

Heuristic Search

- **Idea:** don't ignore the goal when selecting paths.
- Often there is extra knowledge that can be used to guide the search: **heuristics.**
- **$h(n)$** is an estimate of the cost of the shortest path from node n to a goal node.
- $h(n)$ uses only readily obtainable information (that is easy to compute) about a node.
- h can be extended to paths: $h(\langle n_0, \dots, n_k \rangle) = h(n_k)$.
- $h(n)$ is an underestimate if there is no path from n to a goal that has path length less than $h(n)$.

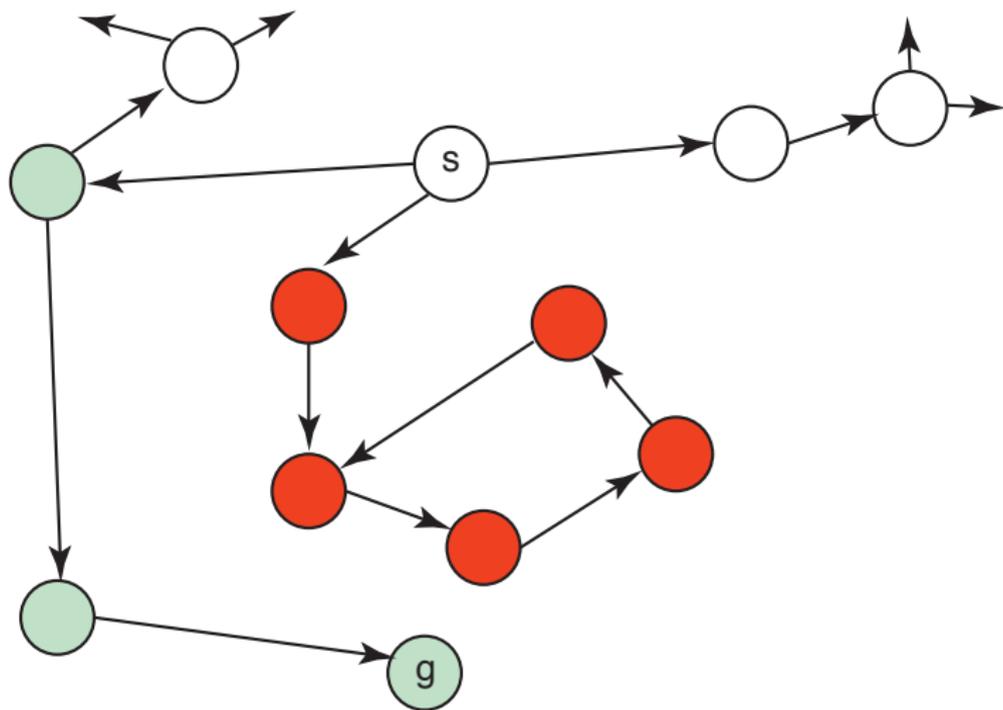
Example Heuristic Functions

- If the nodes are points on a Euclidean plane and the cost is the distance, we can use the straight-line distance from n to the closest goal as the value of $h(n)$.
- If the nodes are locations and cost is time, we can use the distance to a goal divided by the maximum speed.
- If the goal is to collect all of the coins and not run out of fuel, the cost is an estimate of how many steps it will take to collect the rest of the coins, refuel when necessary, and return to goal position.

Best-first Search

- **Idea:** select the path whose end is closest to a goal according to the heuristic function.
- Best-first search selects a path on the frontier with minimal h -value.
- It treats the frontier as a priority queue ordered by h .

Illustrative Graph — Best-first Search



Complexity of Best-first Search

- It uses space exponential in path length.
- It isn't guaranteed to find a solution, even if one exists.
- It doesn't always find the shortest path.

A* Search Algorithm

- A* is a mix of lowest-cost-first and best-first search.
- It treats the frontier as a priority queue ordered by $f(p)$.
- It always selects the node on the frontier with the lowest estimated distance from the start to a goal node constrained to go via that node.

If there is a solution, A^* always finds an optimal solution —the first path to a goal selected— if

- the branching factor is finite
- arc costs are bounded above zero (there is some $\epsilon > 0$ such that all of the arc costs are greater than ϵ), and
- $h(n)$ is an underestimate of the length of the shortest path from n to a goal node.

Why is A^* admissible?

- If a path p to a goal is selected from a frontier, can there be a shorter path to a goal?
- Suppose path p' is on the frontier. Because p was chosen before p' , and $h(p) = 0$:

$$\text{cost}(p) \leq \text{cost}(p') + h(p').$$

- Because h is an underestimate

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

for any path p'' to a goal that extends p'

- So $\text{cost}(p) \leq \text{cost}(p'')$ for any other path p'' to a goal.

Why is A^* admissible?

- There is always an element of an optimal solution path on the frontier before a goal has been selected. This is because, in the abstract search algorithm, there is the initial part of every path to a goal.
- A^* halts, as the costs of the paths on the frontier keeps increasing, and will eventually exceed any finite number.