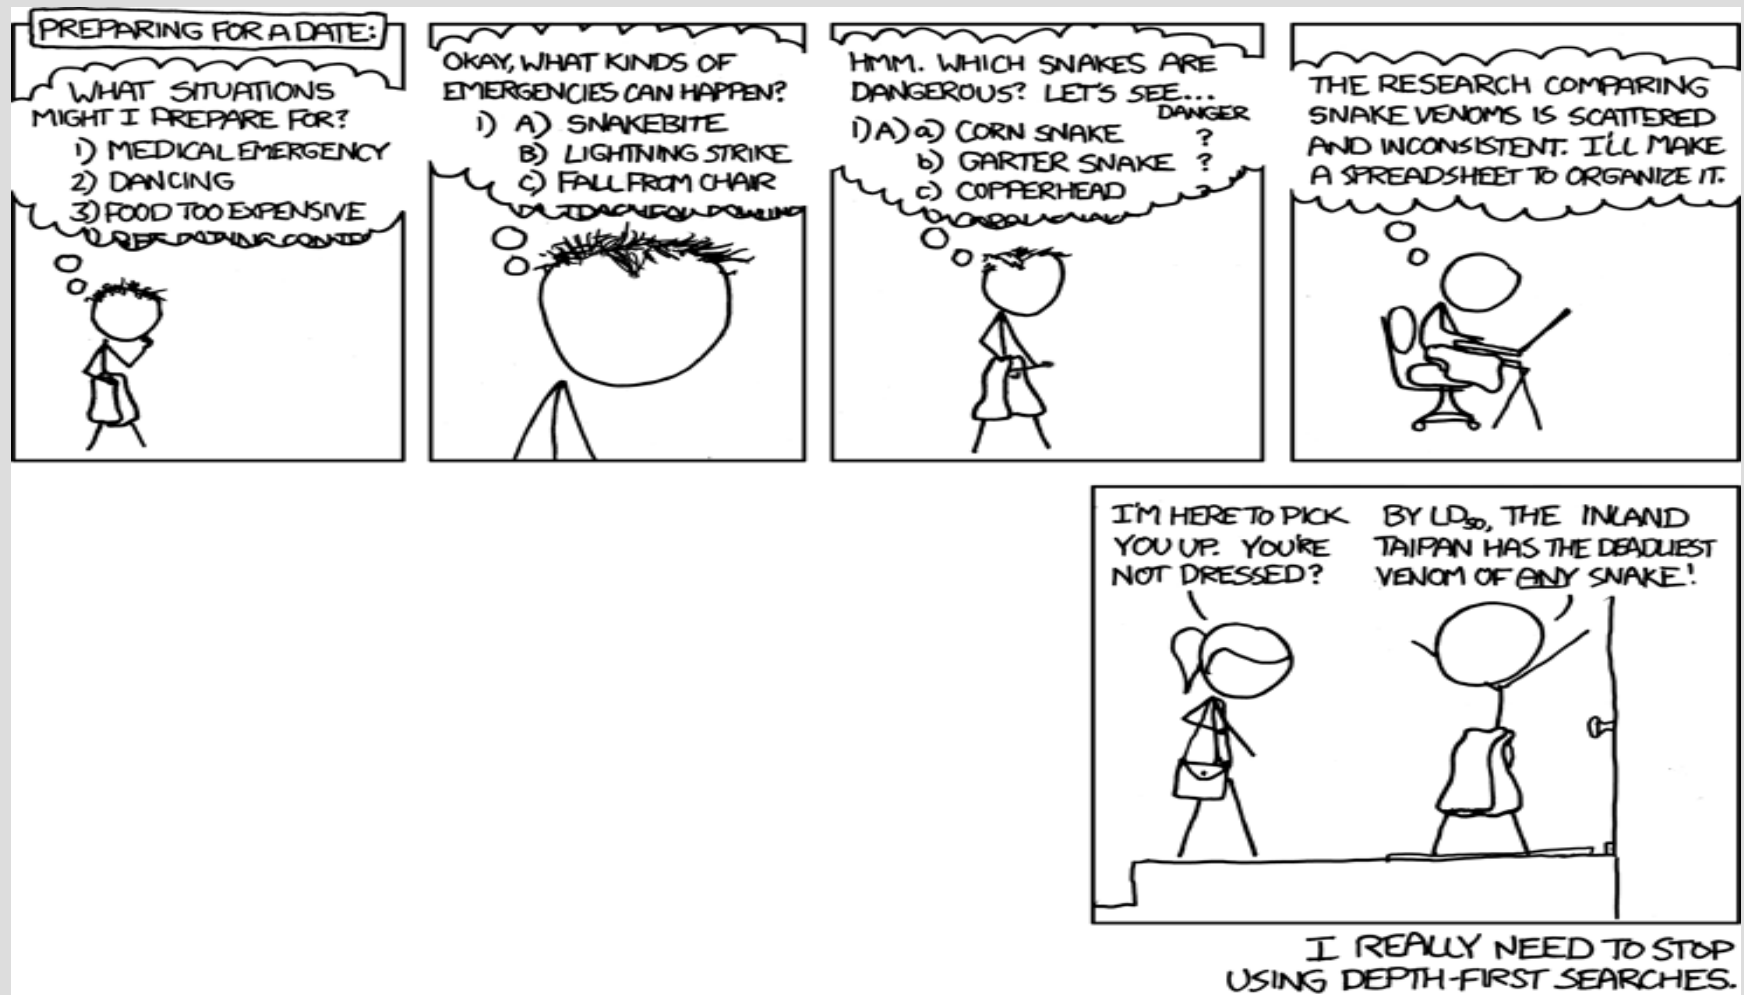


Uninformed Search (Ch. 3-3.4)



Search

Goal based agents need to search to find a path from their start to the goal (a path is a sequence of actions, not states)

For now we consider problem solving agents who search on atomically structured spaces

Today we will focus on uninformed searches, which only know cost between states but no other extra information

Search

In the vacuum example, the states and actions I gave upfront (so only one option)

In more complex environments, we have a choice of how to abstract the problem into simple (yet expressive) states and actions

The solution to the abstracted problem should be able to serve as the basis of a more detailed problem (i.e. fit the detailed solution inside)

Search

Example: Google maps gives direction by telling you a sequence of roads and does not dictate speed, stop signs/lights, road lane

← from Smith Hall, 207 Pleasant St SE, Minneapolis, MN ...
to Kenneth H. Keller Hall, 200 Union St SE, Minneapolis...

4 min (0.2 mile)
via Scholars Walk

⚠ Use caution—walking directions may not always reflect real-world conditions

Smith Hall
207 Pleasant St SE, Minneapolis, MN 55455

- ↑ Head north toward Scholars Walk
125 ft
- ➡ Turn right onto Scholars Walk
0.1 mile
- ↗ Slight right
75 ft
- ↗ Slight right
Destination will be on the left
282 ft

Kenneth H. Keller Hall
200 Union St SE, Minneapolis, MN 55455

These directions are for planning purposes only. You may find that construction projects, traffic, weather, or other events may cause

Map data ©2018 Google, United States, Terms, Send feedback

Search

In deterministic environments the search solution is a single sequence (list of actions)

Stochastic environments need multiple sequences to account for all possible outcomes of actions

It can be costly to keep track of all of these and might be better to keep the most likely and search again when off the main sequences

Search

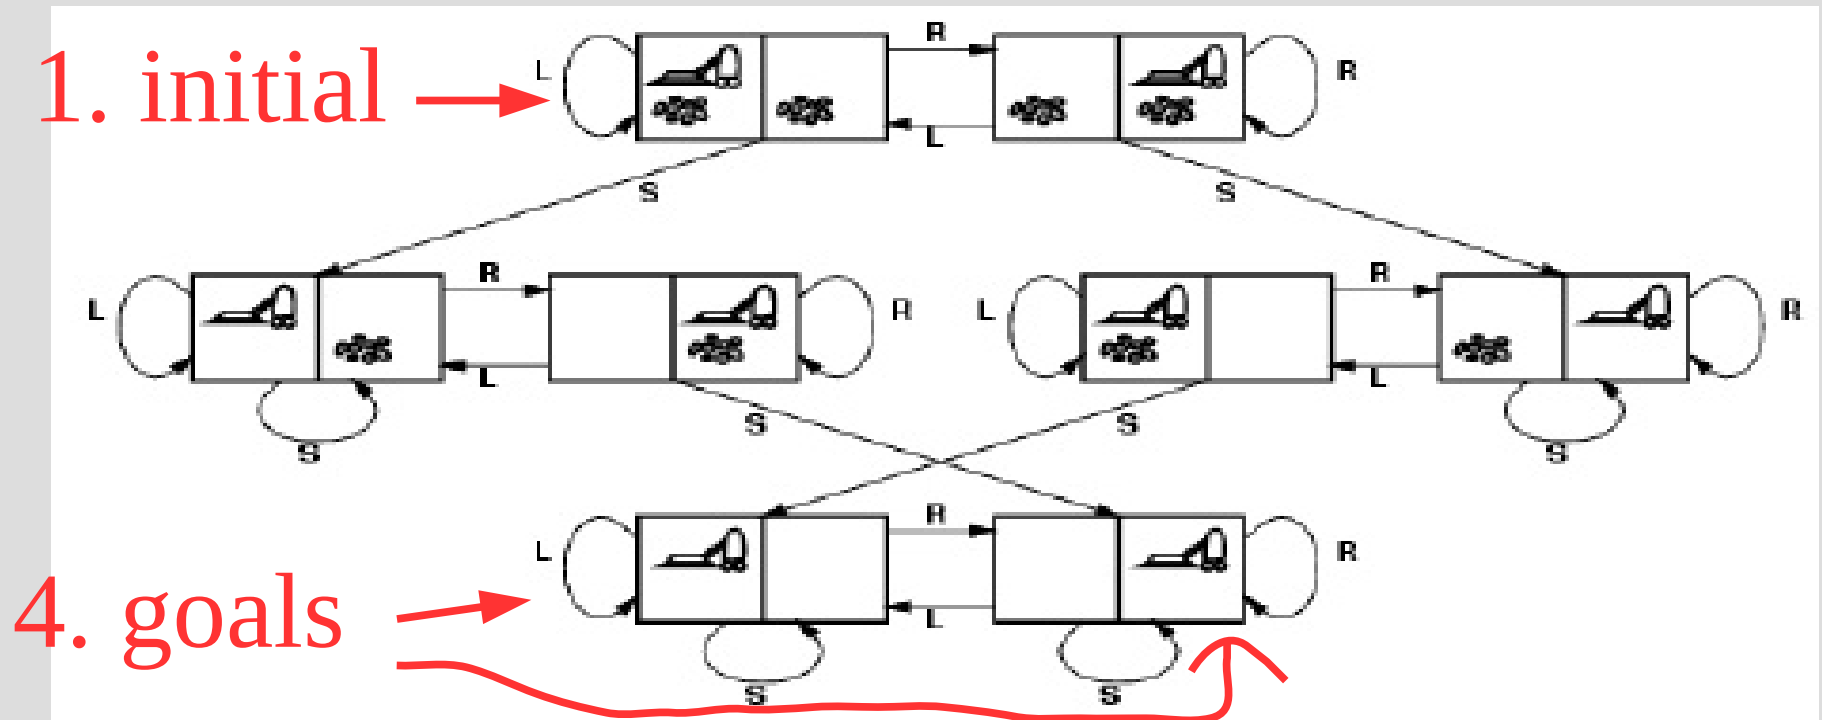
There are 5 parts to search:

1. Initial state
2. Actions possible at each state
3. Transition model (result of each action)
4. Goal test (are we there yet?)
5. Path costs/weights (not stored in states)
(related to performance measure)

In search we normally fully see the problem and the initial state and compute all actions

Small examples

Here is our vacuum world again:




2. For all states, we have actions: L, R or S

3. Transition model = black arrows

5. Path cost = ??? (from performance measure)

Small examples

8-Puzzle

1. (semi) Random
2. All states: U,D,L,R
4. As shown **here** 
5. Path cost = 1 (move count)
3. Transition model (example):

1	2	3
4	5	6
7	8	

$$\text{Result}\left(\begin{array}{|c|c|c|} \hline 1 & 2 & 3 \\ \hline 4 & 5 & \\ \hline 7 & 8 & 6 \\ \hline \end{array}, D \right) = \begin{array}{|c|c|c|} \hline 1 & 2 & 3 \\ \hline 4 & 5 & 6 \\ \hline 7 & 8 & \\ \hline \end{array}$$

(see: <https://www.youtube.com/watch?v=DfVjTkzk2Ig>)

Small examples

8-Puzzle is NP complete so to find the best solution, we must brute force

$$3 \times 3 \text{ board} = \begin{array}{|c|c|c|} \hline 1 & 2 & 3 \\ \hline 4 & 5 & 6 \\ \hline 7 & 8 & \\ \hline \end{array} = 181,440 \text{ states}$$

4x4 board = 1.3 trillion states

Solution time: milliseconds

5x5 board = 10^{25} states

Solution time: hours


Small examples

8-Queens: how to fit 8 queens on a 8x8 board so no 2 queens can capture each other

Two ways to model this:

Incremental = each action is to add a queen to the board
(1.8×10^{14} states)

Complete state formulation = all 8 queens start on board, action = move a queen
(2057 states)



Real world examples


Directions/traveling (land or air)



Model choices: only have interstates?
Add smaller roads, with increased cost?
(pointless if they are never taken)

Real world examples

Traveling salesperson problem (TSP): Visit each location exactly once and return to start



Goal: Minimize distance traveled

Search algorithm

To search, we will build a tree with the root as the initial state


```
function tree-search(root-node)
  fringe ← successors(root-node)
  while ( notempty(fringe) )
    {node ← remove-first(fringe)
     state ← state(node)
     if goal-test(state) return solution(node)
     fringe ← insert-all(successors(node),fringe) }
  return failure
end tree-search
```

(Use same procedure for multiple algorithms)

Search algorithm

What are states/actions for this problem?


Can you help Curious George find the man with the yellow hat?



Search algorithm

Multiple options, but this is a good choice

Can you help Curious George find the man with the yellow hat?

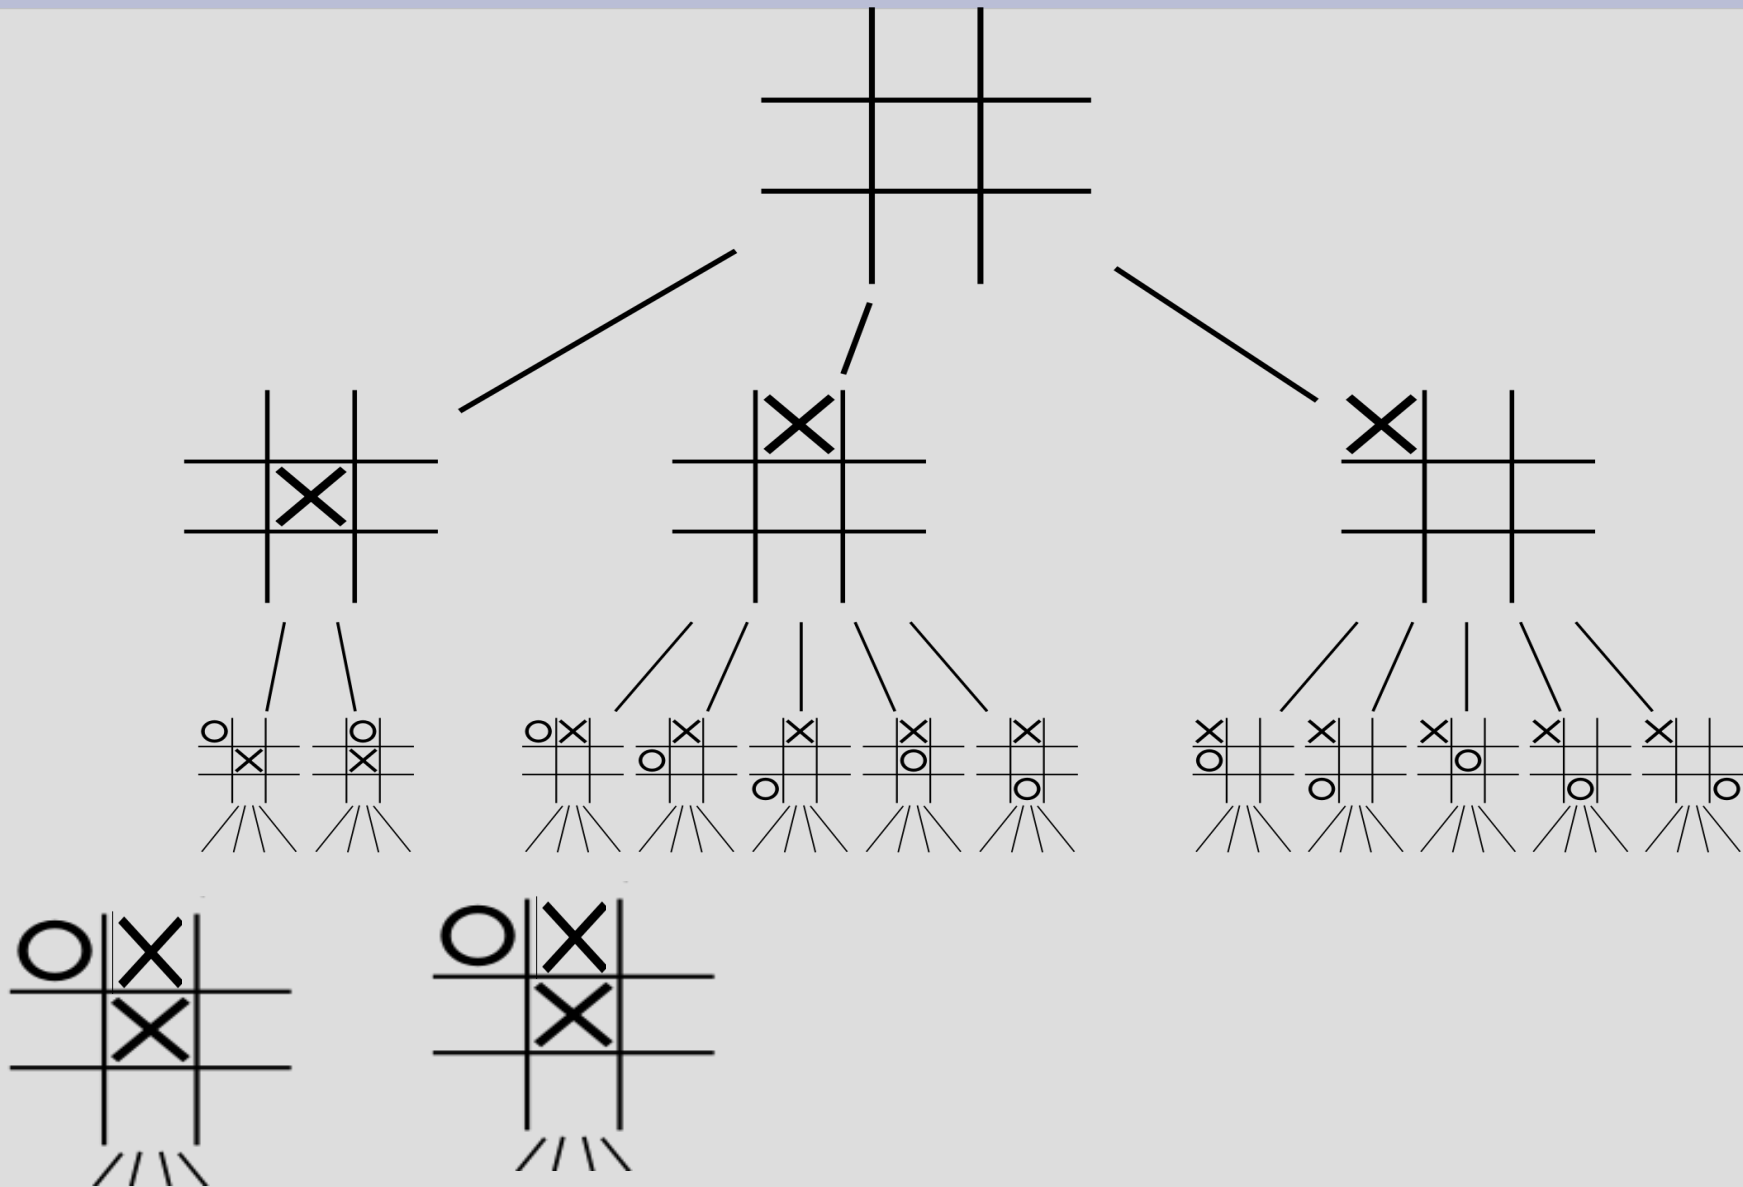


Search algorithm

What are the problems with this?

```
function tree-search(root-node)
  fringe ← successors(root-node)
  while ( notempty(fringe) )
    {node ← remove-first(fringe)
     state ← state(node)
     if goal-test(state) return solution(node)
     fringe ← insert-all(successors(node),fringe) }
  return failure
end tree-search
```


Search algorithm



Search algorithm

We can remove visiting states multiple times by doing this:

```
function tree-search(root-node)
  fringe ← successors(root-node)
  explored ← empty
  while ( notempty(fringe) )
    {node ← remove-first(fringe)
     state ← state(node)
     if goal-test(state) return solution(node)
     explored ← insert(node, explored)
     fringe ← insert-all(successors(node), fringe, if node not in explored)
    }
  return failure
end tree-search
```

But this is still not necessarily all that great...

Search algorithm

When we find a goal state, we can back track via the parent to get the sequence

To keep track of the unexplored nodes, we will use a queue (of various types)

The explored set is probably best as a hash table for quick lookup (have to ensure similar states reached via alternative paths are the same in the has, can be done by sorting)

Search algorithm

The search algorithms metrics/criteria:

1. Completeness (does it terminate with a valid solution)
2. Optimality (is the answer the best solution)
3. Time (in big-O notation)
4. Space (big-O)

b = maximum branching factor


d = minimum depth of a goal

m = maximum length of any path


Breadth first search

Breadth first search checks all states which are reached with the fewest actions first

(i.e. will check all states that can be reached by a single action from the start, next all states that can be reached by two actions, then three...)



Breadth first search



(see: <https://www.youtube.com/watch?v=5UfMU9TsoEM>)

(see: <https://www.youtube.com/watch?v=nI0dT288VLs>)

Breadth first search

BFS can be implemented by using a simple FIFO (first in, first out) queue to track the fringe/frontier/unexplored nodes

Metrics for BFS:

Complete (i.e. guaranteed to find solution if exists)

Non-optimal (unless uniform path cost)

Time complexity = $O(b^d)$

Space complexity = $O(b^d)$

Breadth first search



Exponential problems are not very fun, as seen in this picture:

Depth	Nodes	Time	Memory
2	110	.11 milliseconds	107 kilobytes
4	11,110	11 milliseconds	10.6 megabytes
6	10^6	1.1 seconds	1 gigabyte
8	10^8	2 minutes	103 gigabytes
10	10^{10}	3 hours	10 terabytes
12	10^{12}	13 days	1 petabyte
14	10^{14}	3.5 years	99 petabytes
16	10^{16}	350 years	10 exabytes


Figure 3.13 Time and memory requirements for breadth-first search. The numbers shown assume branching factor $b = 10$; 1 million nodes/second; 1000 bytes/node.

Uniform-cost search

Uniform-cost search also does a queue, but uses a priority queue based on the cost (the lowest cost node is chosen to be explored)



(a)



(b)

Uniform-cost search

The only modification is when exploring a node we cannot disregard it if it has already been explored by another node

We might have found a shorter path and thus need to update the cost on that node

We also do not terminate when we find a goal, but instead when the goal has the lowest cost in the queue.

Uniform-cost search

UCS is..



1. Complete (if costs strictly greater than 0)
2. Optimal

However....

3&4. Time complexity = space complexity
= $O(b^{1+C^*/\min(\text{path cost})})$, where C^* cost of
optimal solution (much worse than BFS)

Depth first search

DFS is same as BFS except with a FILO (or LIFO) instead of a FIFO queue



Depth first search

Metrics:

1. Might not terminate (not correct) (e.g. in vacuum world, if first expand is action L)
2. Non-optimal (just... no)
3. Time complexity = $O(b^d)$
4. Space complexity = $O(b * d)$

Only way this is better than BFS is the space complexity...



Depth limited search

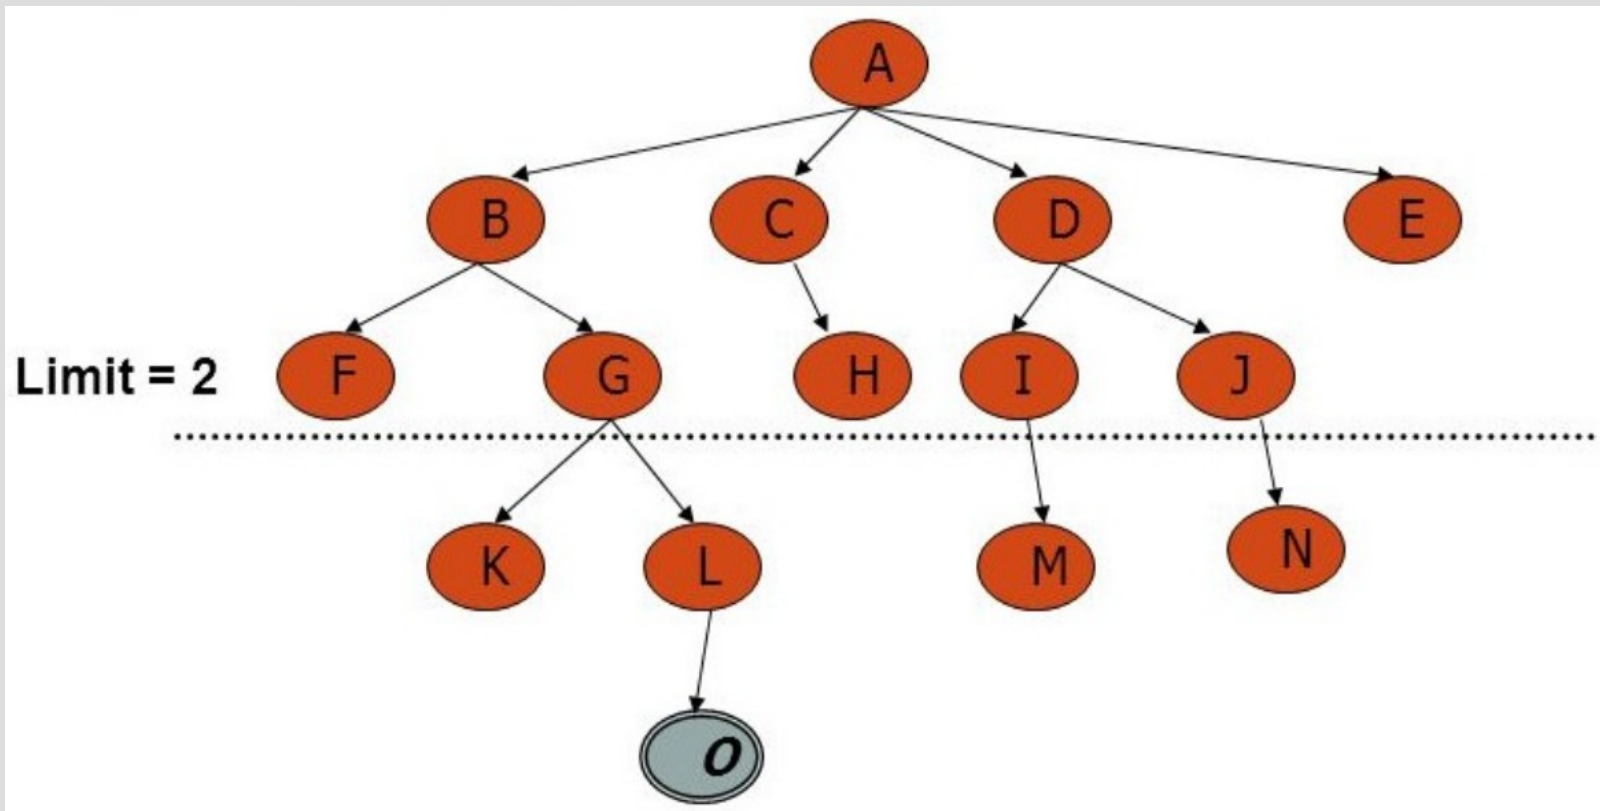
DFS by itself is not great, but it has two (very) useful modifications

Depth limited search runs normal DFS, but if it is at a specified depth limit, you cannot have children (i.e. take another action)

Typically with a little more knowledge, you can create a reasonable limit and makes the algorithm correct

Depth limited search

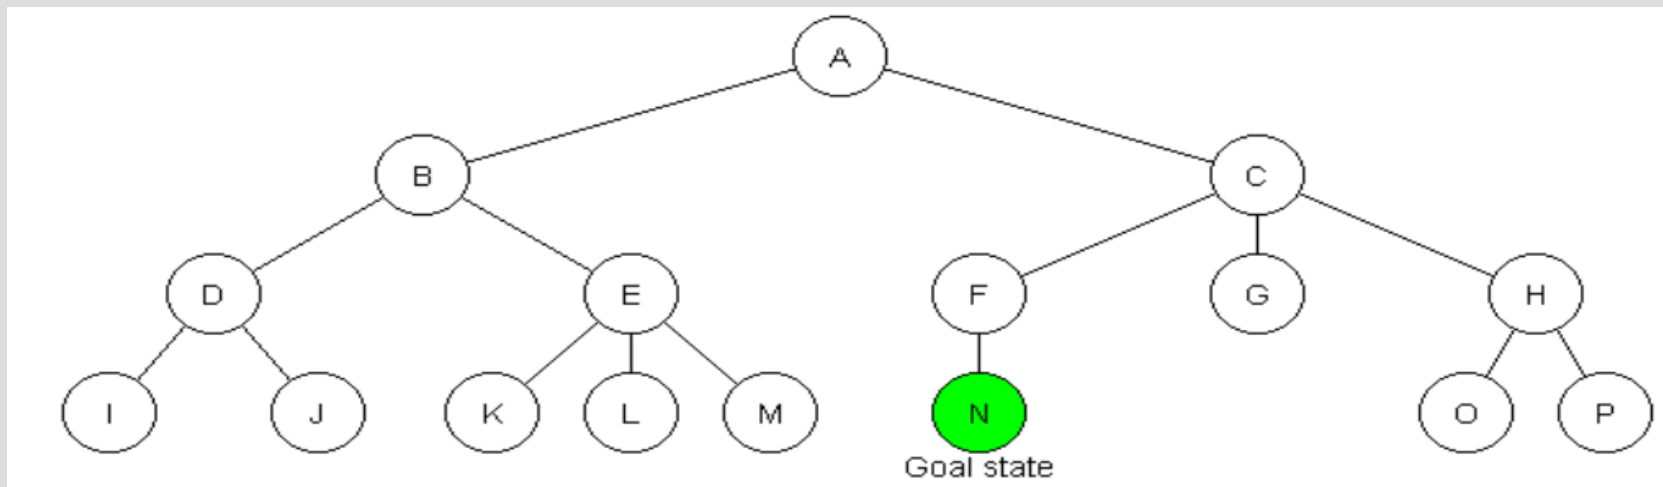
However, if you pick the depth limit before d , you will not find a solution (not correct, but will terminate)




Iterative deepening DFS

Probably the most useful uninformed search is iterative deepening DFS

This search performs depth limited search with maximum depth 1, then maximum depth 2, then 3... until it finds a solution



Iterative deepening DFS



Iterative deepening DFS

The first few states do get re-checked multiple times in IDS, however it is not too many

When you find the solution at depth d , depth 1 is expanded d times (at most b of them)

The second depth are expanded $d-1$ times (at most b^2 of them)

Thus $d \cdot b + (d - 1) \cdot b^2 + \dots + 1 \cdot b^d = O(b^d)$

Iterative deepening DFS

Metrics:

1. Complete
2. Non-optimal (unless uniform cost)
3. $O(b^d)$
4. $O(b * d)$


Thus IDS is better in every way than BFS (asymptotically)

Best uninformed we will talk about

Bidirectional search

Bidirectional search starts from both the goal and start (using BFS) until the trees meet

This is better as $2 * (b^{d/2}) < b^d$
(the space is much worse than IDS, so only applicable to small problems)



Summary of algorithms

Fig. 3.21, p. 91

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

There are a number of footnotes, caveats, and assumptions.
See Fig. 3.21, p. 91.

[a] complete if b is finite

[b] complete if step costs $\geq \epsilon > 0$

[c] optimal if step costs are all identical
(also if path cost non-decreasing function of depth only)

[d] if both directions use breadth-first search

(also if both directions use uniform-cost search with step costs $\geq \epsilon > 0$)

Generally the preferred
uninformed search strategy