6 common properties to distinguish tasks (not exhaustive)

- Fully observable vs Partially observable
- Single agent vs Multiagent
- Deterministic vs Stochastic
- Episodic vs Sequential
- Static vs Dynamic
- Discrete vs Continuous

- Col 1 Poker
- Col 2 Self-driving taxi
- Col 3 Spam classifier
- Col 4 Pacman with ghosts
- Col 5 Oil refinery control system
- Col 6 Automatic speech transcription

Search Problems



Search?



Information Retrieval vs Search









Definition of Search

Finding a (best) sequence of actions to solve a problem

For now, assume the problem is

- Deterministic
- Fully observable
- Known

Search Problem Mechanics

- A search problem consists of:
 - A state space



• A successor function (with actions, costs)



- A start state and a goal test
- A solution is a sequence of actions (a plan) which transforms the start state to a goal state

Example: Traveling in Romania



- State space:
 - Cities
- Successor function:
 - Roads: Go to adjacent city with cost = distance
- Start state:
 - Arad
- Goal test:
 - Is state == Bucharest?
- Solution?

What's in a State Space?



A search state keeps only the details needed for planning (abstraction)

- Problem: Pathing
 - States: (x,y) location
 - Actions: NSEW
 - Successor: update location only
 - Goal test: is (x,y)=END

- Problem: Eat-All-Dots
 - States: {(x,y), dot booleans}
 - Actions: NSEW
 - Successor: update location and possibly a dot boolean
 - Goal test: dots all false

State Space Sizes?

- World state:
 - Agent positions: 120
 - Food count: 30
 - Ghost positions: 12
 - Agent facing: NSEW
- How many
 - World states? 120x(2³⁰)x(12²)x4
 - States for pathing?
 120
 - States for eat-all-dots? 120x(2³⁰)



Search Problem Mechanics

- A search problem consists of:
 - A state space



- A successor function (with actions, costs)
- A start state and a goal test



What are some problems that <u>can't</u> be formulated as search?

 A solution is a sequence of actions (a plan) which transforms the start state to a goal state

State Space Graphs and Search Trees



State Space Graphs

- State space graph: A mathematical representation of a search problem
 - Nodes are (abstracted) world configurations
 - Arcs represent successors (action results)
 - The goal test is a set of goal nodes (maybe only one)
- In a state space graph, each state occurs only once!
- We can rarely build this full graph in memory (it's too big), but it's a useful idea



State Space Graphs

- State space graph: A mathematical representation of a search problem
 - Nodes are (abstracted) world configurations
 - Arcs represent successors (action results)
 - The goal test is a set of goal nodes (maybe only one)
- In a state space graph, each state occurs only once!
- We can rarely build this full graph in memory (it's too big), but it's a useful idea



Tiny search graph for a tiny search problem

Search Trees



- A search tree:
 - A "what if" tree of plans and their outcomes
 - The start state is the root node
 - Children correspond to successors
 - Nodes show states, but correspond to ACTION SEQUENCES that achieve those states
 - For most problems, we can never actually build the whole tree

State Space Graphs vs. Search Trees



Each NODE in in the search tree corresponds to an entire PATH in the state space graph.

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



Quiz: State Space Graphs vs. Search Trees

Consider this 4-state graph:

How big is its search tree (from S)?





Important: Lots of repeated structure in the search tree!

Tree Search



Search Example: Romania



Searching with a Search Tree



- Search:
 - Expand out potential plans (tree nodes)
 - Maintain a fringe of partial plans under consideration
 - 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

Search Algorithm Properties

- Complete: Guaranteed to find a solution if one exists?
- Optimal: Guaranteed to find the least cost path?
- Time complexity?
- Space complexity?
- Cartoon of search tree:
 - b is the branching factor
 - m is the maximum depth
 - solutions at various depths
- Number of nodes in entire tree?
 - $1 + b + b^2 + \dots b^m = O(b^m)$



Breadth-First Search



Breadth-First Search

Strategy: expand a shallowest node first

Implementation: Fringe is a FIFO queue





Breadth-First Search (BFS) Properties

- What nodes does BFS expand?
 - Processes all nodes above shallowest solution
 - Let depth of shallowest solution be s
 - Search takes time O(b^s)
- How much space does the fringe take?
 - Has roughly the last tier, so O(b^s)
- Is it complete?
 - s must be finite if a solution exists, so yes!
- Is it optimal?
 - Only if costs are all 1 (more on costs later)



Python code for BFS

```
import collections
def bfs(graph, root):
    seen, queue = set([root]), collections.deque([root])
   while queue:
        vertex = queue.popleft()
        visit(vertex)
        for node in graph[vertex]:
            if node not in seen:
                seen.add(node)
                queue.append(node)
def visit(n):
    print(n)
if __name__ == '__main__':
    graph = {0: [1, 2], 1: [2, 0], 2: []}
    bfs(graph, 0)
```

Next time

Other search strategies