

Informed Searches

CHAPTER 3 CONTINUED
 COSC 370
 SPRING 2013
 ALAN C. JAMIESON
 SOME SLIDE CONTENT FROM RUSSELL &
 NORVIG PROVIDED SLIDES

- Greedy Best-First Search
- A*
- Improvements to A*

Review: Tree Search

```
function TREE-SEARCH(problem, fringe) returns a solution, or failure
  fringe ← INSERT(MAKE-NODE(INITIAL-STATE[problem]), fringe)
  loop do
    if fringe is empty then return failure
    node ← REMOVE-FRONT(fringe)
    if GOAL-TEST[problem] applied to STATE(node) succeeds return node
    fringe ← INSERT-ALL(EXPAND(node, problem), fringe)
```

A strategy is defined by picking the **order of node expansion**

Informed Searches

- a.k.a Heuristic Search
- Expansion based on a function $f(n)$ – evaluation function.
- A part of this function may be a heuristic $h(n)$:
 $h(n)$ = estimated cost of cheapest path from n to goal
- Constant: $h(n) = 0$ for goal state

Best-First Search


- Evaluation Function for each node – estimating desirability.
- Expand most desirable unexpanded node
- Implementation – use a queue! Sort by desirability.
- Specific Cases:
 - Greedy
 - A*

Greedy Best-first Search

- Greedy – expand the node that is closest to the goal.
- $f(n) = h(n) = \text{estimate}$ of cheapest path
- Example: Romanian Vacation Problem

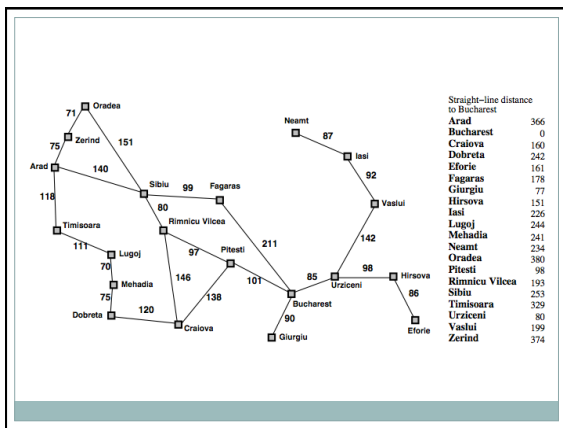
h_{sld} = Straight line distance between two cities

- Take home – Greedy expands the node that appears to be the closest to the goal.



Analysis of Technique

- Complete – nope, can run into loops (example: take RVP with Oradea as goal)



Analysis of Technique

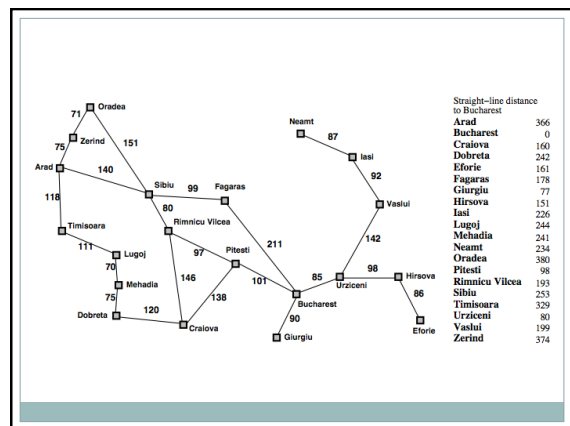
- Complete – nope, can run into loops (example: take RVP with Oradea as goal)
- Time – $O(b^m)$, though heuristics can improve this
- Space – $O(b^m)$, all states in memory
- Optimal – not necessarily

A*

- An extension of Dijkstra's shortest path algorithm developed by Hart, Nilsson, and Raphael (SRI).
- General Idea: try not to go down the expensive paths!
- Evaluation function:

$$f(n) = h(n) + g(n)$$

$h(n)$ = estimated cost from n to goal
 $g(n)$ = cost so far to reach n
- A* uses an admissible heuristic – one that never overestimates the cost to reach the goal.



Exercise!

- Here's a starter 8-puzzle. How would A* solve this puzzle? What's the heuristic?

4	3	7
2		1
5	6	8

Analysis of Technique

- Complete – yes, assuming finite number of states
- Time – potentially exponential – why so long?
 - while efficient and optimal, the number of potential states (think: paths) that have to be expanded/searched is still exponential.
- Space – $O(b^m)$, all states in memory
- Optimal – yes!

Some additional terminology

- Dominance – assuming that we have admissible heuristics, if $h_2(n) \geq h_1(n)$ for any node n , then $h_2(n)$ dominates $h_1(n)$ and is “better”.
- Consistency (aka monotonicity) – a heuristic is consistent if, for every node n and every successor n' of n generated by any action a , the estimated cost of reaching the goal from n is no greater than the step cost of getting to n' plus the estimated cost of reaching the goal from n' .
- Relaxed problems

Memory-bounded Heuristic

- One issue with A* - memory usage!
- Briefly – Iterative-Deepening A* (IDA*)
 - Iterative Deepening Search (see last set of slides)
 - Cutoff used – $f(n)$ rather than the depth
- Recursive best-first search (RBFS) – simple recursive algorithm for BFS in linear space.
- Memory-Bounded A* (MA*) & Simplified MA* (SMA*) – A* until memory is full, then expand by dropping the “worst” leaf node.