

**CS 4700:
Foundations of Artificial Intelligence**

Bart Selman

Search Techniques

R&N: Chapter 3

Search: tree search and graph search

Uninformed search: very briefly (covered before in other prerequisite courses – recommendation: review these techniques at home)

Informed search ← focus of the lecture

Uninformed search strategies

Uninformed (blind) search 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**
- **Bidirectional search**
-

Key issue: type of queue used for the **fringe of the search tree** (collection of tree nodes that have been generated but not yet expanded)

Tree-search algorithms

Searching for a (shortest / least cost) path to goal state(s).

Search through the state space.

We will consider search techniques that use an explicit **search tree** that is generated by the **initial state + successor function**.

initialize (initial node)



Loop

choose a node for expansion
according to a **strategy**

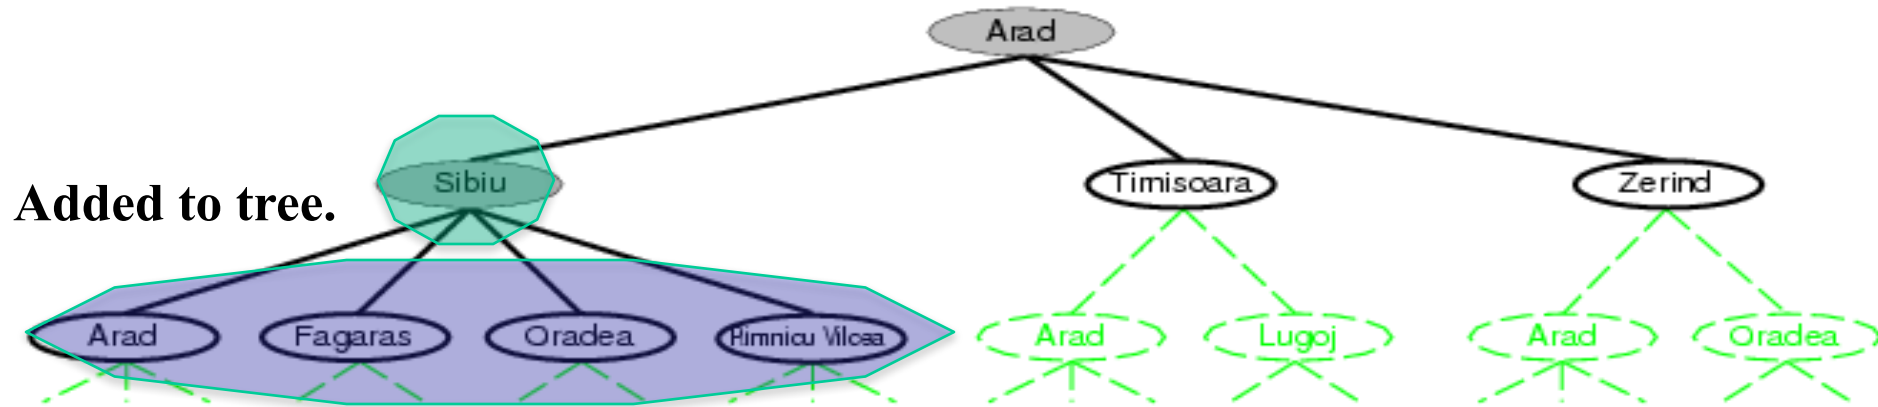
goal node? → done

expand node with successor function

Nodes added to tree.

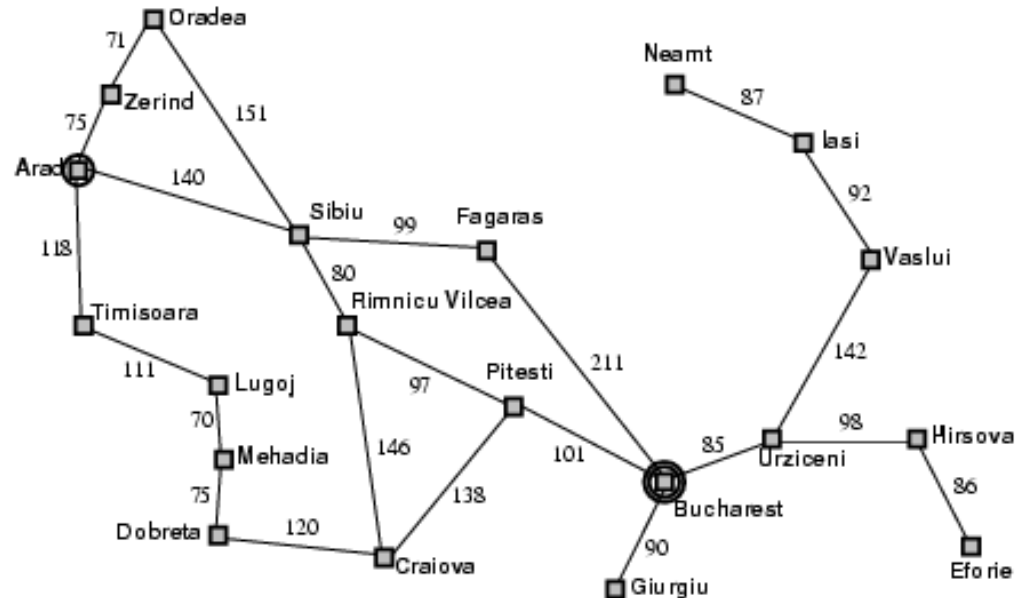


Selected for expansion.



Note: Arad added (again) to tree!
(reachable from Sibiu)

Not necessarily a problem, but
in **Graph-Search**, we will avoid
this by maintaining an
“explored” list.



Tree-search algorithms

Basic idea:

- 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
```

Fig. 3.7 R&N, p. 77

Note: 1) Here we only check a node for possibly being a goal state, *after* we select the node for expansion.

2) A “node” is a data structure containing state + additional info (parent node, etc.

Graph-search

```
function GRAPH-SEARCH(problem) returns a solution, or failure
  initialize the frontier using the initial state of problem
  initialize the explored set to be empty
  loop do
    if the frontier is empty then return failure
    choose a leaf node and remove it from the frontier
    if the node contains a goal state then return the corresponding solution
    add the node to the explored set
    expand the chosen node, adding the resulting nodes to the frontier
      only if not in the frontier or explored set
```

Fig. 3.7 R&N, p. 77. See also exercise 3.13.

Note:

- 1) Uses “explored” set to avoid visiting already explored states.
- 2) Uses “frontier” set to store states that remain to be explored and expanded.
- 3) *However, with eg uniform cost search, we need to make a special check when node (i.e. state) is on frontier. Details later.*

Search strategies

A search *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
- **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 ∞)
-

Uninformed search strategies

Uninformed (blind) search 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**
- **Bidirectional search**
-


Key issue: type of queue used for the **fringe of the search tree** (collection of tree nodes that have been generated but not yet expanded)

Breadth-first search

Expand shallowest unexpanded node.

Implementation:


- *fringe* is a FIFO queue, i.e., new nodes go at end **Fringe queue: <A>**
(First In First Out queue.)



Select A from queue and expand.

Gives
<B, C>

Queue: <B, C>




Select B from front, and expand.


Put children at the end.

**Gives
<C, D, E>**

Fringe queue: <C, D, E>



Fringe queue: <D, E, F, G>



Assuming no further children, queue becomes <E, F, G>, <F, G>, <G>, <>. Each time node checked for goal state.

Properties of breadth-first search

(check them at home!!!!)

Note: check for goal only when node is expanded.

Complete? Yes (if b is finite)

Time? $1 + b + b^2 + b^3 + \dots + b^d + b(b^d - 1) = O(b^{d+1})$ Depth d , goal may be last node (only checked when expanded).

Space? $O(b^{d+1})$ (keeps every node in memory; needed also to reconstruct soln. path expanded.).

Optimal soln. found?

Yes (if all step costs are identical)

Space is the bigger problem (more than time)

b : maximum branching factor of the search tree

d : depth of the least-cost solution

Uniform Cost

We factor in the **cost of each step** (e.g., distance from current state to the neighbors). Assumption: costs are non-negative.

$$g(n) = \text{cost so far to reach } n$$

Queue → ordered by cost

If all the steps are the same → **breadth-first search** is optimal since it always expands the **shallowest (least cost)**!

Uniform-cost search → expand first the nodes with lowest cost (instead of depth).

Does it ring a bell?

Uniform-cost search

Two subtleties: (bottom p. 83 Norvig)

- 1) Do goal state test, only when a node is selected for expansion.
(Reason: Bucharest may occur on frontier with a longer than optimal path. It won't be selected for expansion yet. Other nodes will be expanded first, leading us to uncover a shorter path to Bucharest. See also point 2).
- 2) Graph-search alg. says “don't add child node to frontier if **already on explored list or already on frontier.**” BUT, child may give a shorter path to a state already on frontier. Then, we need to modify the existing node on frontier with the shorter path. See fig. 3.14 (else-if part).

Uniform-cost search

Expand **least-cost** (of path to) unexpanded node
(e.g. useful for finding shortest path on map)

Implementation:

– *fringe* = queue **ordered by path cost**

– g – cost of reaching a node

Complete? Yes, if step cost $\geq \epsilon$ (>0)

Time? # of nodes with $g \leq$ cost of optimal solution (C^*),
 $O(b^{(1+\lceil C^*/\epsilon \rceil)})$ **Note: search is guided by cost not depth!!!**

Space? # of nodes with $g \leq$ cost of optimal solution,
 $O(b^{(1+\lceil C^*/\epsilon \rceil)})$ **Note: search is guided by cost not depth!!!**

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

Note: Some subtleties (e.g. checking for goal state).


See p 84 R&N. Also, next slide.

Depth-first search

“Expand deepest unexpanded node”

Implementation:

- *fringe* = LIFO queue, i.e., put successors at front (“push on stack”)
Last In First Out



Fringe stack:


A

Expanding A,
gives stack:

B

C


So, B next.



**Expanding B,
gives stack:**

**D
E
C**


So, D next.





**Expanding D,
gives stack:**


**H
I
E
C**


**So, H next.
etc.**




















What is main advantage over breadth first search?

What information is stored? How much storage required?

The stack. $O(\text{depth} \times \text{branching})$.

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

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

Note: Can also reconstruct soln. path from single stored branch.

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

Guarantee that opt. soln. is found? No

b : max. branching factor of the search tree
 d : depth of the shallowest (least-cost) soln.
 m : maximum depth of state space

Note: In “backtrack search” only one successor is generated
→ only $O(m)$ memory is needed; also successor is modification of the current state, but we have to be able to undo each modification.
More when we talk about Constraint Satisfaction Problems (CSP).

Iterative deepening search $l = 0$

Limit = 0




Iterative deepening search $l = 1$

Limit = 1



Iterative deepening search $l=2$

Limit = 2



Why would one do that?

Combine **good memory requirements** of depth-first with the **completeness** of breadth-first when branching factor is finite and is **optimal** when the path cost is a non-decreasing function of the depth of the node.

Idea was a breakthrough in game playing. All game tree search uses iterative deepening nowadays. What's the added advantage in games?

“Anytime” nature.

Iterative deepening search

```
function ITERATIVE-DEEPENING-SEARCH(problem) returns a solution, or fail-  
ure  
  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
```

Iterative deepening search

Number of nodes generated in an **iterative deepening search** to depth d with branching factor b :

Looks quite wasteful, is it?

“One search depth d ”

$$N_{IDS} = d b^1 + (d-1)b^2 + \dots + 3b^{d-2} + 2b^{d-1} + 1b^d$$

Nodes generated in a **breadth-first search** with branching factor b :

$$N_{BFS} = b^1 + b^2 + \dots + b^{d-2} + b^{d-1} + b^d$$

For $b = 10$, $d = 5$,

– $N_{BFS} = 10 + 100 + 1,000 + 10,000 + 100,000 = 111,110$

–

– $N_{IDS} = 50 + 400 + 3,000 + 20,000 + 100,000 = 123,456$



Iterative deepening is the preferred uninformed search method when there is a large search space and the depth of the solution is not known.

Properties of iterative deepening search

Complete? Yes

(b finite)


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

Space? $O(bd)$

Optimal? Yes, if step costs identical

Bidirectional Search

- Simultaneously:
 - Search forward from start
 - Search backward from the goal
- Stop when the two searches meet.




- If branching factor = b in each direction, with solution at depth d
 - only $O(2 b^{d/2}) = O(2 b^{d/2})$
- Checking a node for membership in the other search tree can be done in constant time (hash table)
- Key limitations:
 - Space $O(b^{d/2})$
 - Also, how to search backwards can be an issue (e.g., in Chess)? What's tricky?
 - Problem: lots of states satisfy the goal; don't know which one is relevant.

Aside: The predecessor of a node should be easily computable (i.e., actions are easily reversible).

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

Repeated states



Problems in which actions are reversible (e.g., routing problems or sliding-blocks puzzle). Also, in eg Chess; uses hash tables to check for repeated states. Huge tables 100M+ size but very useful.

See Tree-Search vs. Graph-Search in Fig. 3.7 R&N. But need to be careful to maintain (path) optimality and completeness.

Summary: General, uninformed search

Original search ideas in AI were inspired by studies of human problem solving in, eg, puzzles, math, and games, but a great many AI tasks now require some form of search (e.g. find optimal agent strategy; active learning; constraint reasoning; NP-complete problems require search).

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.

Avoid repeating states / cycles.

Verify these results at home!!!!

Criterion	Breadth-First	Uniform-Cost	Depth-First	Depth-Limited	Iterative Deepening	Bidirectional (if applicable)
Complete?	Yes ^a	Yes ^{a,b}	No	No	Yes ^a	Yes ^{a,d}
Time	$O(b^{d+1})$	$O(b^{1+\lceil C^*/\epsilon \rceil})$	$O(b^m)$	$O(b^l)$	$O(b^d)$	$O(b^{d/2})$
Space	$O(b^{d+1})$	$O(b^{1+\lceil C^*/\epsilon \rceil})$	$O(bm)$	$O(bl)$	$O(bd)$	$O(b^{d/2})$
Optimal?	Yes ^c	Yes	No	No	Yes ^c	Yes ^{c,d}

Figure 3.17 Evaluation of search strategies. b is the branching factor; d is the depth of the shallowest solution; m is the maximum depth of the search tree; l is the depth limit. Superscript caveats are as follows: ^a complete if b is finite; ^b complete if step costs $\geq \epsilon$ for positive ϵ ; ^c optimal if step costs are all identical; ^d if both directions use breadth-first search.