

Informed Search

Day 2/3 of Search

Chap. 4, Russel & Norvig


Material in part from <http://www.cs.cmu.edu/~awm/tutorials>

Uninformed Search Complexity

- N = Total number of states
- B = Average number of successors (branching factor)
- L = Length for start to goal with smallest number of steps
- Q = Average size of the priority queue
- L_{max} = Length of longest path from $START$ to any state


Algorithm	Complete	Optimal	Time	Space
BFS Breadth First Search	Y	Y, If all trans. have same cost	$O(\text{Min}(N, B^L))$	$O(\text{Min}(N, B^L))$
BIBFS Bi- Direction. BFS	Y	Y, If all trans. have same cost	$O(\text{Min}(N, 2B^{L/2}))$	$O(\text{Min}(N, 2B^{L/2}))$
UCS Uniform Cost Search	Y, If cost > 0	Y, If cost > 0	$O(\log(Q) * B^{C/q})$	$O(\text{Min}(N, B^{C/q}))$
PCDFS Path Check DFS	Y	N	$O(B^{L_{max}})$	$O(B^{L_{max}})$
MEMD FS Memorizing DFS	Y	N	$O(\text{Min}(N, B^{L_{max}}))$	$O(\text{Min}(N, B^{L_{max}}))$
IDS Iterative Deepening	Y	Y, If all trans. have same cost	$O(B^L)$	$O(BL)$

Search Revisited



1. Store a value $f(s)$ at each state s
2. Choose the state with lowest f to expand next
3. Insert its successors

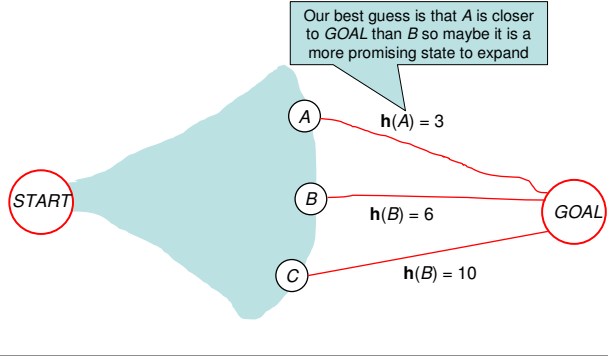
If $f(\cdot)$ is chosen carefully, we will eventually find the lowest-cost sequence



Example:

- UCS (Uniform Cost Search): $f(A) = g(A)$ = total cost of current shortest path from $START$ to A
- Store states awaiting expansion in a priority queue for efficient retrieval of minimum f
- Optimal \rightarrow Guaranteed to find lowest cost sequence, *but*.....


- Problem: No guidance as to how “far” any given state is from the goal
- Solution: Design a function $h(\cdot)$ that gives us an estimate of the distance between a state and the goal



Heuristic Functions


- $h(\cdot)$ is a heuristic function for the search problem
- $h(s)$ = estimate of the cost of the shortest path from s to *GOAL*
- $h(\cdot)$ cannot be computed solely from the states and transitions in the current problem → If we could, we would already know the optimal path!
- $h(\cdot)$ is based on external knowledge about the problem → *informed* search
- Questions:
 1. Typical examples of h ?
 2. How to use h ?
 3. What are desirable/necessary properties of h ?

Heuristic Functions Example



- $h(s)$ = Linear-geometric distance to *GOAL*

Heuristic Functions Example



- How could we define $h(s)$?



First Attempt: Greedy Best First Search

- Simplest use of heuristic function: Always select the node with smallest $h(\cdot)$ for expansion (i.e., $f(s) = h(s)$)

Initialize PQ

Insert $START$ with value $h(START)$ in PQ

While (PQ not empty and no goal state is in PQ)


Pop the state s with the minimum value of h from PQ

For all s' in $\text{succs}(s)$

If s' is not already in PQ and has not already been visited


Insert s' in PQ with value $h(s')$

Problem




- What solution do we find in this case?
- Greedy search clearly not optimal, even though the heuristic function is non-stupid

Trying to Fix the Problem



- $g(s)$ is the cost from $START$ to s only
- $h(s)$ estimates the cost from s to $GOAL$
- Key insight: $g(s) + h(s)$ estimates the **total** cost of the cheapest path from $START$ to $GOAL$ going through s
- **A* algorithm**

Can A* Fix the Problem?



$\{(START,4)\}$

$\{(A,5)\}$

$$(f(A) = h(A) + g(A) = 3 + g(START) + \text{cost}(START, A) = 3 + 0 + 2)$$

$\{(B,5) (C,7)\}$

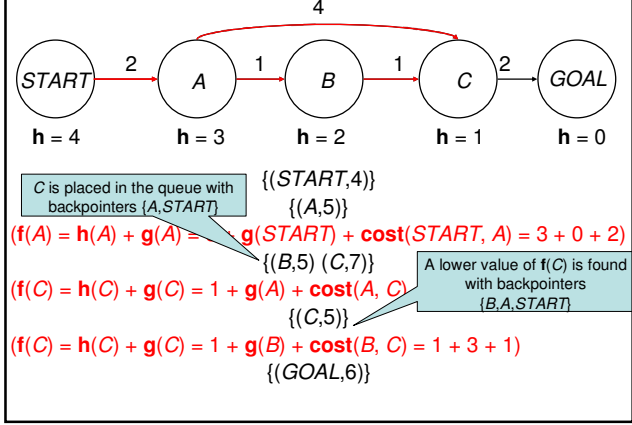
$$(f(C) = h(C) + g(C) = 1 + g(A) + \text{cost}(A, C) = 1 + 2 + 4)$$

$\{(C,5)\}$


$$(f(C) = h(C) + g(C) = 1 + g(B) + \text{cost}(B, C) = 1 + 3 + 1)$$

$\{(GOAL,6)\}$

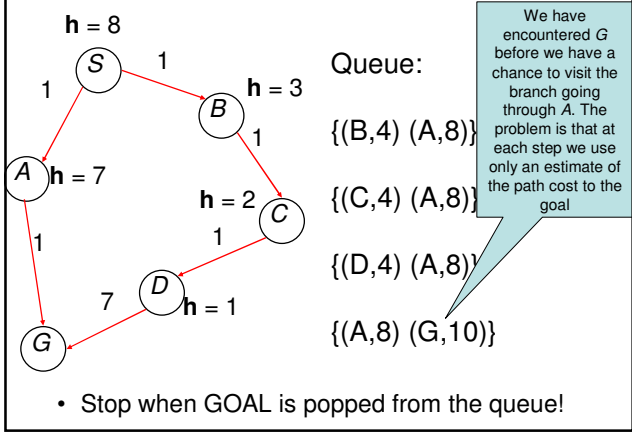
Can A* Fix the Problem?




A* Termination Condition

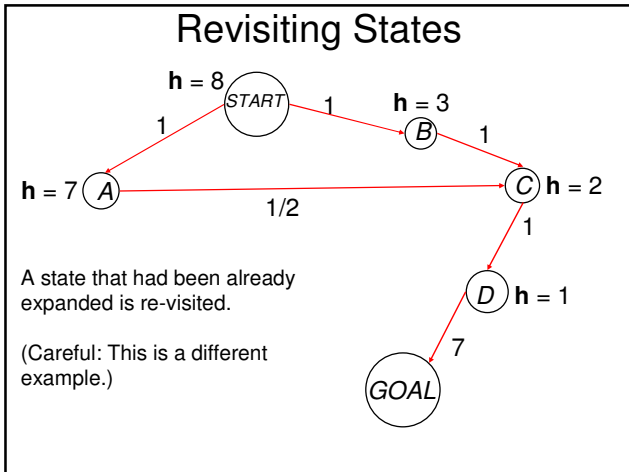


A* Termination Condition



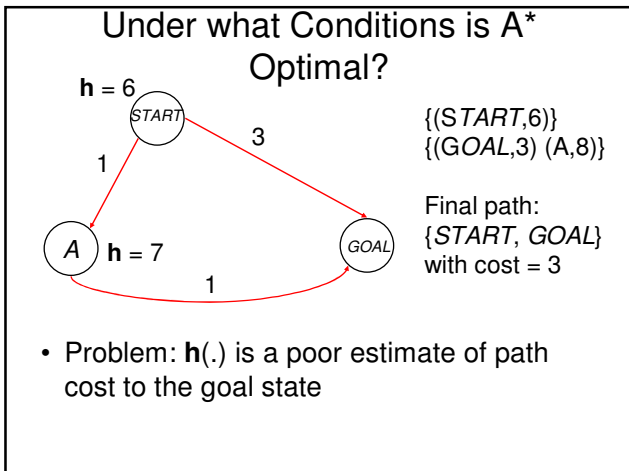
Revisiting States





Pop state s with lowest $f(s)$ in queue
 If $s = GOAL$
 return *SUCCESS*
 Else expand s :
 For all s' in **succs** (s):
 $f' = g(s') + h(s') = g(s) + \text{cost}(s,s') + h(s')$
 If (s' not seen before OR
 s' previously expanded with $f(s') > f'$ OR
 s' in PQ with $f(s') > f'$)
 Promote/Insert s' with new value f' in PQ
 previous(s') $\leftarrow s$
 Else
 Ignore s' (because it has been visited and its current path cost $f(s')$ is still the lowest path cost from *START* to s')

A* Algorithm
(inside loop)




Admissible Heuristics

- Define $h^*(s)$ = the true minimal cost to the goal from s
- h is admissible if

$h(s) \leq h^*(s)$ for all states s
- In words: An admissible heuristic never overestimates the cost to the goal.
 "Optimistic" estimate of cost to goal.


A* is guaranteed to find the optimal path if h is admissible

Consistent (Monotonic) Heuristics



$$h(s) \leq h(s') + \text{cost}(s,s')$$

Consistent (Monotonic) Heuristics



Sort of triangular inequality implies that path cost always increases + need to expand node only once

$$h(s) \leq h(s') + \text{cost}(s,s')$$

Pop state s with lowest $f(s)$ in queue

If $s = \text{GOAL}$

return *SUCCESS*

Else expand s :

For all s' in **succs** (s):

$$f' = g(s') + h(s') = g(s) + \text{cost}(s,s') + h(s')$$

If (s' not seen before OR

~~s' previously expanded with $f(s') > f'$ OR~~
 s' in PQ with with $f(s') > f'$)


Promote/Insert s' with new value f' in PQ
previous(s') $\leftarrow s$

Else

Ignore s' (because it has been visited and its current path cost $f(s')$ is still the lowest path cost from *START* to s')


If h is consistent

Examples



For the navigation problem:
 The length of the shortest path is at least the distance between s and *GOAL* \rightarrow Euclidean distance is an admissible heuristic

What about the puzzle?



s

GOAL

$h_1(s) = 7$

$h_2(s) = 2 + 3 + 3 + 2 + 4 + 2 + 0 + 2 = 18$

Comparing Heuristics

	L = 4 steps	L = 8 steps	L = 12 steps
$h_1 =$ misplaced tiles			
Iterative Deepening	112	6,300	3.6×10^6
$h_2 =$ Manhattan distance			
A* with heuristic h_1	13	39	227
A* with heuristic h_2	12	25	73

- Overestimates A* performance because of the tendency of IDS to expand states repeatedly
- Number of states expanded does not include $\log()$ time access to queue

Example from Russell&Norvig

Comparing Heuristics

s

GOAL

$h_1(s) = 7$

$h_2(s) = 2 + 3 + 3 + 2 + 4 + 2 + 0 + 2 = 18$

h_2 is larger than h_1 and, at same time, A* seems to be more efficient with h_2 .

Is there a connection between these two observations?

h_2 dominates h_1 if $h_2(s) \geq h_1(s)$ for all s

For any two heuristics h_2 and h_1 :
 If h_2 dominates h_1 then A* is more efficient (expands fewer states) with h_2

Intuition: since $h \leq h^*$, a larger h is a better approximation of the true path cost

Limitations

- Computation: In the worst case, we may have to explore all the states $\rightarrow O(N)$
- The good news: A* is optimally efficient \rightarrow For a given $h(\cdot)$, no other optimal algorithm will expand fewer nodes
- The bad news: Storage is also potentially exponential $\rightarrow O(N)$

IDA* (Iterative Deepening A*)

- Same idea as Iterative Deepening DFS except use $f(s)$ to control depth of search instead of the number of transitions
- Example, assuming integer costs:
 1. Run DFS, stopping at states s such that $f(s) > 0$
Stop if goal reached
 2. Run DFS, stopping at states s such that $f(s) > 1$
Stop if goal reached
 3. Run DFS, stopping at states s such that $f(s) > 2$
Stop if goal reached.....Keep going by increasing the limit on f by 1 every time
- Complete (assuming we use loop-avoiding DFS)
- Optimal
- More expensive in computation cost than A*
- Memory order L as in DFS

Summary

- Informed search and heuristics
- First attempt: Best-First Greedy search
- A* algorithm
 - Optimality
 - Condition on heuristic functions
 - Completeness
 - Limitations, space complexity issues
 - Extensions

Nils Nilsson. Problem Solving Methods in Artificial Intelligence. McGraw Hill (1971)
Judea Pearl. Heuristics: Intelligent Search Strategies for Computer Problem Solving (1984)
Chapters 3&4 Russel & Norvig