

CSE 473: Artificial Intelligence Spring 2012

Search

Dan Weld

With slides from
Dan Klein, Stuart Russell, Andrew Moore, Luke Zettlemoyer

Announcements

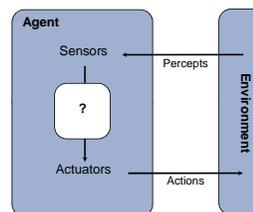
- Project 0: Python Tutorial
 - Online, but not graded
- Project 1: Search
 - On the web by tomorrow.
 - Start early and ask questions. **It's longer than most!**

Outline

- Agents that Plan Ahead
- Search Problems
- Uninformed Search Methods (part review for some)
 - Depth-First Search
 - Breadth-First Search
 - Uniform-Cost Search
- Heuristic Search Methods (new for all)
 - Best First / Greedy Search

Review: Rational Agents

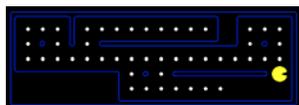
- An **agent** is an entity that *perceives* and *acts*.
- A **rational agent** selects actions that maximize its **utility function**.
- Characteristics of the **percepts, environment, and action space** dictate techniques for selecting rational actions.



Search -- the environment is:
fully observable, single agent, deterministic, episodic, discrete

Reflex Agents

- Reflex agents:
 - Choose action based on current percept (and maybe memory)
 - Do not consider the future consequences of their actions
 - **Act on how the world IS**
- Can a reflex agent be rational?
- Can a non-rational agent achieve goals?

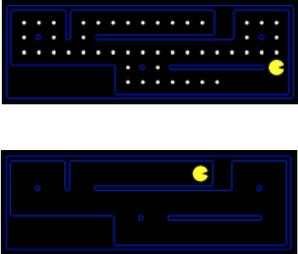


Famous Reflex Agents



Goal Based Agents

- Goal-based agents:
 - Plan ahead
 - Ask "what if"
 - Decisions based on (hypothesized) consequences of actions
 - Must have a model of how the world evolves in response to actions
 - Act on how the world **WOULD BE**



Search thru a Problem Space / State Space

- Input:
 - Set of states
 - Operators [and costs]
 - Start state
 - Goal state [test]
- Output:
 - Path: start \Rightarrow a state satisfying goal test
 - [May require shortest path]
 - [Sometimes just need state passing test]

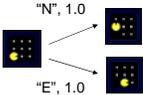
Example: Simplified Pac-Man

- Input:
 - A state space
 - A successor function
 - A start state
 - A goal test
- Output:



"N", 1.0

"E", 1.0



Ex: Route Planning: Romania \rightarrow Bucharest

- Input:
 - Set of states
 - Operators [and costs]
 - Start state
 - Goal state (test)
- Output:



Example: N Queens

- Input:
 - Set of states
 - Operators [and costs]
 - Start state
 - Goal state (test)
- Output

		Q	
Q			
			Q
	Q		

Ex: Blocks World

- Input:
 - Set of states
 - Partially specified plans
 - Operators [and costs]
 - Plan modification operators
 - Start state
 - The null plan (no actions)
 - Goal state (test)
 - A plan which provably achieves
 - The desired world configuration
- Output:



Multiple Problem Spaces

Real World

States of the world (e.g. block configurations)

Actions (take one world-state to another)



Robot's Head

- Problem Space 1
 - PS states =
 - models of world states
 - Operators =
 - models of actions

- Problem Space 2
 - PS states =
 - partially spec. plan
 - Operators =
 - plan modificat'n ops



Algebraic Simplification

$$\partial_t^2 u = - \left[E' - \frac{l(l+1)}{r^2} - \mu^2 \right] u(r)$$

$$e^{-2\mu} (\partial_t^2 - \partial_t) u(s) = - [E' - l(l+1)e^{-2\mu} - \mu^2] u(s)$$

$$e^{-2\mu} \left[e^{2\mu} (e^{-2\mu} u(s))'' - \frac{1}{s} u \right] = - [E' - l(l+1)e^{-2\mu} - \mu^2] u(s)$$

$$e^{-2\mu} \left[e^{2\mu} (e^{-2\mu} u(s))'' \right] = - \left[E' - \left(l + \frac{1}{2} \right)^2 e^{-2\mu} - \mu^2 \right] u(s)$$

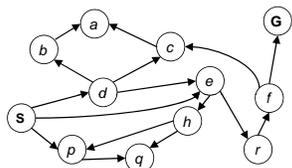
$$v'' = - e^{2\mu} \left[E' - \left(l + \frac{1}{2} \right)^2 e^{-2\mu} - \mu^2 \right] v$$

- **Input:**
 - Set of states
 - Operators [and costs]
 - Start state
 - Goal state (test)
- **Output:**

14

State Space Graphs

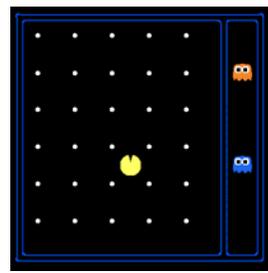
- State space graph:
 - Each node is a state
 - The successor function is represented by arcs
 - Edges may be labeled with costs
- We can rarely build this graph in memory (so we don't)



Ridiculously tiny search graph for a tiny search problem

State Space Sizes?

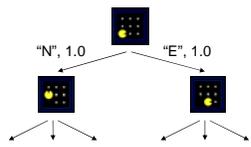
- Search Problem: Eat all of the food
- Pacman positions: 10 x 12 = 120
- Pacman facing: up, down, left, right
- Food Count: 30
- Ghost positions: 12



Search Strategies

- **Blind Search**
 - Depth first search
 - Breadth first search
 - Iterative deepening search
 - Uniform cost search
- **Informed Search**
- **Constraint Satisfaction**
- **Adversary Search**

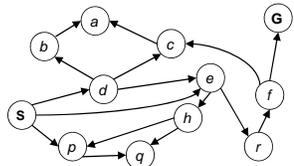
Search Trees



- A search tree:
 - Start state at the root node
 - Children correspond to successors
 - Nodes contain states, correspond to PLANS to those states
 - Edges are labeled with actions and costs
 - For most problems, we can never actually build the whole tree

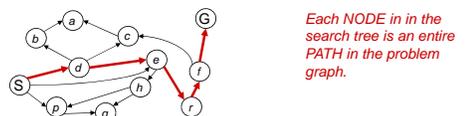
Example: Tree Search

State Graph:



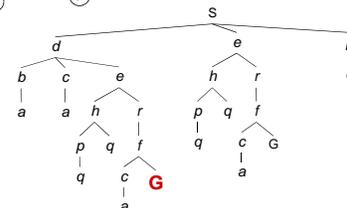
What is the search tree?

State Graphs vs. Search Trees

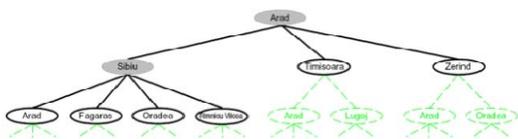


Each NODE in the search tree is an entire PATH in the problem graph.

We construct both on demand – and we construct as little as possible.



Building Search Trees



- Search:
 - Expand out possible plans
 - Maintain a **fringe** of unexpanded plans
 - Try to expand as few tree nodes as possible

General Tree Search

```

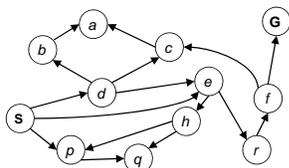
function TREE-SEARCH( problem, strategy) returns a solution, or failure
  initialize the search tree using the initial state of problem
  loop do
    if there are no candidates for expansion then return failure
    choose a leaf node for expansion according to strategy
    if the node contains a goal state then return the corresponding solution
    else expand the node and add the resulting nodes to the search tree
  end
    
```

- Important ideas:
 - Fringe
 - Expansion
 - Exploration strategy
- Main question: which fringe nodes to explore?

Detailed pseudocode is in the book!

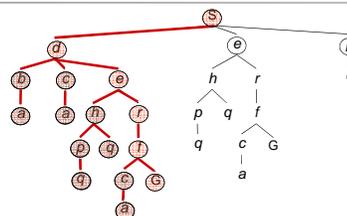
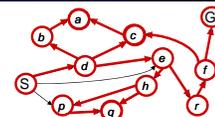
Review: Depth First Search

Strategy: expand deepest node first
Implementation: Fringe is a LIFO queue (a stack)



Review: Depth First Search

Expansion ordering:
 (d,b,a,c,a,e,h,p,q,q,r,f,c,a,G)



Review: Breadth First Search

Strategy: expand shallowest node first

Implementation: Fringe is a FIFO queue

Review: Breadth First Search

Expansion order:
 (S, d, e, p, b, c, e, h, r, q, a, a, h, r, p, q, f, p, q, f, q, c, G)

Search Algorithm Properties

- **Complete?** Guaranteed to find a solution if one exists?
- **Optimal?** Guaranteed to find the least cost path?
- **Time complexity?**
- **Space complexity?**

Variables:

n	Number of states in the problem
b	The maximum branching factor B (the maximum number of successors for a state)
C^*	Cost of least cost solution
d	Depth of the shallowest solution
m	Max depth of the search tree

DFS

Algorithm	Complete	Optimal	Time	Space
DFS Depth First Search	No	No	Infinite	Infinite

- Infinite paths make DFS incomplete...
 - How can we fix this?
 - Check new nodes against path from S
- Infinite search spaces still a problem

DFS

Algorithm	Complete	Optimal	Time	Space
DFS w/ Path Checking	Y if finite	N	$O(b^m)$	$O(bm)$

* Or graph search – next lecture.

BFS

Algorithm	Complete	Optimal	Time	Space
DFS w/ Path Checking	Y	N	$O(b^m)$	$O(bm)$
BFS	Y	Y*	$O(b^d)$	$O(b^d)$

Memory a Limitation?

- **Suppose:**
 - 4 GHz CPU
 - 6 GB main memory
 - 100 instructions / expansion
 - 5 bytes / node
- 400,000 expansions / sec
 - Memory filled in 300 sec ... 5 min

Comparisons

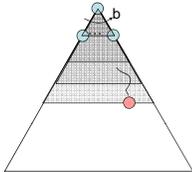
- When will BFS outperform DFS?
- When will DFS outperform BFS?

Iterative Deepening

Iterative deepening uses DFS as a subroutine:

1. Do a DFS which only searches for paths of length 1 or less.
2. If "1" failed, do a DFS which only searches paths of length 2 or less.
3. If "2" failed, do a DFS which only searches paths of length 3 or less.

...and so on.



Algorithm	Complete	Optimal	Time	Space
DFS <small>w/ Path Checking</small>	Y	N	$O(b^m)$	$O(bm)$
BFS	Y	Y*	$O(b^d)$	$O(b^d)$
ID	Y	Y*	$O(b^d)$	$O(bd)$

Cost of Iterative Deepening

b	ratio ID to DFS
2	3
3	2
5	1.5
10	1.2
25	1.08
100	1.02

Speed

Assuming 10M nodes/sec & sufficient memory

	BFS		Iter. Deep.	
	Nodes	Time	Nodes	Time
8 Puzzle	10^5	.01 sec	10^5	.01 sec
2x2x2 Rubik's	10^6	.2 sec	10^6	.2 sec
15 Puzzle	10^{13}	6 days	10^{17}	20k yrs
3x3x3 Rubik's	10^{19}	68k yrs	10^{20}	574k yrs
24 Puzzle	10^{25}	12B yrs	10^{37}	10^{23} yrs

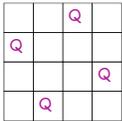
Why the difference?
Rubik has higher branch factor
15 puzzle has greater depth

of duplicates

Slide adapted from Richard Korf presentation

When to Use Iterative Deepening

- N Queens?



© Daniel S. Weld 37