Search Heuristics

- A heuristic is:
  - A function that *estimates* how close a state is to a goal
  - Designed for a particular search problem
  - Examples: Manhattan distance, Euclidean distance for pathing

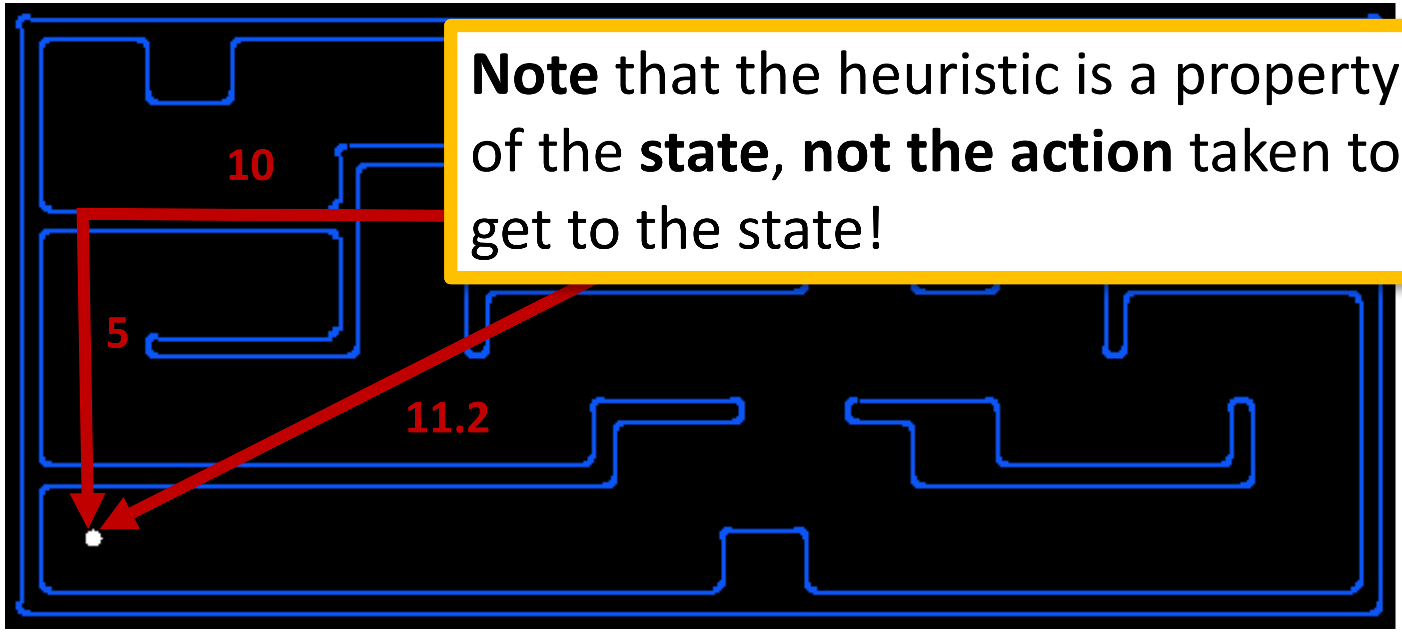

Note that the heuristic is a property of the **state**, **not the action** taken to get to the state!

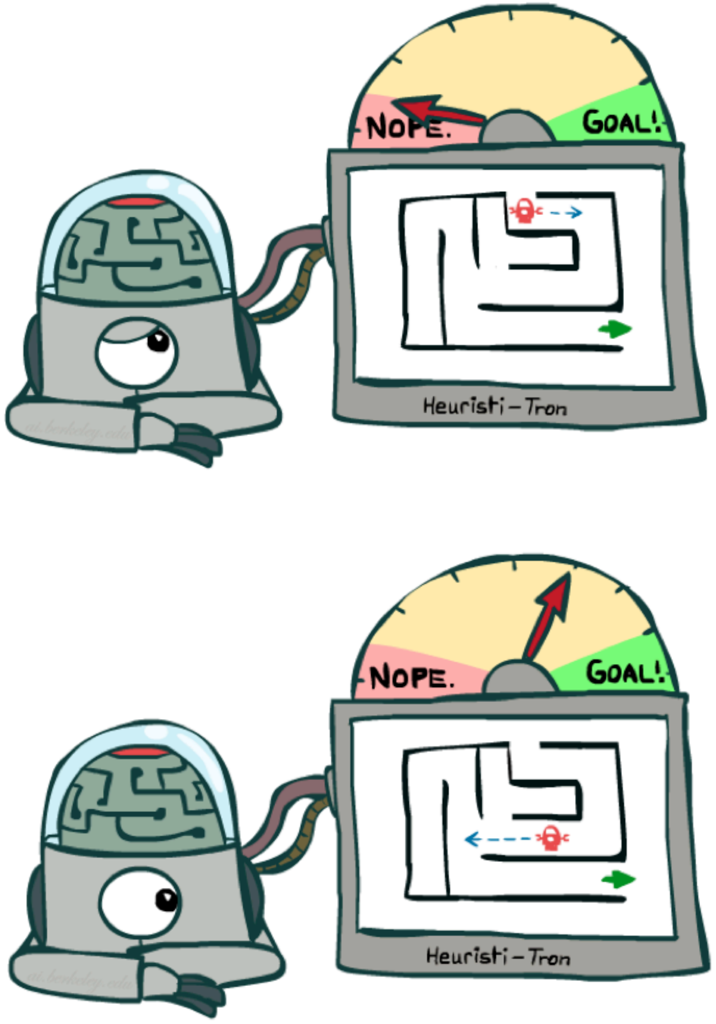

Pancake Heuristics

Heuristic 1: the number of the largest pancake that is still out of place

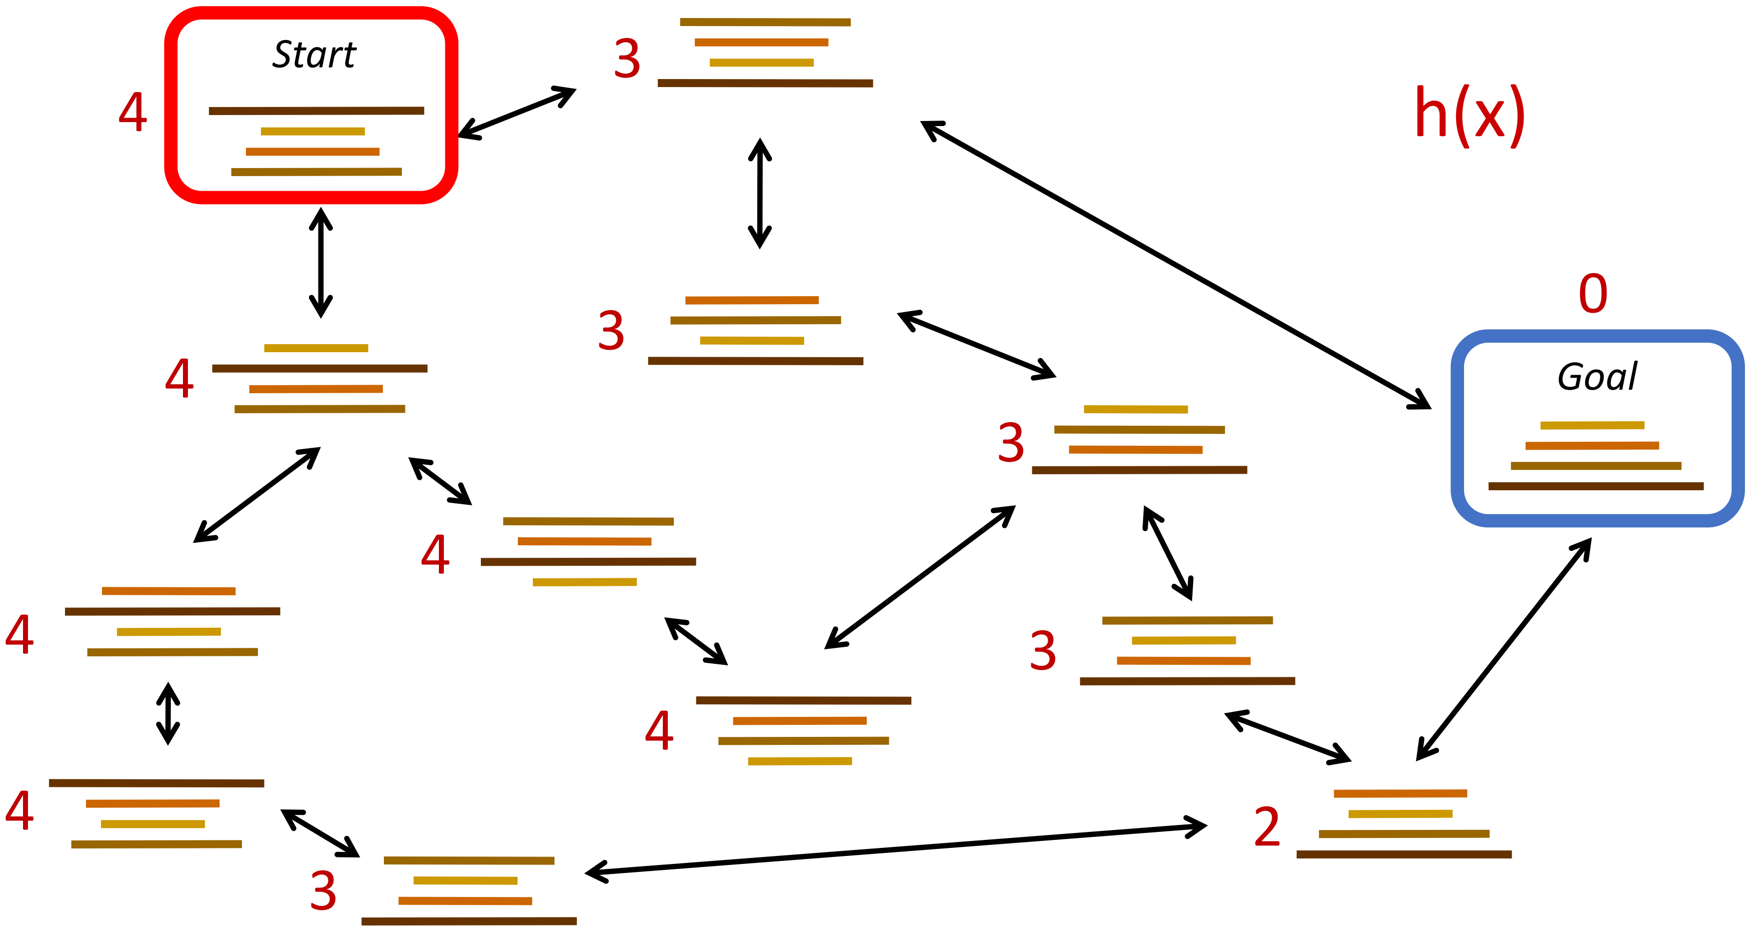

Pancake Heuristics

Heuristic 2: how many pancakes are on top of a smaller pancake?

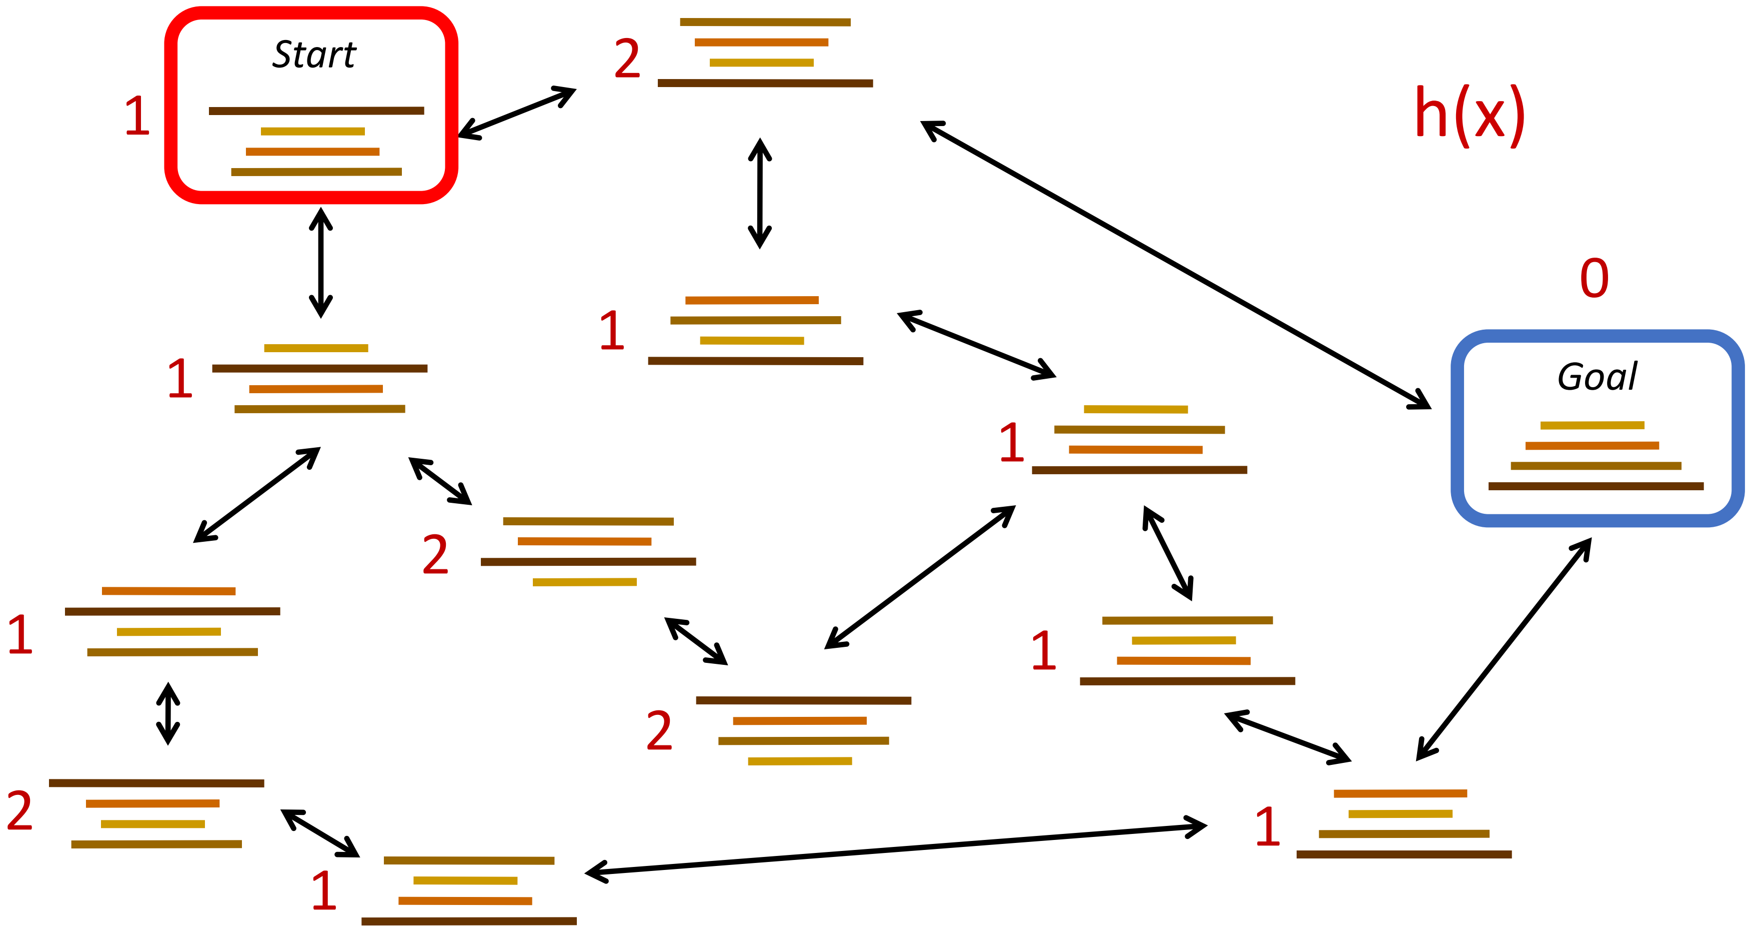

Pancake Heuristics

<u>Heuristic 3</u>: All zeros (aka *null heuristic*, or ”I like waffles better anyway”)

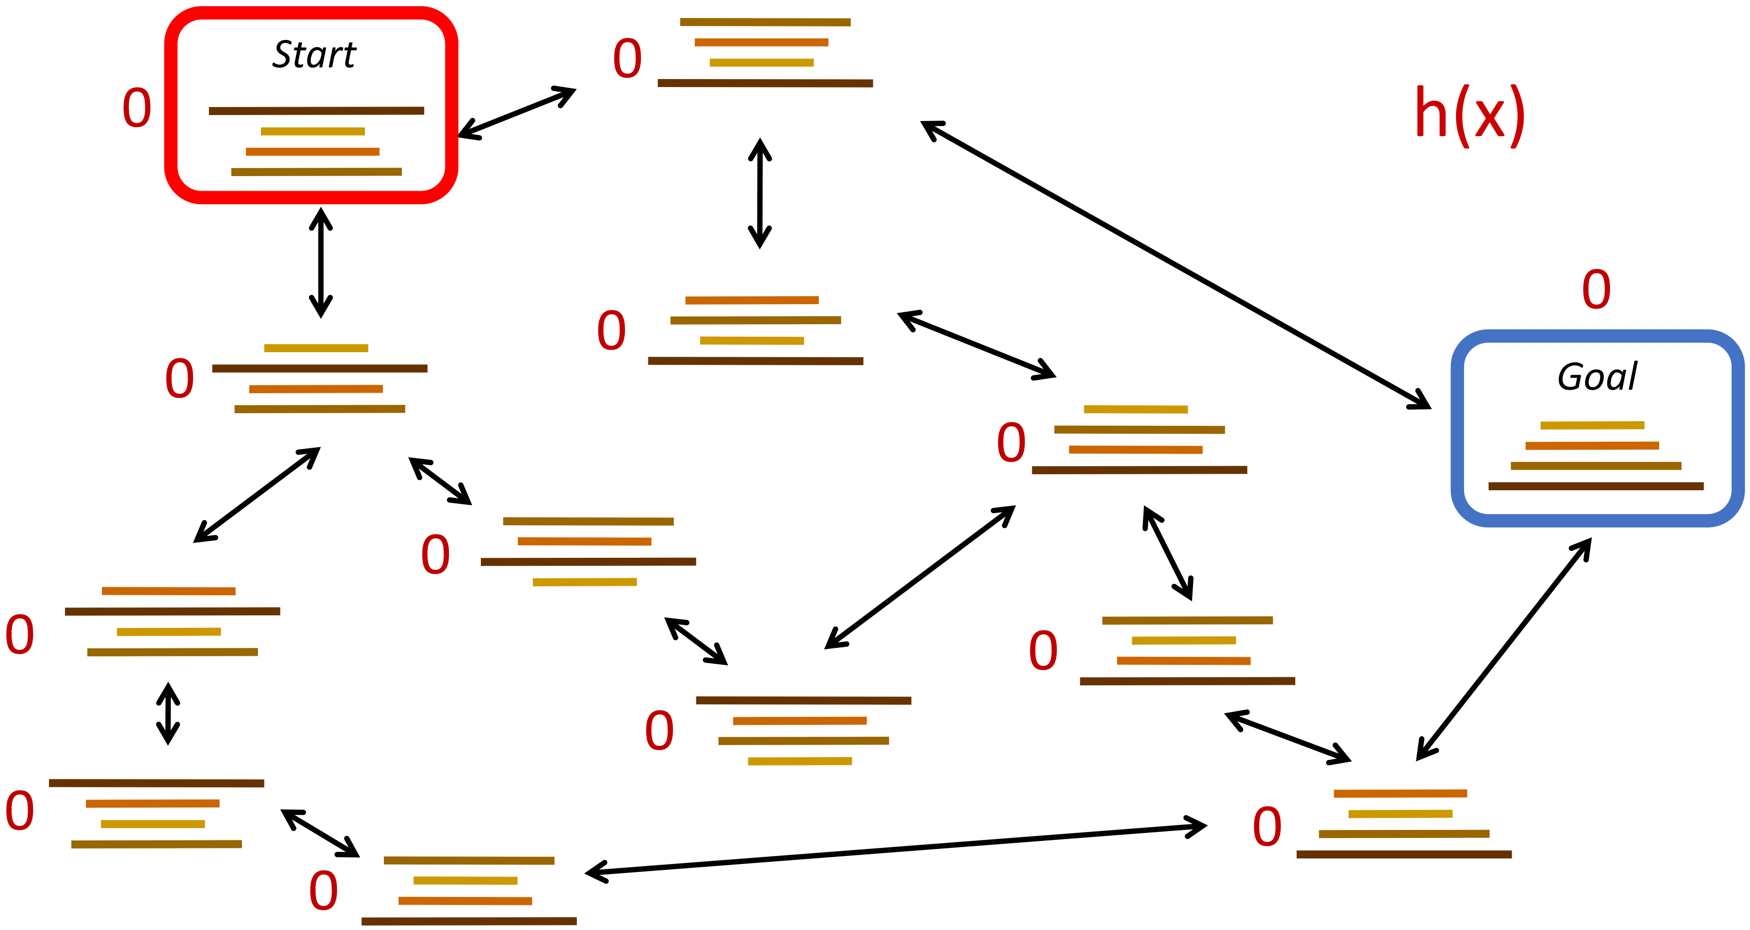

Greedy Search

Straight-line Heuristic in Romania

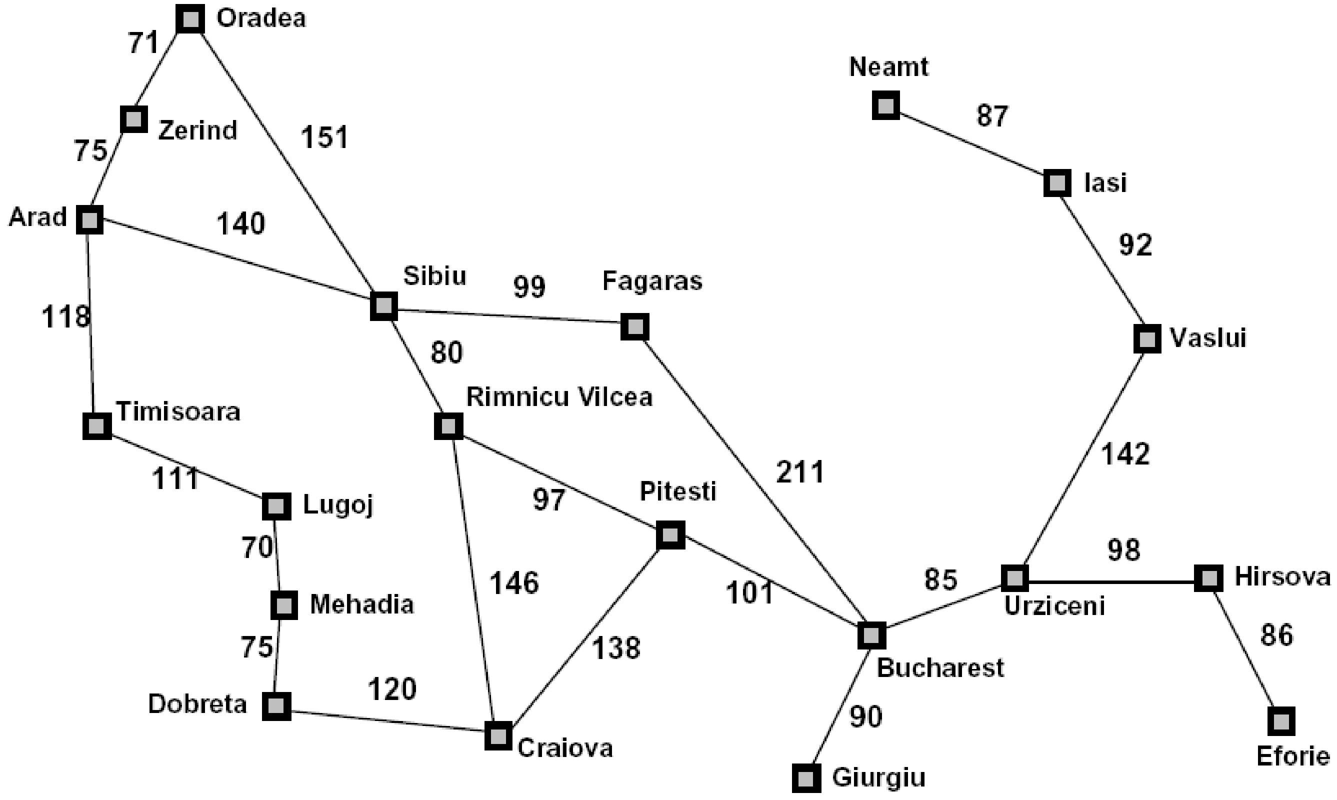

Straight−line distance to Bucharest

Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	178
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	98
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

h(x)

Greedy Straight-Line Search in Romania

- Expand the node that seems closest…

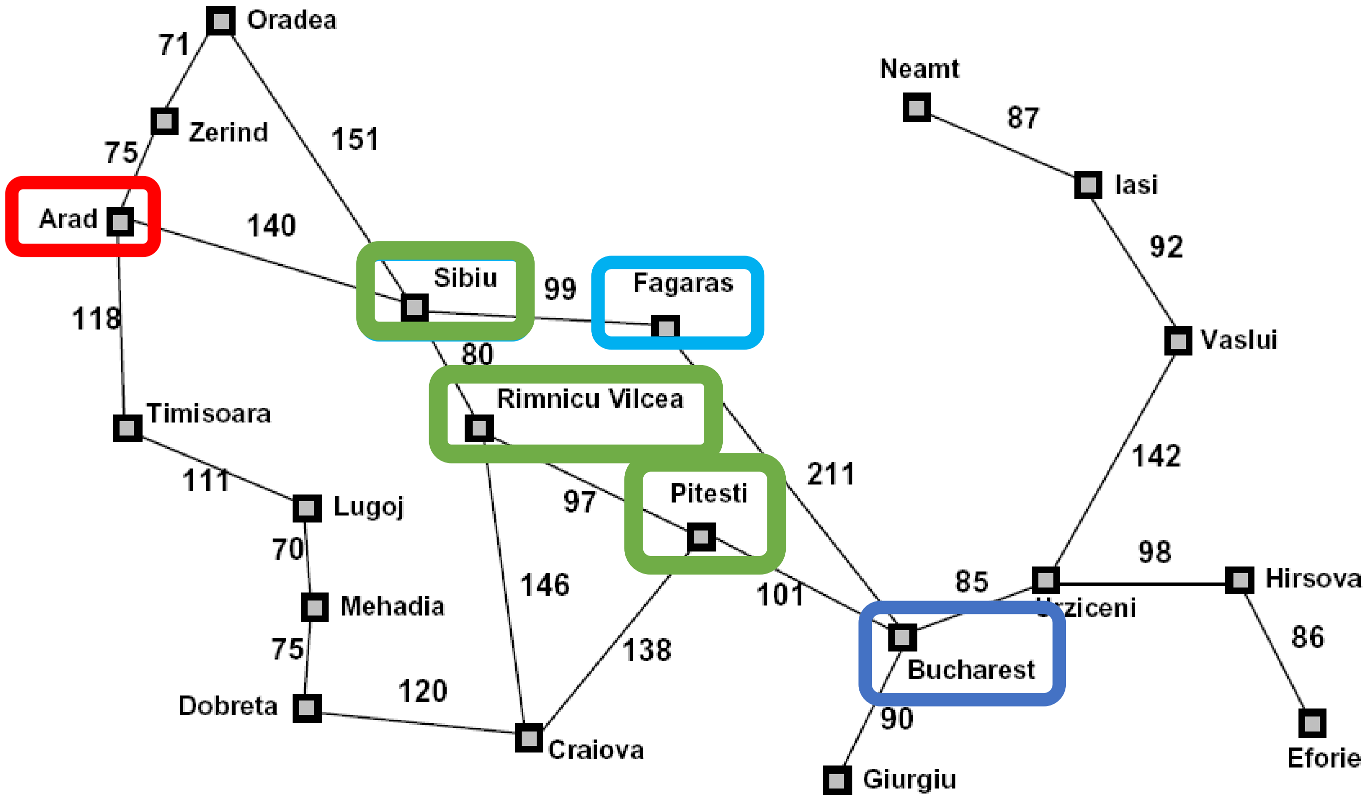

Greedy

- Cost: 450

Optimal

- Cost: 418

Straight−line distance to Bucharest

Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	178
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	98
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

h(x)

Greedy Search

- Strategy: expand a node that you think is closest to a goal state
  - Heuristic: estimate of distance to nearest goal for each state

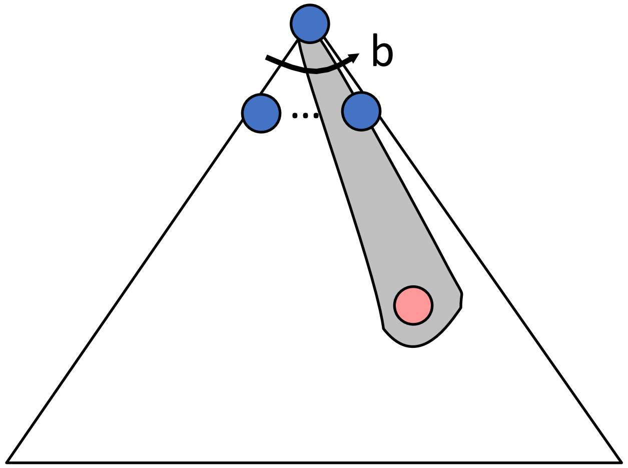

- A common case:
  - Best-first takes you straight to the (<u>non-optimal</u>) goal

- Worst-case: like a badly-guided DFS

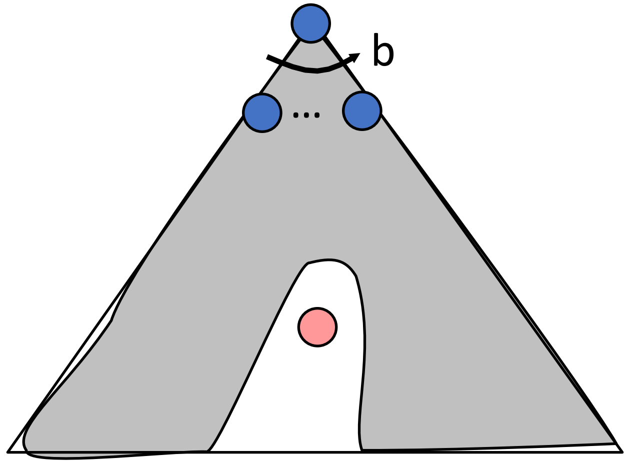

- What goes wrong?
  - Doesn't take <u>real</u> path cost into account

A* Search

Combining UCS and Greedy

- Uniform-cost orders by path cost, or *backward cost* g(n)
- Greedy orders by goal proximity, or *forward cost* h(n)

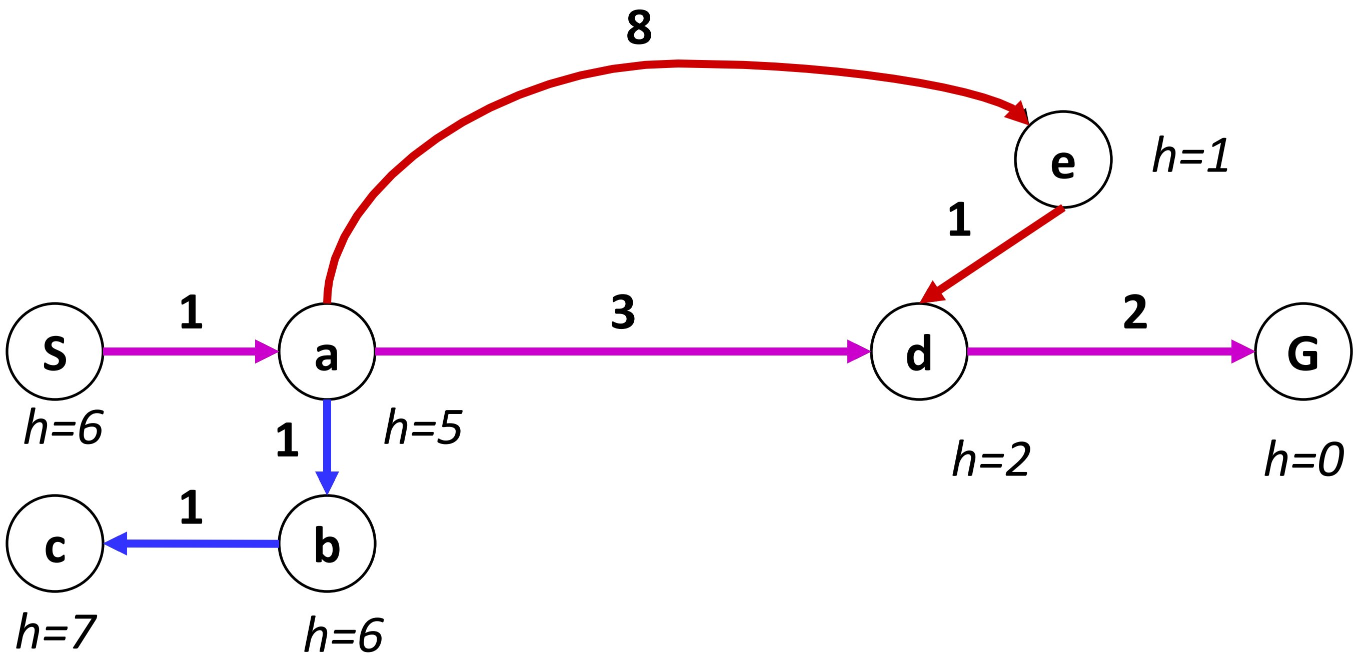

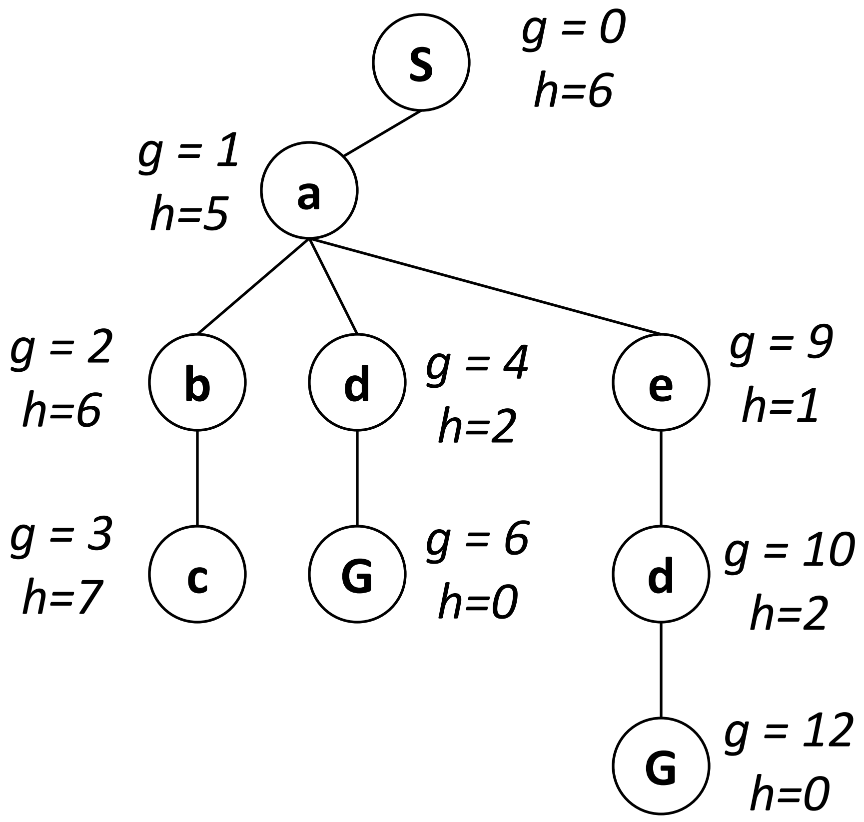

- A* Search orders by the sum: $f(n) = g(n) + h(n)$

Example: Teg Grenager

When should A* terminate?

- Should we stop when we <u>enqueue</u> a goal?

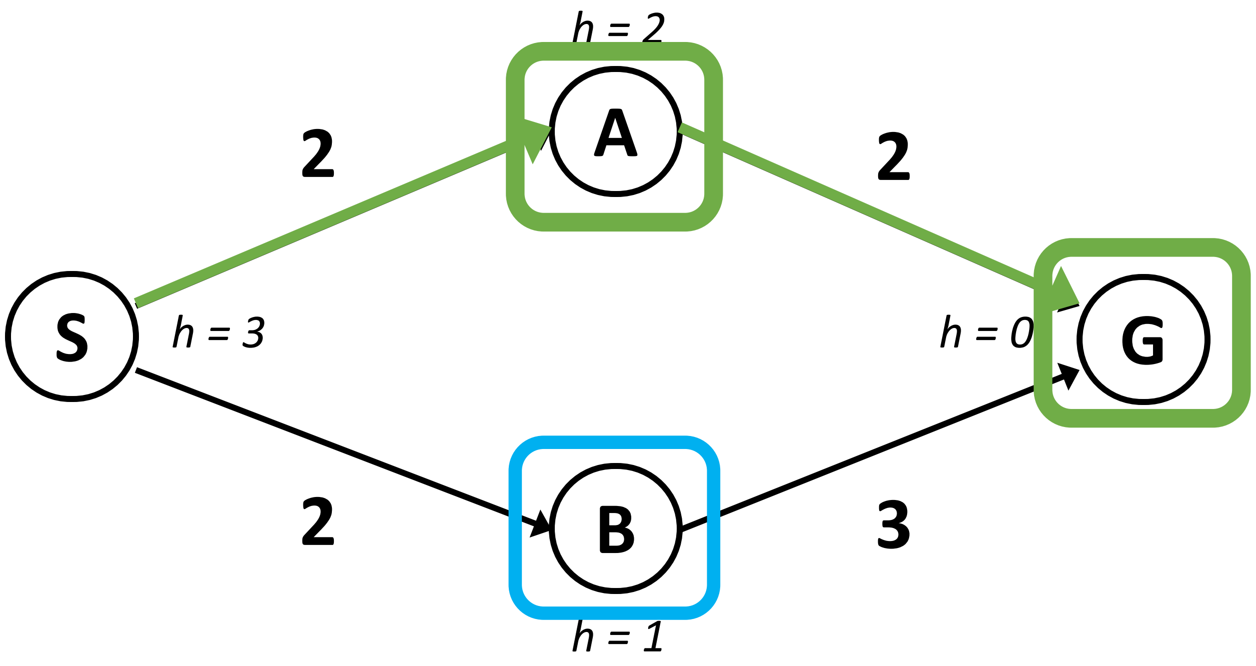

- No: only stop when we <u>dequeue</u> a goal

Is A* Optimal?

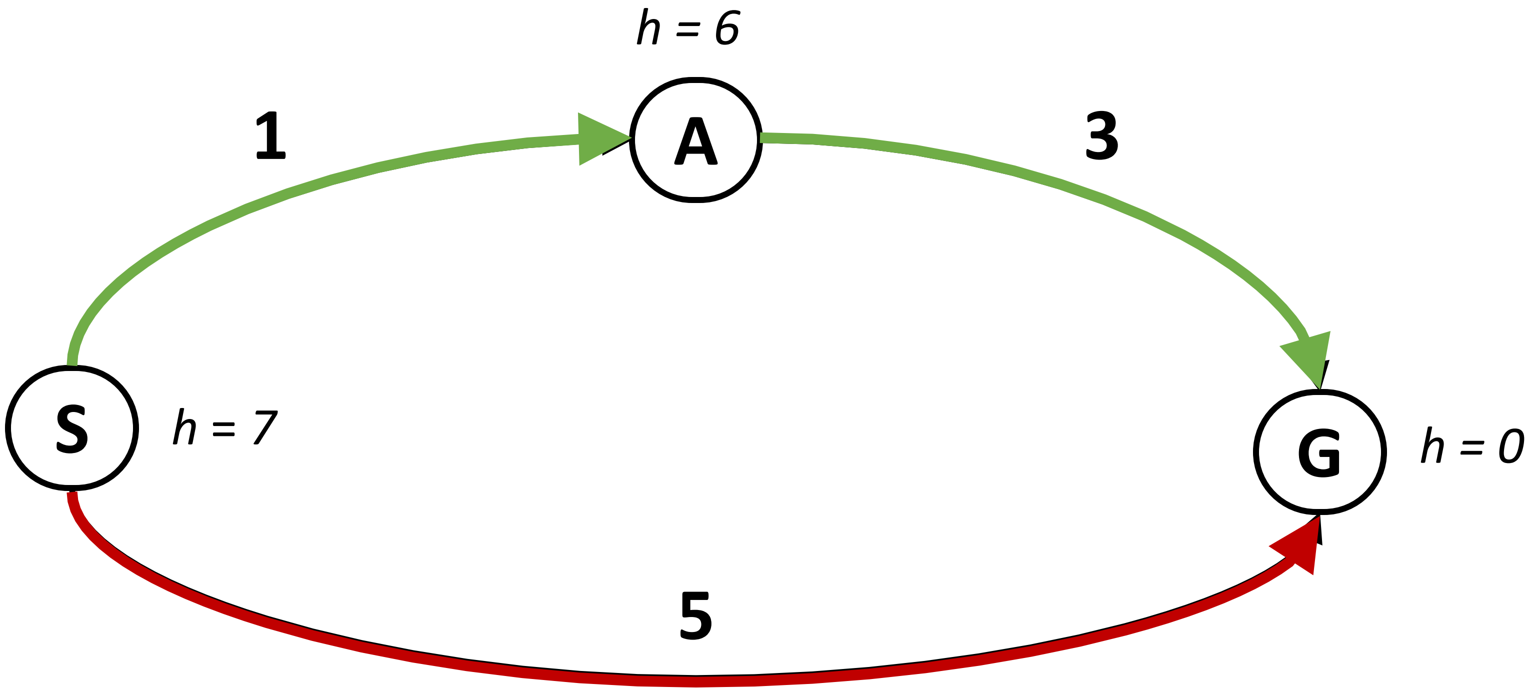

- What went wrong?
- Actual cost of bad path < estimated cost of optimal path
- We need estimates to be less than actual costs!

Admissible Heuristics

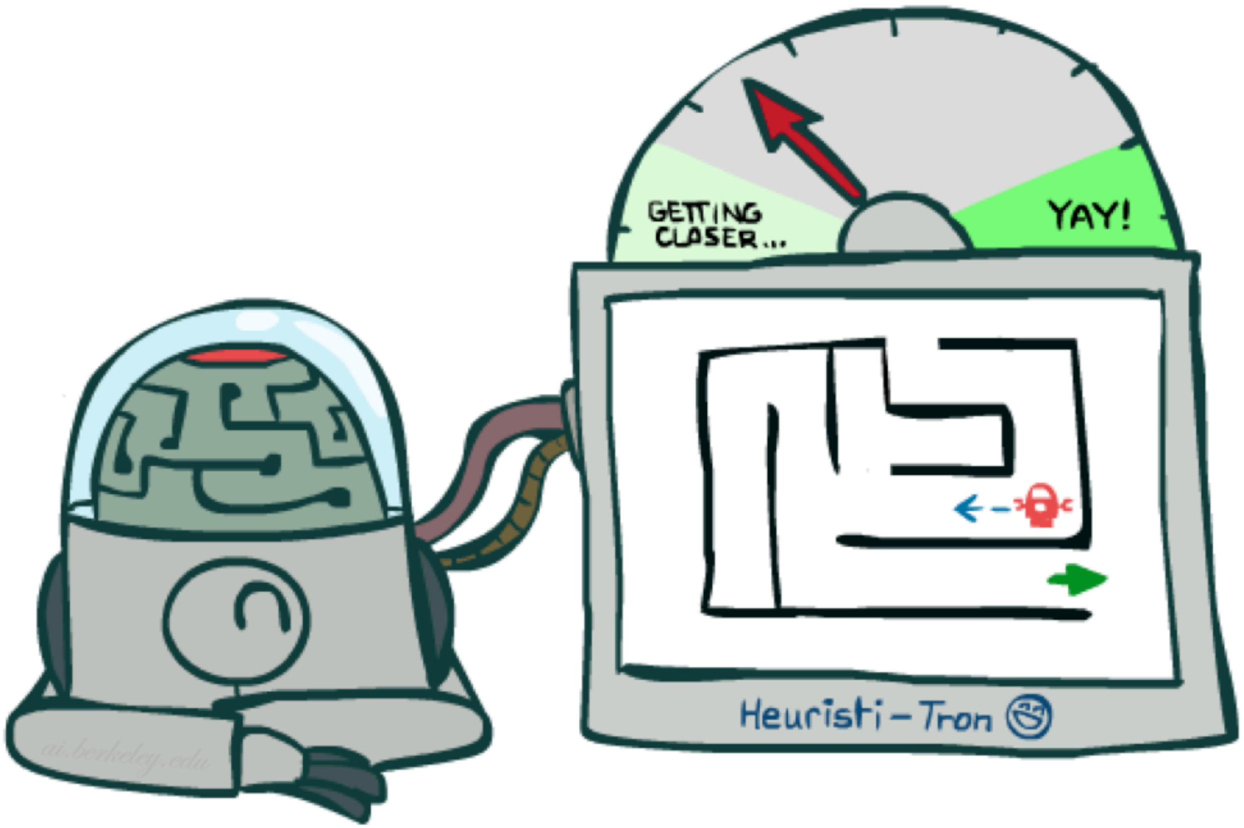

Idea: Admissibility

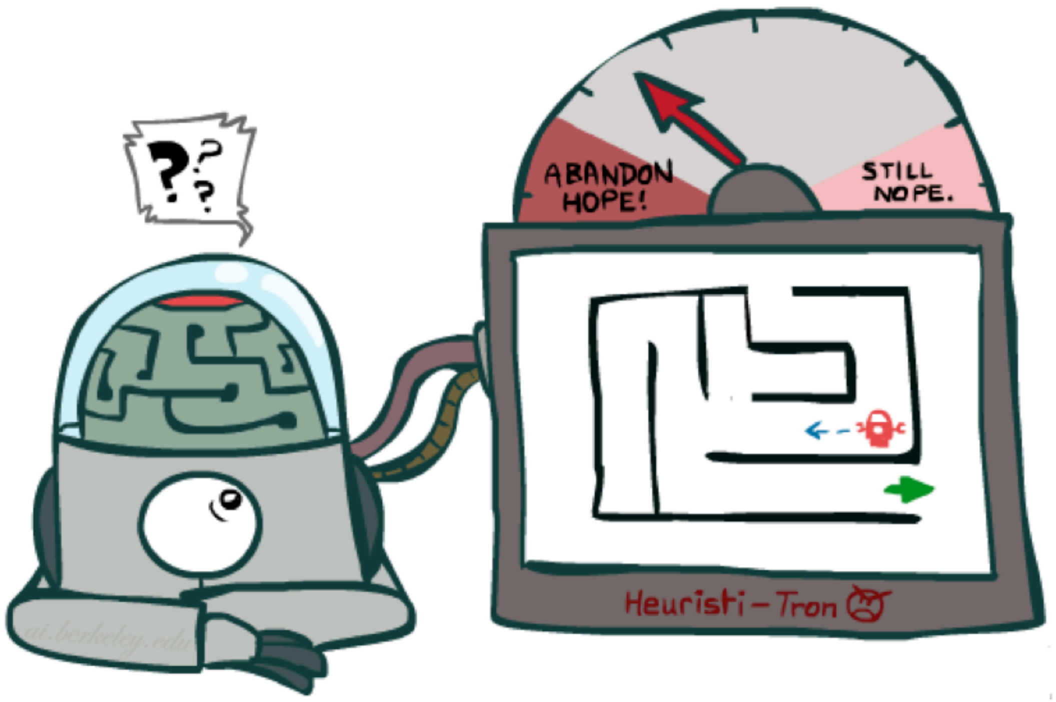

Inadmissible (pessimistic) heuristics break optimality by trapping good plans on the fringe

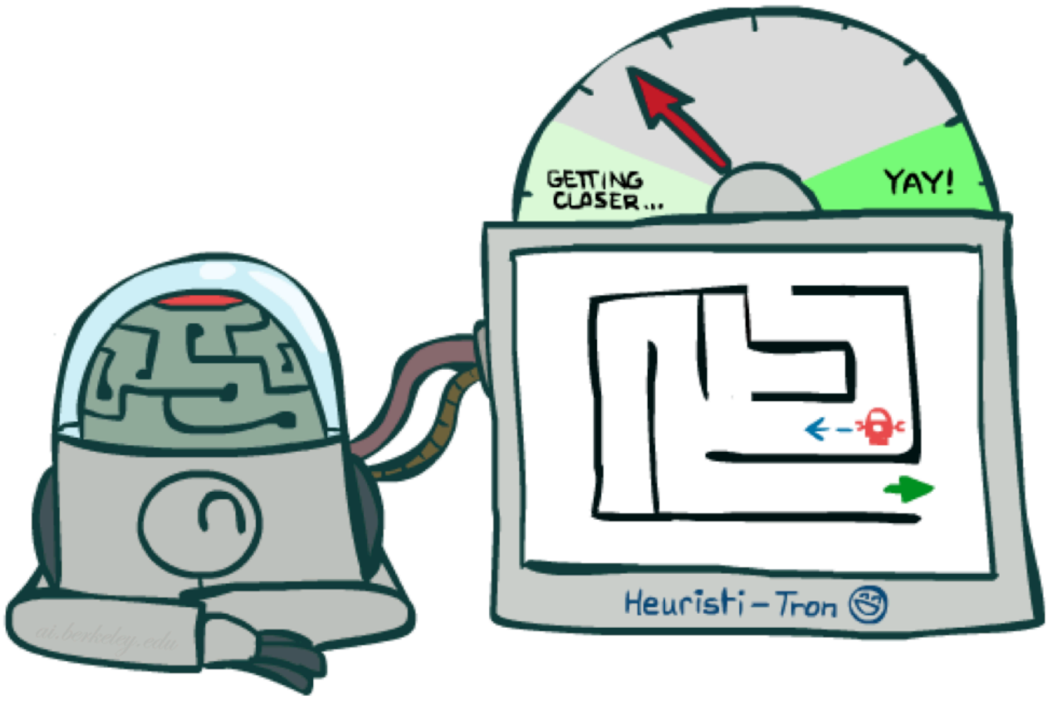

Admissible (optimistic) heuristics slow down bad plans but never outweigh true costs

Admissible Heuristics

- A heuristic h is *admissible* (optimistic) if:

$$0 \leq h(n) \leq h^*(n)$$

where $h^*(n)$ is the true cost to a nearest goal

- Examples:

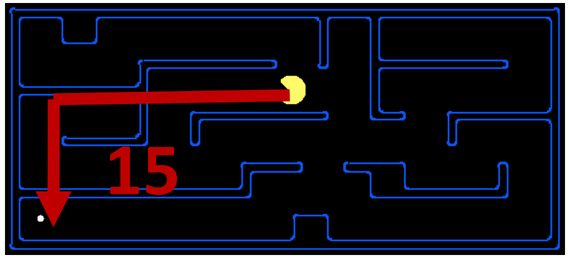

4

- Coming up with admissible heuristics is most of what's involved in using A* in practice.

Optimality of A* Tree Search

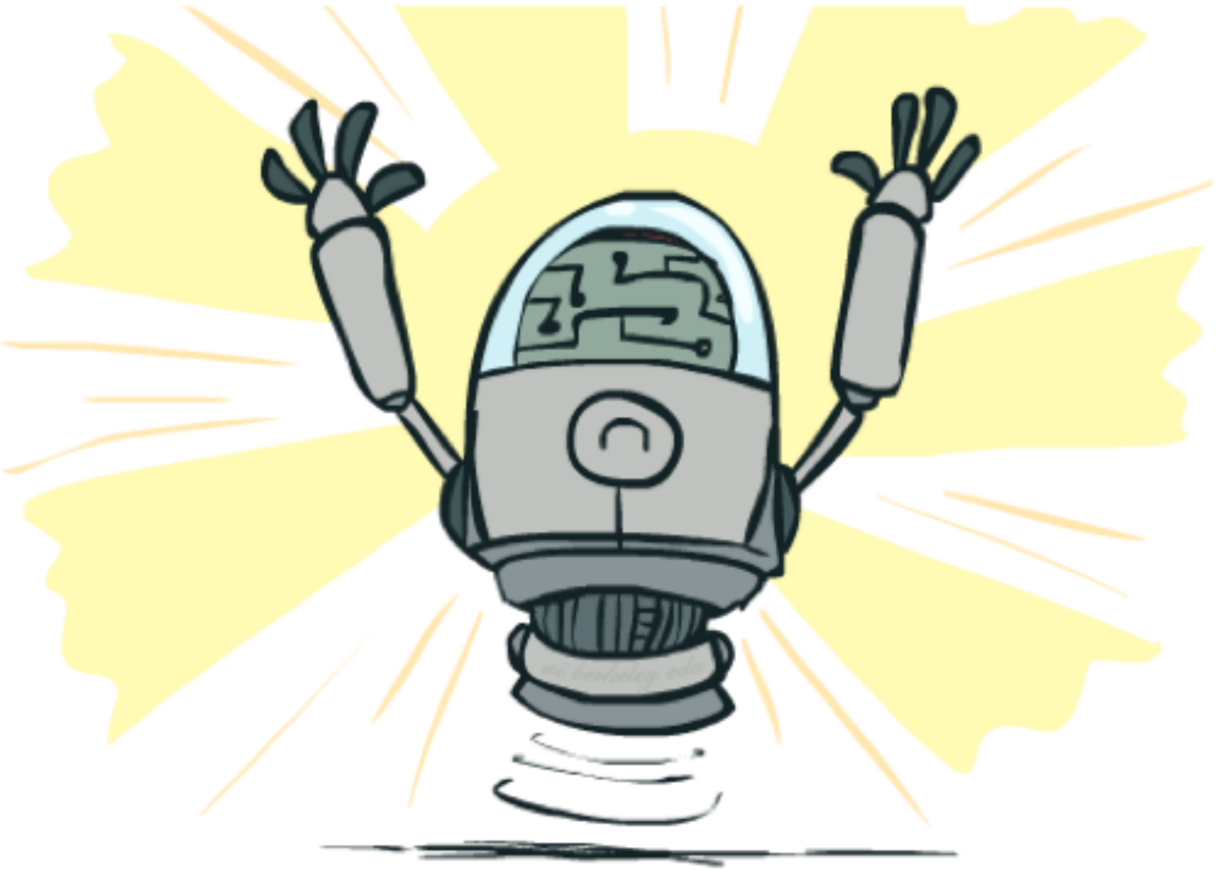

Optimality of A* Tree Search

Assume:

- A is an optimal goal node
- B is a suboptimal goal node
- h is admissible

Claim:

- A will exit the fringe before B

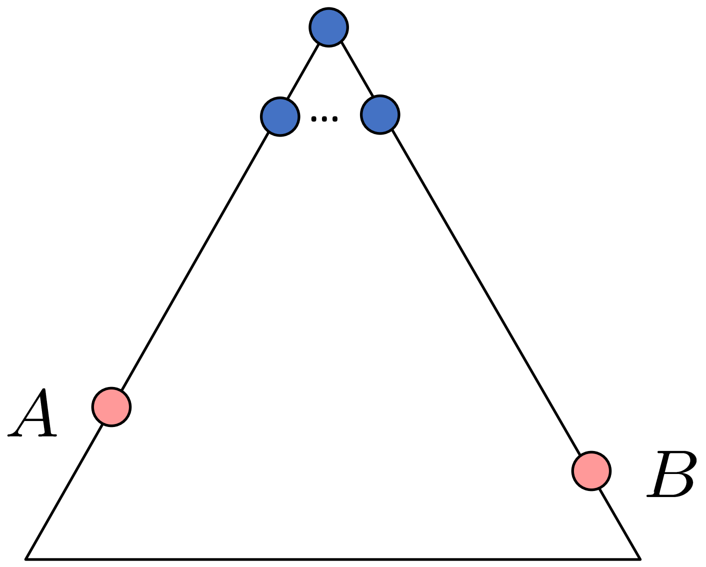

Optimality of A* Tree Search

Proof:

- Imagine B is on the fringe
- Some ancestor *n* of A is on the fringe, too (maybe A!)
- Claim: *n* will be expanded before B
    1. f(n) is less or equal to f(A)

g(n) = backward (path) cost
h(n) = forward (heuristic) cost

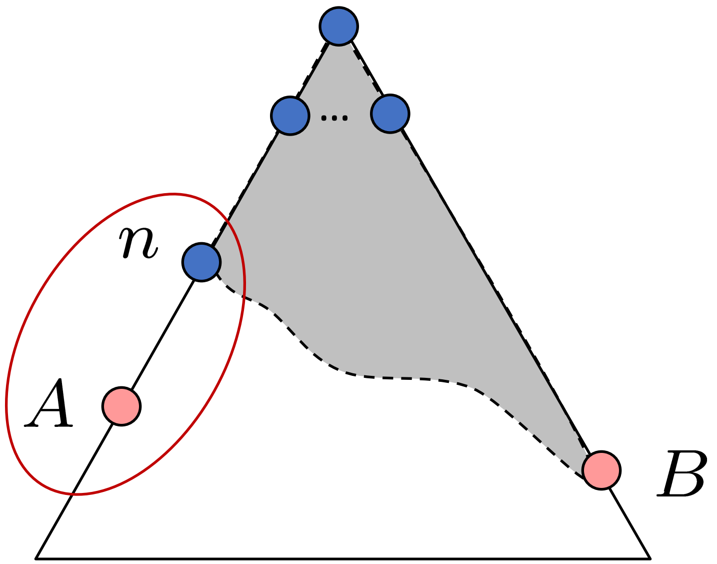

$f(n) = g(n) + h(n)$	Definition of f-cost
$f(n) \leq g(A)$	Admissibility of h
$g(A) = f(A)$	h = 0 at a goal

Optimality of A* Tree Search

Proof:

- Imagine B is on the fringe
- Some ancestor *n* of A is on the fringe, too (maybe A!)
- Claim: *n* will be expanded before B
    1. f(n) is less or equal to f(A)
    2. f(A) is less than f(B)

g(n) = backward (path) cost
h(n) = forward (heuristic) cost

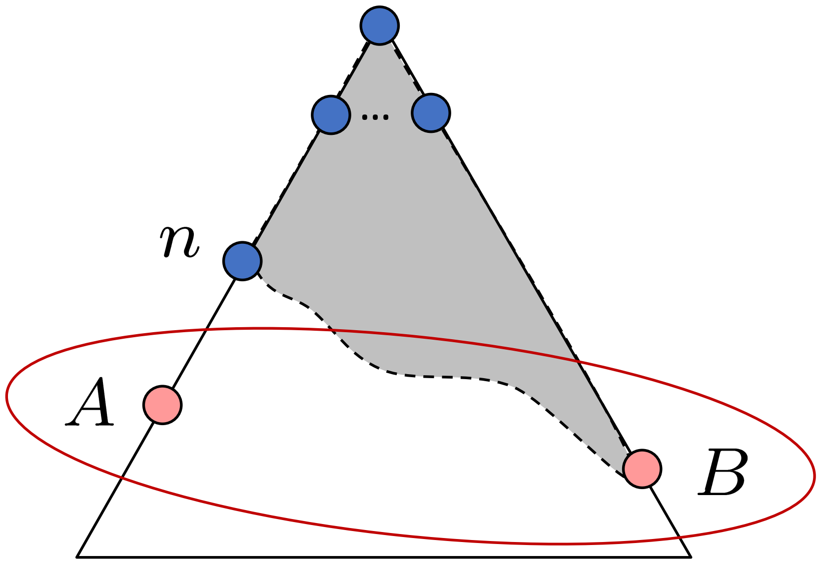

$g(A) < g(B)$	B is suboptimal
$f(A) < f(B)$	h = 0 at a goal

Optimality of A* Tree Search

Proof:

- Imagine B is on the fringe
- Some ancestor *n* of A is on the fringe, too (maybe A!)
- Claim: *n* will be expanded before B
    1. f(n) is less or equal to f(A)
    2. f(A) is less than f(B)
    3. *n* expands before B
- All ancestors of A expand before B
- A expands before B
- A* search is optimal

g(n) = backward (path) cost

h(n) = forward (heuristic) cost

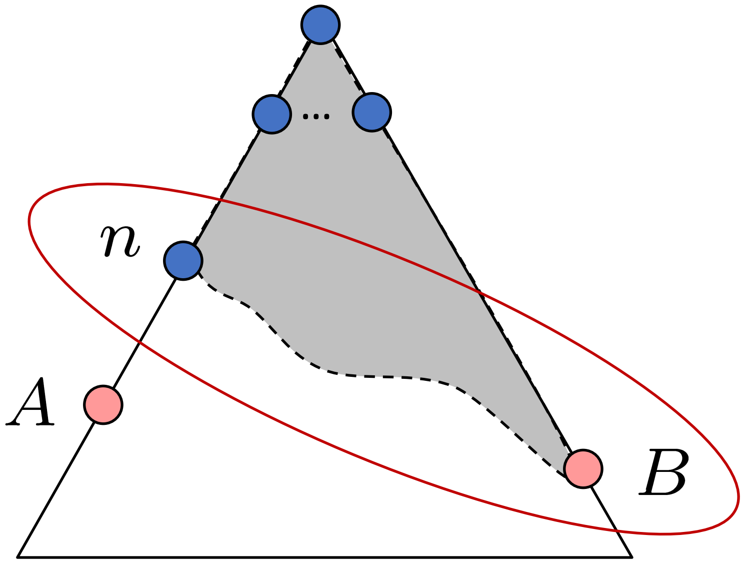

$$f(n) \leq f(A) < f(B)$$

Corollary: Optimality of UCS

A* search is optimal, given an admissible heuristic *h*

$$0 \leq h(n) \leq h^*(n)$$

UCS is equivalent to A* with null heuristic $h(n) = 0$

- ✓ Definitely admissible!

Therefore, UCS is also optimal.

Next Class

A* applications