

PROBLEM SOLVING AND SEARCH

CHAPTER 3

Chapter 3 1

Outline

- ◇ Problem-solving agents
- ◇ Problem types
- ◇ Problem formulation
- ◇ Example problems
- ◇ Basic search algorithms

Example: Romania

On holiday in Romania; currently in Arad.
Flight leaves tomorrow from Bucharest

Formulate goal:

be in Bucharest

Formulate problem:

states: various cities

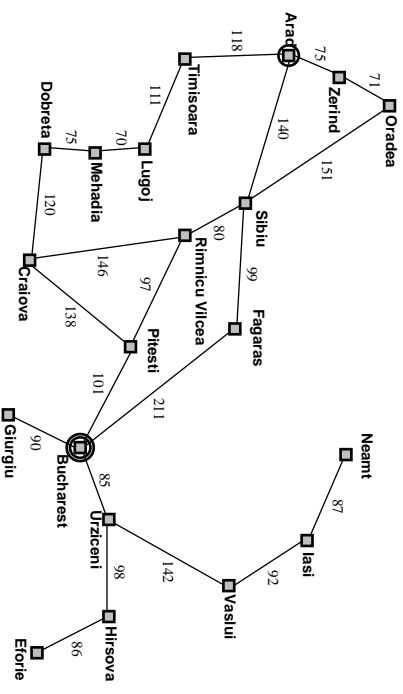
actions: drive between cities

Find solution:

sequence of cities, e.g., Arad, Sibiu, Fagaras, Bucharest

Chapter 3 4

Example: Romania



Chapter 3 2

Chapter 3 5

Problem-solving agents

Restricted form of general agent:

function SIMPLE-PROBLEM-SOLVING-AGENT (*percept*) **returns** an action

static: seq, an action sequence, initially empty

state, some description of the current world state

goal, a goal, initially null

problem, a problem formulation

state ← UPDATE-STATE(*state*, *percept*)

if *seq* is empty **then**

goal ← FORMULATE-GOAL(*state*)

problem ← FORMULATE-PROBLEM(*state*, *goal*)

seq ← SEARCH(*problem*)

action ← FIRST(*seq*)

seq ← REST(*seq*)

return *action*

Note: this is offline problem solving; solution executed "eyes closed."

Online problem solving involves acting without complete knowledge.

Chapter 3 3

Problem types

Deterministic, fully observable \implies single-state problem

Agent knows exactly which state it will be in; solution is a sequence

Non-observable \implies conformant problem

Agent may have no idea where it is; solution (if any) is a sequence

Nondeterministic and/or partially observable \implies contingency problem

percepts provide **new** information about current state

solution is a contingent plan or a policy


often **interleave** search, execution

Unknown state space \implies exploration problem ("online")

Chapter 3 6


Example: vacuum world

Single-state, start in #5. **Solution??**



Example: vacuum world

Single-state, start in #5. **Solution??**




Example: vacuum world

Single-state, start in #5. **Solution??**

[Right, Suck]

Conformant, start in $\{1, 2, 3, 4, 5, 6, 7, 8\}$

e.g., Right goes to $\{2, 4, 6, 8\}$. **Solution??**



Single-state problem formulation

A problem is defined by four items:

initial state e.g., "at Arad"

successor function $S(x)$ = set of action–state pairs

e.g., $S(\text{Arad}) = \{(\text{Arad} \rightarrow \text{Zerind}, \text{Zerind}), \dots\}$

goal test, can be

explicit, e.g., $x = \text{"at Bucharest"}$

implicit, e.g., $\text{NoDirt}(x)$

path cost (additive)

e.g., sum of distances, number of actions executed, etc.

$c(x, a, y)$ is the step cost, assumed to be ≥ 0

A solution is a sequence of actions leading from the initial state to a goal state

Example: vacuum world

Single-state, start in #5. **Solution??**

[Right, Suck]

Conformant, start in $\{1, 2, 3, 4, 5, 6, 7, 8\}$

e.g., Right goes to $\{2, 4, 6, 8\}$. **Solution??**


[Right, Suck, Left, Suck]

Contingency, start in #5 or #7

Murphy's Law: Suck can dirty a clean carpet

Local sensing: dirt, location only.

Solution??



Selecting a state space

Real world is absurdly complex

\Rightarrow state space must be **abstracted** for problem solving

(Abstract) state = set of real states

(Abstract) action = complex combination of real actions

e.g., "Arad \rightarrow Zerind" represents a complex set of possible routes, detours, rest stops, etc.


For guaranteed realizability, **any** real state "in Arad" must get to some real state "in Zerind"

(Abstract) solution =

set of real paths that are solutions in the real world


Each abstract action should be "easier" than the original problem!

Example: vacuum world state space graph




- states??
- actions??
- goal test??
- path cost??

Example: vacuum world state space graph




- states??: integer dirt and robot locations (ignore dirt amounts etc.)
- actions??: *Left, Right, Suck, NoOp*
- goal test??: no dirt
- path cost??

Example: vacuum world state space graph




- states??: integer dirt and robot locations (ignore dirt amounts etc.)
- actions??
- goal test??
- path cost??

Example: vacuum world state space graph




- states??: integer dirt and robot locations (ignore dirt amounts etc.)
- actions??: *Left, Right, Suck, NoOp*
- goal test??: no dirt
- path cost??: 1 per action (0 for *NoOp*)

Example: vacuum world state space graph



- states??: integer dirt and robot locations (ignore dirt amounts etc.)
- actions??: *Left, Right, Suck, NoOp*
- goal test??
- path cost??

Example: The 8-puzzle



- states??
- actions??
- goal test??
- path cost??

Example: The 8-puzzle

Start State

7	2	4
5	6	6
8	3	1

Goal State

1	2	3
4	5	6
7	8	

states??: integer locations of tiles (ignore intermediate positions)

actions??:

goal test??:

path cost??:

Example: The 8-puzzle

Start State

7	2	4
5	6	6
8	3	1

Goal State

1	2	3
4	5	6
7	8	

states??: integer locations of tiles (ignore intermediate positions)

actions??: move blank left, right, up, down (ignore unjamming etc.)

goal test??: = goal state (given)

path cost??: 1 per move

[Note: optimal solution of n -Puzzle family is NP-hard]

Example: The 8-puzzle

Start State

7	2	4
5	6	6
8	3	1

Goal State

1	2	3
4	5	6
7	8	


states??: integer locations of tiles (ignore intermediate positions)

actions??: move blank left, right, up, down (ignore unjamming etc.)

goal test??:


path cost??:

Example: robotic assembly



states??:

Example: robotic assembly



states??: real-valued coordinates of robot joint angles

actions??: parts of the object to be assembled

actions??:


states??: integer locations of tiles (ignore intermediate positions)

actions??: move blank left, right, up, down (ignore unjamming etc.)

goal test??: = goal state (given)

path cost??:

Example: robotic assembly



states??: real-valued coordinates of robot joint angles
parts of the object to be assembled

actions??: continuous motions of robot joints

goal test??:

Tree search algorithms

Basic idea:

offline, simulated exploration of state space
by generating successors of already-explored states
(a.k.a. **expanding states**)

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

Example: robotic assembly




states??: real-valued coordinates of robot joint angles
parts of the object to be assembled

actions??: continuous motions of robot joints


goal test??: complete assembly **with no robot included!**

path cost??:

Tree search example



Example: robotic assembly




states??: real-valued coordinates of robot joint angles
parts of the object to be assembled

actions??: continuous motions of robot joints


goal test??: complete assembly **with no robot included!**

path cost??: time to execute

Tree search example



Tree search example




Chapter 3 31

Breadth-first search

Expand shallowest unexpanded node

Implementation:
fringe is a FIFO queue, i.e., new successors go at end



Chapter 3 34

Search strategies

A strategy is defined by picking the **order of node expansion**

Strategies are evaluated along the following dimensions:

- completeness**—does it always find a solution if one exists?
- time complexity**—number of nodes generated/expanded
- space complexity**—maximum number of nodes in memory
- optimality**—does it always find a least-cost solution?

Time and space complexity are measured in terms of

- b —maximum branching factor of the search tree
- d —depth of the least-cost solution
- m —maximum depth of the state space (may be ∞)

Chapter 3 32

Uninformed search strategies

Uninformed strategies use only the information available in the problem definition


- Breadth-first search
- Uniform-cost search
- Depth-first search
- Depth-limited search
- Iterative deepening search

Chapter 3 33

Breadth-first search

Expand shallowest unexpanded node

Implementation:
fringe is a FIFO queue, i.e., new successors go at end




Chapter 3 35

Breadth-first search

Expand shallowest unexpanded node

Implementation:
fringe is a FIFO queue, i.e., new successors go at end




Chapter 3 36

Breadth-first search

Expand shallowest unexpanded node

Implementation:

fringe is a FIFO queue, i.e., new successors go at end



Chapter 3 37

Properties of breadth-first search

Complete??? Yes (if b is finite)

Time?? $1 + b + b^2 + b^3 + \dots + b^d + b(b^{d-1} - 1) = O(b^{d+1})$, i.e., exp. in d

Space??

Chapter 3 40

Properties of breadth-first search

Complete???

Properties of breadth-first search

Complete??? Yes (if b is finite)

Time?? $1 + b + b^2 + b^3 + \dots + b^d + b(b^{d-1} - 1) = O(b^{d+1})$, i.e., exp. in d

Space?? $O(b^{d+1})$ (keeps every node in memory)

Optimal???

Chapter 3 38

Chapter 3 41

Properties of breadth-first search

Complete??? Yes (if b is finite)

Time???

Properties of breadth-first search

Complete??? Yes (if b is finite)

Time?? $1 + b + b^2 + b^3 + \dots + b^d + b(b^{d-1} - 1) = O(b^{d+1})$, i.e., exp. in d

Space?? $O(b^{d+1})$ (keeps every node in memory)

Optimal?? Yes (if cost = 1 per step); not optimal in general

Space is the big problem: can easily generate nodes at 100MB/sec
so 24hrs = 8640GB.

Chapter 3 39

Chapter 3 42

Uniform-cost search

Expand least-cost unexpanded node

Implementation:

fringe = queue ordered by path cost, lowest first

Equivalent to breadth-first if step costs all equal

Complete?? Yes, if step cost $\geq \epsilon$

Time?? # of nodes with $g \leq$ cost of optimal solution, $O(b^{(C^*/\epsilon)})$
where C^* is the cost of the optimal solution

Space?? # of nodes with $g \leq$ cost of optimal solution, $O(b^{(C^*/\epsilon)})$


Optimal?? Yes—nodes expanded in increasing order of $g(n)$

Depth-first search

Expand deepest unexpanded node

Implementation:

fringe = LIFO queue, i.e., put successors at front




Depth-first search

Expand deepest unexpanded node

Implementation:

fringe = LIFO queue, i.e., put successors at front




Depth-first search

Expand deepest unexpanded node

Implementation:

fringe = LIFO queue, i.e., put successors at front




Depth-first search

Expand deepest unexpanded node

Implementation:

fringe = LIFO queue, i.e., put successors at front




Depth-first search

Expand deepest unexpanded node

Implementation:

fringe = LIFO queue, i.e., put successors at front




Depth-first search

Expand deepest unexpanded node

Implementation:

fringe = LIFO queue, i.e., put successors at front




Chapter 3 49

Depth-first search

Expand deepest unexpanded node

Implementation:

fringe = LIFO queue, i.e., put successors at front




Chapter 3 52

Depth-first search

Expand deepest unexpanded node

Implementation:

fringe = LIFO queue, i.e., put successors at front




Chapter 3 56

Depth-first search

Expand deepest unexpanded node

Implementation:

fringe = LIFO queue, i.e., put successors at front




Chapter 3 53

Depth-first search

Expand deepest unexpanded node

Implementation:

fringe = LIFO queue, i.e., put successors at front




Chapter 3 51

Depth-first search

Expand deepest unexpanded node

Implementation:

fringe = LIFO queue, i.e., put successors at front




Chapter 3 54

Depth-first search

Expand deepest unexpanded node

Implementation:

fringe = LIFO queue, i.e., put successors at front



Chapter 3 55

Properties of depth-first search

Complete?? No: fails in infinite-depth spaces, spaces with loops

Modify to avoid repeated states along path
⇒ complete in finite spaces

Time?? $O(b^m)$: terrible if m is much larger than d

but if solutions are dense, may be much faster than breadth-first

Space??

Chapter 3 56

Properties of depth-first search

Complete??

Properties of depth-first search

Complete?? No: fails in infinite-depth spaces, spaces with loops

Modify to avoid repeated states along path
⇒ complete in finite spaces

Time?? $O(b^m)$: terrible if m is much larger than d

but if solutions are dense, may be much faster than breadth-first

Space?? $O(bm)$, i.e., linear space!

Optimal??

Chapter 3 56

Chapter 3 56

Properties of depth-first search

Complete?? No: fails in infinite-depth spaces, spaces with loops

Modify to avoid repeated states along path
⇒ complete in finite spaces

Time??

Properties of depth-first search

Complete?? No: fails in infinite-depth spaces, spaces with loops

Modify to avoid repeated states along path
⇒ complete in finite spaces

Time?? $O(b^m)$: terrible if m is much larger than d

but if solutions are dense, may be much faster than breadth-first

Space?? $O(bm)$, i.e., linear space!

Optimal?? No

Chapter 3 57

Chapter 3 56

Depth-limited search

= depth-first search with depth limit l ,
i.e., nodes at depth l have no successors

Recursive implementation:


```

function DEPTH-LIMITED-SEARCH(problem, limit) returns soln/fail/cutoff
  RECURSIVE-DLS(MAKE-NODE(INITIAL-STATE(problem)), problem, limit)
function RECURSIVE-DLS(node, problem, limit) returns soln/fail/cutoff
  cutoff-occurred?  $\leftarrow$  false
  if GOAL-TEST(problem, STATE(node)) then return node
  else if DEPTH(node) = limit then return cutoff
  else for each successor in EXPAND(node, problem) do
    result  $\leftarrow$  RECURSIVE-DLS(successor, problem, limit)
    if result = cutoff then cutoff-occurred?  $\leftarrow$  true
    else if result  $\neq$  failure then return result
  if cutoff-occurred? then return cutoff else return failure
  
```

Iterative deepening search $l = 1$



Iterative deepening search $l = 2$



Iterative deepening search


```

function ITERATIVE-DEEPENING-SEARCH(problem) returns a solution
  inputs: problem, a problem
  for depth  $\leftarrow$  0 to  $\infty$  do
    result  $\leftarrow$  DEPTH-LIMITED-SEARCH(problem, depth)
    if result  $\neq$  cutoff then return result
  end
  
```

Iterative deepening search $l = 0$



Iterative deepening search $l = 3$



Properties of iterative deepening search

Complete??

Properties of iterative deepening search

Complete?? Yes

Time?? $(d + 1)b^0 + db^1 + (d - 1)b^2 + \dots + b^d = O(b^d)$

Space?? $O(bd)$

Optimal??

Properties of iterative deepening search

Complete?? Yes

Time??

Properties of iterative deepening search

Complete?? Yes

Time?? $(d + 1)b^0 + db^1 + (d - 1)b^2 + \dots + b^d = O(b^d)$

Space?? $O(bd)$

Optimal?? Yes, if step cost = 1

Can be modified to explore uniform-cost tree

Properties of iterative deepening search

Complete?? Yes

Time?? $(d + 1)b^0 + db^1 + (d - 1)b^2 + \dots + b^d = O(b^d)$

Space??

Summary of algorithms

Criterion	Breadth-First	Uniform-Cost	Depth-First	Depth-Limited	Iterative Deepening
Complete?	Yes*	Yes*	No	Yes, if $l \geq d$	Yes
Time	b^{d+1}	$b^{\lceil C^*/\epsilon \rceil}$	b^m	b^l	b^d
Space	b^{d+1}	$b^{\lceil C^*/\epsilon \rceil}$	bm	b	bd
Optimal?	Yes*	Yes	No	No	Yes*



Repeated state checking

Depth-first search: Is checking current node w.r.t. path stored in memory enough?
i.e. Is linear space sufficient?

Chapter 3 73

Repeated states

Failure to detect repeated states can turn a linear problem into an exponential one!



Chapter 3 74

Repeated state checking

Depth-first search: Is checking current node w.r.t. path stored in memory enough?
i.e. Is linear space sufficient?

No! Can only detect looping paths, not all repeated states.

Need exponential space to store all visited nodes.

Chapter 3 75

Graph search

```
function GRAPH-SEARCH(problem, fringe) returns a solution, or failure
  closed ← an empty set
  fringe ← INSERT(MAKE-NODE(INITIAL-STATE[problem]), fringe)
  loop do
    if fringe is empty then return failure
    node ← REMOVE-FRONT(fringe)
    if GOAL-TEST(problem, STATE[node]) then return node
    if STATE[node] is not in closed then
      add STATE[node] to closed
      fringe ← INSERT-ALL(EXPAND(node, problem), fringe)
  end
```

Is this optimal?

- BFS in INSERTALL
- DFS in INSERTALL

Chapter 3 76

Summary

Problem formulation usually requires abstracting away real-world details to define a state space that can feasibly be explored

Variety of uninformed search strategies

Iterative deepening search uses only linear space and not much more time than other uninformed algorithms

Graph search can be exponentially more efficient than tree search

Chapter 3 77

Complexity of BFS and DFS

Complexity of BFS and DFS is linear in the number of states V

In particular, Dijkstra's algorithm for single source shortest paths is $\Theta(E + V \log V)$, i.e. polynomial in V

However, V is b^m in many cases.

e.g. chess, theorem proving, scheduling problems

Chapter 3 78