



CS 540 Introduction to Artificial Intelligence

Informed Search

Yingyu Liang
University of Wisconsin-Madison
Nov 16, 2021

Based on slides by Fred Sala

Outline

- Uninformed continued
- A* Search
 - Heuristic properties, stopping rules, analysis

General State-Space Search Algorithm

```
function general-search(problem, QUEUEING-FUNCTION)
;; problem describes the start state, operators, goal test, and
;; operator costs
;; queueing-function is a comparator function that ranks two states
;; general-search returns either a goal node or "failure"

nodes = MAKE-QUEUE(MAKE-NODE(problem.INITIAL-STATE))
loop
  if EMPTY(nodes) then return "failure"
  node = REMOVE-FRONT(nodes)
  if problem.GOAL-TEST(node.STATE) succeeds then return node
  nodes = QUEUEING-FUNCTION(nodes, EXPAND(node,
                                     problem.OPERATORS))
;; succ(s)=EXPAND(s, OPERATORS)
;; Note: The goal test is NOT done when nodes are generated
;; Note: This algorithm does not detect loops
end
```

Recall the bad space complexity of BFS

Four measures of search algorithms:




- **Completeness** (not finding all goals): find a goal.
- **Optimality**: yes if edges cost 1 (more generally positive non-decreasing with depth), **no otherwise**.
- **Time** complexity (goal is the last node at radius d):
 - Have to generate all nodes at radius d .
 - $b + b^2 + \dots + b^d \sim O(b^d)$
- **Space** complexity (bad, see the Figure)
 - Back points for all generated nodes $O(b^d)$
 - The queue (smaller, but still $O(b^d)$)

Solution:
Uniform-cost
search


Solution:
Depth-first
search

Depth-first search

Expand the deepest node first

1. Select a direction, go deep to the end 
2. Slightly change the end 
3. Slightly change the end some more... 


fan



Depth-first search (DFS)

Use a **stack** (First-in Last-out)

1. push(Initial states)
2. While (stack not empty)
3. s = pop()
4. if (s==goal) success!
5. T = succs(s)
6. push(T)
7. endwhile




stack (**fringe**)


[] ⇔

What's in the fringe for DFS?

- m = maximum depth of graph from start
- $m(b-1) \sim O(mb)$
(Space complexity)




c.f. BFS $O(b^d)$




- “backtracking search” even less space
 - generate siblings (if applicable)

What's wrong with DFS?

- Infinite tree: may not find goal (incomplete)
- May not be optimal
- Finite tree: may visit almost all nodes, time complexity $O(b^m)$



c.f. BFS $O(b^d)$



Performance of search algorithms on trees

b: branching factor (assume finite)

d: goal depth


m: graph depth

	Complete	optimal	time	space
Breadth-first search	Y	Y, if ¹	$O(b^d)$	$O(b^d)$
Uniform-cost search ²	Y	Y	$O(b^{C^*/\epsilon})$	$O(b^{C^*/\epsilon})$
Depth-first search	N	N	$O(b^m)$	$O(bm)$

1. edge cost constant, or positive non-decreasing in depth
2. edge costs $\geq \epsilon > 0$. C^* is the best goal path cost.

How about this?

1. DFS, but stop if path length > 1 .
 2. If goal not found, repeat DFS, stop if path length > 2 .
 3. And so on...
- fan within ripple



Iterative deepening

- Search proceeds like BFS, but fringe is like DFS
 - Complete, optimal like BFS
 - Small space complexity like DFS
 - Time complexity like BFS
- Preferred uninformed search method

Nodes expanded by:

- Breadth-First Search: S A B C D E G

Solution found: S A G

- Uniform-Cost Search: S A D B C E G


Solution found: S B G (This is the only uninformed search that worries about costs.)

- Depth-First Search: S A D E G

Solution found: S A G

- Iterative-Deepening Search: S A B C S A D E G

Solution found: S A G



Performance of search algorithms on trees

b: branching factor (assume finite)

d: goal depth


m: graph depth

	Complete	optimal	time	space
Breadth-first search	Y	Y, if ¹	$O(b^d)$	$O(b^d)$
Uniform-cost search ²	Y	Y	$O(b^{C^*/\epsilon})$	$O(b^{C^*/\epsilon})$
Depth-first search	N	N	$O(b^m)$	$O(bm)$
Iterative deepening	Y	Y, if ¹	$O(b^d)$	$O(bd)$

1. edge cost constant, or positive non-decreasing in depth
2. edge costs $\geq \epsilon > 0$. C^* is the best goal path cost.

If state space graph is not a tree

- The problem: repeated states




- Ignore the danger of repeated states: wasteful (BFS) or impossible (DFS). Can you see why?
- How to prevent it?

If state space graph is not a tree

- We have to remember already-expanded states (**CLOSED**).
- When we take out a state from the fringe (OPEN), check whether it is in CLOSED (already expanded).
 - If yes, throw it away.
 - If no, expand it (add successors to OPEN), and move it to CLOSED.

What you should know

- Problem solving as search: state, successors, goal test
 - Uninformed search
 - Breadth-first search
 - Uniform-cost search
 - Depth-first search
 - Iterative deepening★
- 
- Can you unify them using the same algorithm, with different priority functions?
 - Performance measures
 - Completeness, optimality, time complexity, space complexity

Uninformed vs Informed Search


Uninformed search (all of what we saw). Know:

- Path cost $g(s)$ from start to node s
- Successors.



Informed search. Know:


- All uninformed search properties, plus
- Heuristic $h(s)$ from s to goal



Informed Search

Informed search. Know:

- All uninformed search properties, plus
- Heuristic $h(s)$ from s to goal




- Use information to **speed up search**.

Using the Heuristic

Back to uniform-cost search

- We had the priority queue
- Expand the node with the smallest $g(s)$
 - $g(s)$ “first-half-cost”




- Now let's use the heuristic (“second-half-cost”)
 - Several possible approaches: let's see what works

Attempt 1: Best-First Greedy

One approach: just use $h(s)$ alone

- Specifically, expand node with smallest $h(s)$
- This isn't a good idea. Why?




- Not optimal! **Get** $A \rightarrow C \rightarrow G$. **Want:** $A \rightarrow B \rightarrow C \rightarrow G$

Attempt 2: A Search

Next approach: use both $g(s) + h(s)$

- Specifically, expand node with smallest $g(s) + h(s)$
- Again, use a priority queue
- Called “A” search




- **Still not optimal!** (Does work for former example).

Attempt 3: A* Search

Same idea, use $g(s) + h(s)$, with one **requirement**

- Demand that $0 \leq h(s) \leq h^*(s)$, the actual cost
- If heuristic has this property, “admissible”
 - Optimistic! Never over-estimates
- Still need $h(s) \geq 0$
 - Negative heuristics can lead to strange behavior
- This is **A*** search



Admissible Heuristic Functions

Have to be careful to ensure admissibility (**optimism!**)

- Example: **8-puzzle**

Example State

1		5
2	6	3
7	4	8

Goal State

1	2	3
4	5	6
7	8	

- One useful approach: **relax constraints**
 - $h(s)$ = number of tiles in wrong position
 - allows tiles to fly to destination in a single step

Heuristic Function Tradeoffs

Dominance: h_2 dominates h_1 if for all states s ,


$$h_1(s) \leq h_2(s) \leq h^*(s)$$

- **Idea:** we want to be as close to h^* as possible
 - But not over!
- **Tradeoff:** being very close might require a very complex heuristic, expensive computation
 - Might be better off with cheaper heuristic & expand more nodes.

A* Termination

When should A* **stop**?

- One idea: as soon as we reach goal state?




- *h* admissible, but note that we get $A \rightarrow B \rightarrow G$ (**cost 1000**)!

A* Termination

When should A* stop?


- **Rule:** terminate **when a goal is popped** from queue.



- Note: taking $h = 0$ reduces to uniform cost search rule.

A* Revisiting Expanded States

Possible to revisit an expanded state, get a shorter path:



- Put D back into priority queue, smaller $g+h$

A* Full Algorithm

1. Put the start node S on the priority queue, called $OPEN$
2. If $OPEN$ is empty, exit with failure
3. Remove from $OPEN$ and place on $CLOSED$ a node n for which $f(n)$ is minimum (note that $f(n)=g(n)+h(n)$)
4. If n is a goal node, exit (trace back pointers from n to S)
5. Expand n , generating all successors and attach to pointers back to n . For each successor n' of n
 1. If n' is not already on $OPEN$ or $CLOSED$ estimate $h(n')$, $g(n')=g(n)+c(n,n')$, $f(n')=g(n')+h(n')$, and place it on $OPEN$.
 2. If n' is already on $OPEN$ or $CLOSED$, then check if $g(n')$ is lower for the new version of n' . If so, then:
 1. Redirect pointers backward from n' along path yielding lower $g(n')$.
 2. Put n' on $OPEN$.
 3. If $g(n')$ is not lower for the new version, do nothing.
6. Goto 2.

A* Analysis

Some properties:

- Terminates!
- A* can use **lots of memory**: $O(\# \text{ states})$.
- Will run out on large problems.

The screenshot shows a terminal window with a progress bar for an A* search. The progress bar consists of 48 columns, each representing a step in the search. Each column contains a percentage of completion (e.g., 93.5%, 89.6%, 76.4%) and a number (e.g., 25, 26, 27). The progress bar is followed by system resource usage statistics: Mem: 177G/370G, Swap: 0K/0K, Tasks: 332, 49 thr, 749 kthr, 89 running, Load average: 157.55 143.11 124.84, Uptime: 13:52:44.

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
3231	ubuntu	20	0	2384M	1925M	5820	R	102.	0.5	10h48:36	/usr/lib/R/bin/exec/R
3232	ubuntu	20	0	2384M	1925M	5820	R	102.	0.5	10h38:29	/usr/lib/R/bin/exec/R
3176	ubuntu	20	0	2792M	2252M	5788	R	102.	0.6	11h12:51	/usr/lib/R/bin/exec/R
3179	ubuntu	20	0	2384M	1925M	5800	R	102.	0.5	11h37:07	/usr/lib/R/bin/exec/R
3154	ubuntu	20	0	1978M	1361M	5800	R	102.	0.4	11h19:01	/usr/lib/R/bin/exec/R
3146	ubuntu	20	0	2588M	1943M	5788	R	102.	0.5	11h18:23	/usr/lib/R/bin/exec/R
3208	ubuntu	20	0	2402M	1644M	5788	R	102.	0.4	11h03:11	/usr/lib/R/bin/exec/R
3148	ubuntu	20	0	2922M	2367M	5788	R	102.	0.6	11h38:07	/usr/lib/R/bin/exec/R
3230	ubuntu	20	0	1988M	1522M	5788	R	102.	0.4	10h48:57	/usr/lib/R/bin/exec/R
3150	ubuntu	20	0	2588M	1985M	5788	R	102.	0.5	11h42:22	/usr/lib/R/bin/exec/R
3207	ubuntu	20	0	2615M	2157M	5788	R	102.	0.6	11h13:52	/usr/lib/R/bin/exec/R
3200	ubuntu	20	0	2996M	2337M	5788	R	101.	0.6	11h46:18	/usr/lib/R/bin/exec/R
3157	ubuntu	20	0	2792M	2263M	5800	R	101.	0.6	11h52:37	/usr/lib/R/bin/exec/R
3152	ubuntu	20	0	2591M	2068M	5788	R	101.	0.5	11h37:38	/usr/lib/R/bin/exec/R

Summary

- Informed search: introduce heuristics
 - Not all approaches work: best-first greedy is bad
- A* algorithm
 - Properties of A*, idea of admissible heuristics



Acknowledgements: Adapted from materials by Jerry Zhu, Anthony Gitter, and Fred Sala (University of Wisconsin-Madison).