,  
resp.

defun  $v = \text{maxValue}(s, \text{depth}, \alpha, \beta)$ :

$v = -\infty$

  for  $(a, s')$  in successors( $s$ ):

$v = \max(v, \text{value}(s', \text{depth} + 1, \alpha, \beta))$

**if  $v > \beta$  return  $v$**

$\alpha = \max(\alpha, v)$

# Formalize as Alpha-Beta Pruning

defun value( $s$ ):

  cutoff at depth  $m'$ : eval( $s$ )

  if  $s$  is  $\square$ :  $U(s)$

  if  $s$  is  $\Delta$ : maxValue( $s$ ,  $depth$ ,  $\alpha$ ,  $\beta$ )

  if  $s$  is  $\nabla$ : minValue( $s$ ,  $depth$ ,  $\alpha$ ,  $\beta$ )

where  $\alpha, \beta$  are overall max and min values,  
resp.

defun  $v = \text{maxValue}(s, depth, \alpha, \beta)$ :

$v = -\infty$

  for  $(a, s')$  in successors( $s$ ):

$v = \max(v, \text{value}(s', depth + 1, \alpha, \beta))$

**if  $v > \beta$  return  $v$**

$\alpha = \max(\alpha, v)$

Can cut up to  $O(b^{m/2})!$

# Complexity Reduction by Tree $\rightarrow$ Graph

Convert into graph search problem:

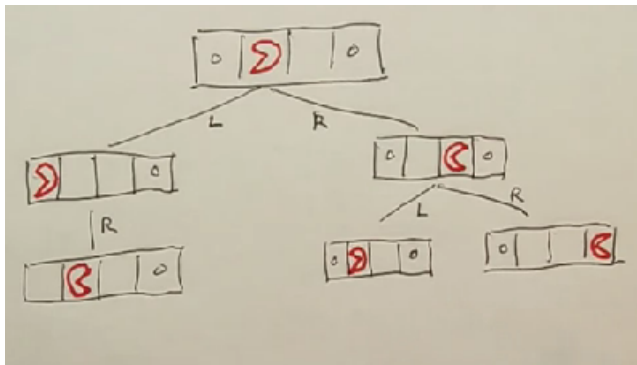
- to reach special **opening and closing states**
- to make and protect from **killer-moves**

# Complexity Reduction by Tree $\rightarrow$ Graph

Convert into graph search problem:

- to reach special **opening and closing states**
- to make and protect from **killer-moves**

Utility: How many food particles can pacman eat?

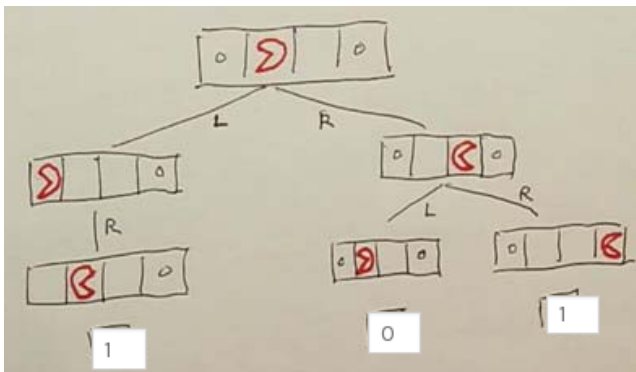


# Complexity Reduction by Tree $\rightarrow$ Graph

Convert into graph search problem:

- to reach special **opening and closing states**
- to make and protect from **killer-moves**

Utility: How many food particles can pacman eat?




# Complexity Reduction by Tree $\rightarrow$ Graph

Convert into graph search problem:

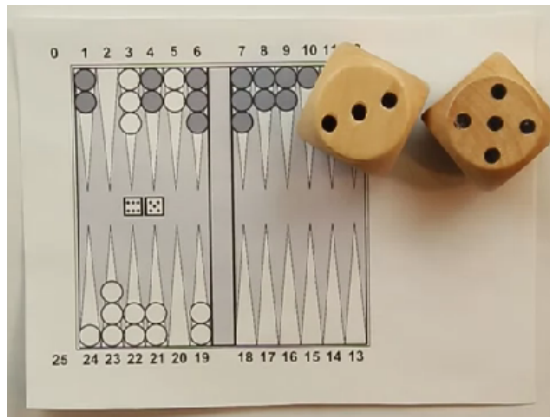
- to reach special **opening and closing states**
- to make and protect from **killer-moves**

Utility: How many food particles can pacman eat?



2-step limit causes **horizon effect**?

# Stochastic Games



# Stochastic Games

defun value( $s$ ):


  cutoff at depth  $m'$ : eval( $s$ )

  if  $s$  is  $\square$ :  $U(s)$

  if  $s$  is  $\Delta$ :  $\text{maxValue}(s, \text{depth}, \alpha, \beta)$

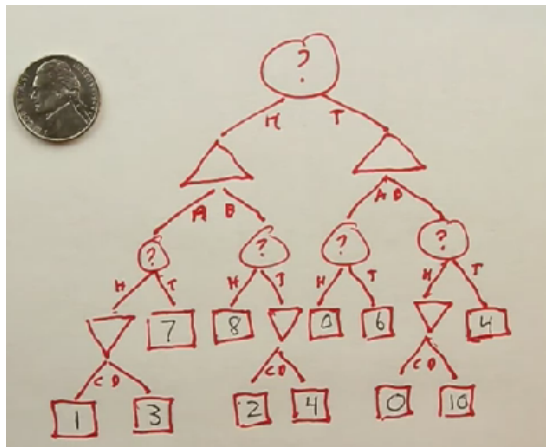
  if  $s$  is  $\nabla$ :  $\text{minValue}(s, \text{depth}, \alpha, \beta)$

  if  $s$  is  $?$ :  $\text{expValue}(s, \text{depth}, \alpha, \beta)$





# Coin-flip Game



# Coin-flip Game

