

Foundations of Artificial Intelligence

CS472/3

Lecture #5

Bart Selman

Slide CS472-1

Today's Lecture

Informed Search

Readings: R&N, Chapter 4.

Slide CS472-2

Informed Methods: Heuristic Search

Informed Methods use problem-specific knowledge.

Heuristic search is an attempt to search the most promising paths first. Uses heuristics, or rules of thumb, to find the best node to expand next.

Relies on an *evaluation function* — indicates the desirability of expanding a node. E.g., *path cost*.

$h(n)$ = estimated cost of the cheapest path from the state at node n to a goal state (*Heuristic Function*: e.g., straight line on a map)

Slide CS472–3

General Principle

Given a list of nodes to be expanded, choose the one that the heuristic function estimates as the most promising.

Slide CS472–4

Best-First Search

1. Set L to be the initial node(s).
2. Let n be the node on L that is “most promising” according to eval function ($h(n) / f(n)$). If L is empty, fail.
3. If n is a goal node, stop and return it (and the path from the initial node to n).
4. Otherwise, remove n from L and add all of n 's children to L (labeling each with its path from the initial node). Return to step 2.

Slide CS472–5

Two Instantiations of Best-First Search

Issue: Best-First depends on *evaluation function*.

Alternatives:

Greedy Search minimize estimated cost to reach the goal, i.e., expand the node “closest” to the goal.

A* minimize total estimated path cost to reach the goal, i.e., expand the node on the “least-cost” solution path to the goal.

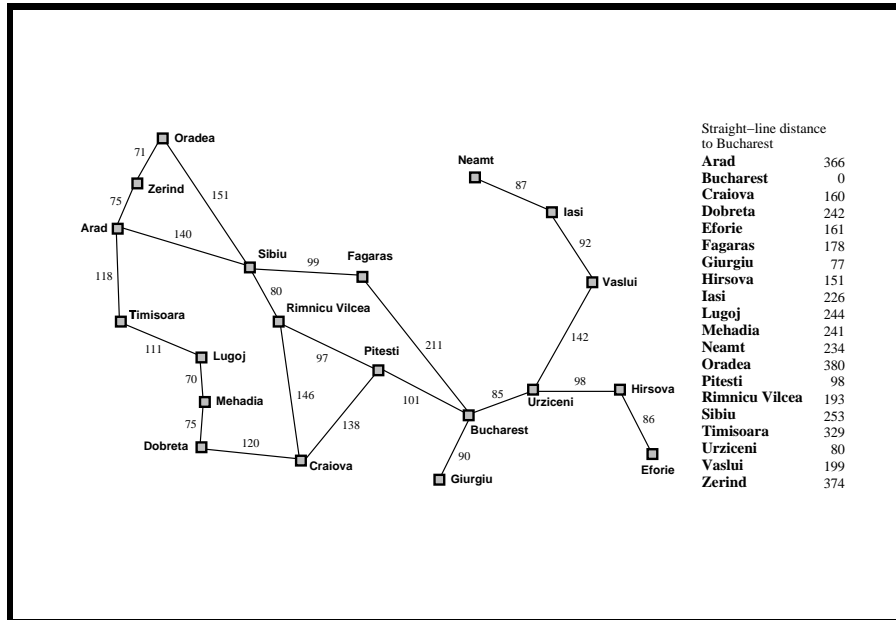
Slide CS472–6

Greedy Search


$h(n)$ = Estimated cost from node n to nearest goal

1. Set L to be the initial node(s).
2. Let n be the node on L that minimizes $h(n)$. If L is empty, fail.
3. If n is a goal node, stop and return it (and the path from the initial node to n).
4. Otherwise, remove n from L and add all of n 's children to L (labeling each with its path from the initial node). Return to step 2.

Slide CS472-7



Slide CS472-8



Slide CS472-9

Problem: Too Greedy

From Arad to Sibiu to Fagaras — but to Rimnicu would have been better.

Need to consider: cost of getting from start node (Arad) to intermediate nodes!

Slide CS472-10

Evaluation Function for A* Algorithm

Goal:

Find the *shallowest* goal as quickly as possible.

$g(n)$ Cost of reaching node n from start node

$h(n)$ Estimated cost from node n to nearest goal

New evaluation function:

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

$f(n)$ Estimated cost of cheapest solution through n

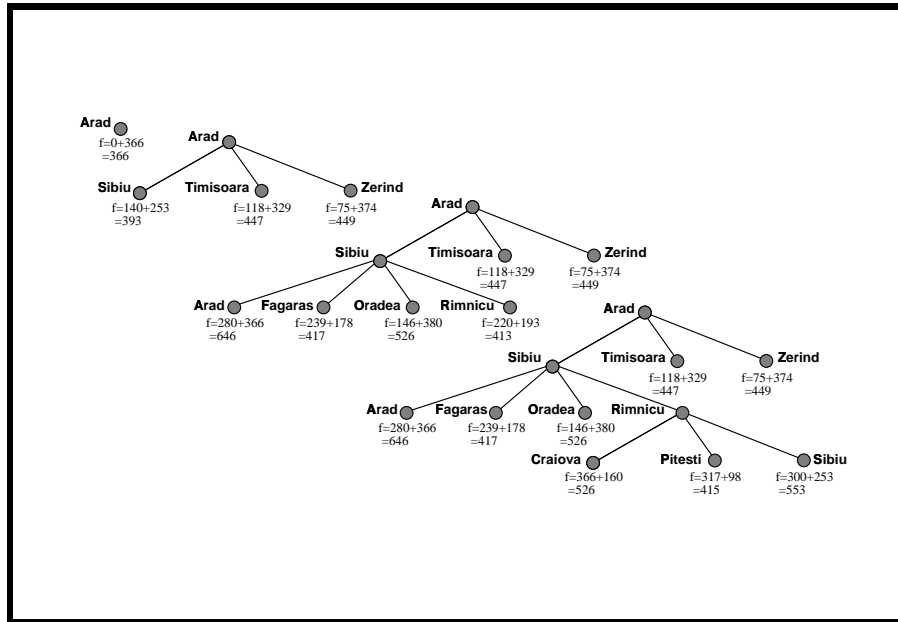
Slide CS472–11

A*

$f(n)$ = Estimated cost of cheapest solution through n

1. Set L to be the initial node(s).
2. Let n be the node on L that minimizes to $f(n)$. If L is empty, fail.
3. If n is a goal node, stop and return it (and the path from the initial node to n).
4. Otherwise, remove n from L and add all of n 's children to L (labeling each with its path from the initial node). Return to step 2.

Slide CS472–12



Slide CS472-13

Finds Optimal Path

Now expands Rimnicu ($f = (140 + 80) + 193 = 413$) over Faragas ($f = (140 + 99) + 178 = 417$).

Q. What if $h(\text{Faragas}) = 170$ (also an underestimate)?

Slide CS472-14

Need Some Conditions

To guarantee that A^* finds an optimal solution,
we need that h **never overestimates** the cost
of reaching the goal.

Called an *admissible heuristics*.

Transfers to f , i.e., f also doesn't overestimate.

Slide CS472–15

Let's also also assume (true for most admissible heuristics)
that f is **monotonic**, i.e., along any path from
the root f never decreases.

Can often modify heuristic to become monotonic.

E.g. let n be parent of n' . Suppose that

$g(n) = 3$ and $h(n) = 4$, so $f(n) = 7$.

and $g(n') = 4$ and $h(n') = 2$, so $f(n') = 6$.

But because any path through n' is also a path through n ,
we can set $f(n') = 7$.

Slide CS472–16

In effect: $f(n') = \max(f(n), g(n') + h(n'))$.

Called the **path-max** equation.

(ignores misleading numbers in heuristic)

Slide CS472–17

Intuition A*


Let f^* be the cost of the *optimal* solution path.

We have:

A* expands *all* nodes with $f(n) < f^*$

A* may then expand some nodes right on “goal contour”,
with $f(n) = f^*$ before selecting a goal node.

Slide CS472–18



Slide CS472-19