

Informed Search and Exploration

Sections 3.5 and 3.6





- Best-first search
- A* search
- Heuristics, pattern databases
- IDA* search
- (Recursive Best-First Search (RBFS), MA* and SMA* search)

Best-first search

- Idea: use an evaluation function for each node
- The evaluation function is an estimate of "desirability"
- Expand the most desirable unexpanded node
- The desirability function comes from domain knowledge
- Implementation: The *frontier* is a queue sorted in decreasing order of desirability
- Special cases:
 - greedy best first search
 - A* search

Romania with step costs in km



Sample straight line distances to Bucharest: Arad: 366, Bucharest: 0, Sibiu: 253, Timisoara: 329.



Greedy best-first search example



After expanding Arad



After expanding Sibiu



After expanding Fagaras



The goal Bucharest is found with a cost of 450. However, there is a better solution through Pitesti (h = 417).

Properties of greedy best-first search

- Complete No can get stuck in loops For example, going from Iasi to Fagaras, Iasi → Neamt → Iasi → Neamt → ... Complete in finite space with repeated-state checking
- Time $O(b^m)$, but a good heuristic can give dramatic improvement (more later)
- Space $O(b^m)$ —keeps all nodes in memory
- Optimal No

(For example, the cost of the path found in the previous slide was 450. The path Arad, Sibiu, Rimnicu Vilcea, Pitesti, Bucharest has a cost of 140+80+97+101 = 418.)



- Idea: avoid expanding paths that are already expensive
- **•** Evaluation function f(n) = g(n) + h(n)
 - $g(n) = exact \operatorname{cost} so far to reach n$
 - h(n) = estimated cost to goal from n
 - f(n) = estimated total cost of path through n to goal
- A* search uses an *admissible* heuristic i.e., $h(n) \le h^*(n)$ where $h^*(n)$ is the *true* cost from *n*. (Also require $h(n) \ge 0$, so h(G) = 0 for any goal *G*.)
- Straight line distance $(h_{SLD}(n))$ is an admissible heuristic because never overestimates the actual road distance.





366=0+366

After expanding Arad



After expanding Sibiu



After expanding Rimnicu Vilcea



Ch. 03 – p.15/4

After expanding Fagaras



Remember that the goal test is performed when a node is selected for expansion, not when it is generated.

After expanding Pitesti



Ch. 03 – p.17/4

Theorem: A* search is optimal.

Note that, A* search uses an admissible heuristic by definition.

Suppose some suboptimal goal G_2 has been generated and is in the queue. Let *n* be an unexpanded node on a shortest path to an optimal goal G_1 .

Optimality of A* for trees (cont'd)



Since $f(n) < f(G_2)$, A^{*} will never select G_2 for expansion.



Note that h is admissible, it never overestimates.

Ch. 03 – p.20/4'



The root node was expanded. Note that f decreased from 6 to 4.



The suboptimal path is being pursued.



Goal found, but we cannot stop until it is selected for expansion.



The node with f = 7 is selected for expansion.



The optimal path to the goal is found.

Consistency

A heuristic is consistent if $h(n) \leq c(n, a, n') + h(n')$

If h is consistent, we have



I.e., f(n) is nondecreasing along any path.

Optimality of A* for graphs

- Lemma: A* expands nodes in order of increasing f value
- Gradually adds "f-contours" of nodes (cf. breadth-first adds layers) Contour i has all nodes with $f = f_i$, where $f_i < f_{i+1}$
- With uniform-cost search (A* search with h(n)=0) the bands are "circular".

With a more accurate heuristic, the bands will stretch toward the goal and become more narrowly focused around the optimal path.





Performance of A*

- The *absolute error* of a heuristic is defined as $\Delta \equiv h^* h$
- The *relative error* of a heuristic is defined as $\epsilon \equiv \frac{h^* h}{h^*}$
- Complexity with constant step costs: $O(b^{\epsilon d})$
- Problem: there can be exponentially many states with $f(n) < C^*$ even if the absolute error is bounded by a constant

Properties of A*

- Complete Yes, unless there are infinitely many nodes with $f \le f(G)$
- Time Exponential in (relative error in $h \times$ length of solution)
- Space Keeps all nodes in memory
- Optimal Yes—cannot expand f_{i+1} until f_i is finished
 - A* expands all nodes with $f(n) < C^*$
 - A* expands some nodes with $f(n) = C^*$
 - A* expands no nodes with $f(n) > C^*$

E.g., for the 8-puzzle: $h_1(n)$ = number of misplaced tiles $h_2(n)$ = total *Manhattan* distance (i.e., no. of squares from desired location of each tile)





Start State

Goal State

 $h_1(S) = ??$ $h_2(S) = ??$

Dominance

If $h_2(n) \ge h_1(n)$ for all *n* (both admissible) then h_2 *dominates* h_1 and is better for search

Typical search costs:

$$\begin{array}{ll} d = 14 & {\sf IDS} = 3,473,941 \mbox{ nodes} \\ & {\sf A}^*(h_1) = 539 \mbox{ nodes} \\ & {\sf A}^*(h_2) = 113 \mbox{ nodes} \\ d = 24 & {\sf IDS} \approx 54,000,000,000 \mbox{ nodes} \\ & {\sf A}^*(h_1) = 39,135 \mbox{ nodes} \\ & {\sf A}^*(h_2) = 1,641 \mbox{ nodes} \end{array}$$

The effect is characterized by the effective branching factor (b^*)

- If the total number of nodes generated by A^* is N and
- the solution depth is d,
- then *b* is branching factor of a uniform tree, such that $N+1 = 1 + b + (b)^2 + (b)^d$

A well designed heuristic has a *b* close to 1.

Using relaxed problems to find heuristics

- Admissible heuristics can be derived from the *exact* solution cost of a *relaxed* version of the problem
- If the rules of the 8-puzzle are relaxed so that a tile can move *anywhere*, then $h_1(n)$ gives the shortest solution
- If the rules are relaxed so that a tile can move to *any adjacent* square, then $h_2(n)$ gives the shortest solution
- Key point: the optimal solution cost of a relaxed problem is no greater than the optimal solution cost of the real problem

Well-known example: *travelling salesperson problem (TSP)* Find the shortest tour visiting all cities exactly once



Minimum spanning tree can be computed in $O(n^2)$ and is a lower bound on the shortest (open) tour



- Admissible heuristics can also be generated from the solution cost of sub- problems.
- For example, in the 8-puzzle problem a sub-problem of getting the tiles 2, 4, 6, and 8 into position is a lower bound on solving the complete problem.
- Pattern databases store the solution costs for all the sub-problem instances.
- The choice of sub-problem is flexible: for the 8-puzzle a subproblem for 2,4,6,8 or 1,2,3,4 or 5,6,7,8, . . . could be created.

Iterative Deepening A* (IDA*)

- Idea: perform iterations of DFS. The cutoff is defined based on the *f*-cost rather than the depth of a node.
- Each iteration expands all nodes inside the contour for the current *f*-cost, peeping over the contour to find out where the contour lies.

function IDA* (problem)
returns a solution sequence

inputs: *problem*, a problem local variables:

f-limit, the current *f*-COST limit *root*, a node

root ← MAKE-NODE(INITIAL-STATE[problem]) f-limit ← f-Cost(root) loop do solution, f-limit ← DFS-Contour(root, f-limit) if solution is non-null then return solution

if *f-limit* = ∞ then return failure

function DFS-CONTOUR (*node, f-limit*) **returns** a solution sequence and a new *f*-COST limit

inputs: node, a node *f-limit*, the current *f*-COST limit local variables: *next-f*, the *f*-COST limit for the next contour, initally ∞

if f-Cost[node] > f-limit then return null, f-Cost[node]
if GOAL-TEST[problem](STATE[node]) then return node, f-limit
for each node s in SUCCESSORS(node) do
 solution, new-f ← DFS-CONTOUR(s, f-limit)
 if solution is non-null then return solution, f-limit
 next-f ← MIN(next-f, new-f)
return null, next-f

Complete Yes, similar to A*.

- Time Depends strongly on the number of different values that the heuristic value can take on. 8-puzzle: few values, good performance TSP: the heuristic value is different for every state. Each contour only includes one more state than the previous contour. If A* expands N nodes, IDA* expands $1 + 2 + ... + N = O(N^2)$ nodes.
- Space It is DFS, it only requires space proportional to the longest path it explores. If δ is the smallest operator cost, and f^* is the optimal solution cost, then IDA* will require bf^*/δ nodes.
- Optimal Yes, similar to A*

Recursive best-first search (RBFS)

- Idea: mimic the operation of standard best-first search, but use only linear space
- Runs similar to recursive depth-first search, but rather than continuing indefinitely down the current path, it uses the *f_limit* variable to keep track of the best alternative path available from any ancestor of the current node.
- If the current node exceeds this limit, the recursion unwinds back to the alternative path.



function RECURSIVE-BEST-FIRST-SEARCH (problem) returns a solution or failure return RBFS(problem, MAKE-NODE(problem.INITIAL-STATE), ∞)

function RBFS (*problem, node, f_limit*)

returns a solution or failure and a new *f*-cost limit

if *problem*.GOAL-TEST(*node*.STATE) **then return** SOLUTION(*node*) *successors* ← []

for each action in problem. ACTIONS (node. STATE) do add CHILD-NODE (problem, node, action) into successors if successors is empty then return failure, ∞ for each s in successors do

/* update f with value from previous search, in any */

 $s.f \gets \max\left(s.g + s.h, node.f\right)\right)$

loop do

best \leftarrow the lowest *f*-value in *successors*

if best.f > f-limit then return failure, best.f

alternative \leftarrow the second lowest *f*-value among successors result, best.*f* \leftarrow RBFS (problem, best, min(*f*-limit, alternative)) if result \neq failure then return result

Progress of RBFS





- Idea: use all the available memory IDA* remembers only the current *f*-cost limit RBFS uses linear space
- Proceeds just like A*, expanding the best leaf until the memory is full. When the memory if full, drop the worst leaf node.



- The evaluation function for a node n is: f(n) = g(n) + h(n)
- If only g(n) is used, we get uniform-cost search
- If only h(n) is used, we get greedy best-first search
- If both g(n) and h(n) are used we get best-first search
- If both g(n) and h(n) are used with an admissible heuristic we get A* search
- A consistent heuristic is admissible but not necessarily vice versa



- Admissibility is sufficient to guarantee solution optimality for tree search
- Consistency is required to guarantee solution optimality for graph search
- If an admissible but not consistent heuristic is used for graph search, we need to adjust path costs when a node is rediscovered
- Heuristic search usually brings dramatic improvement over uninformed search
- Keep in mind that the f-contours might still contain an exponential number of nodes