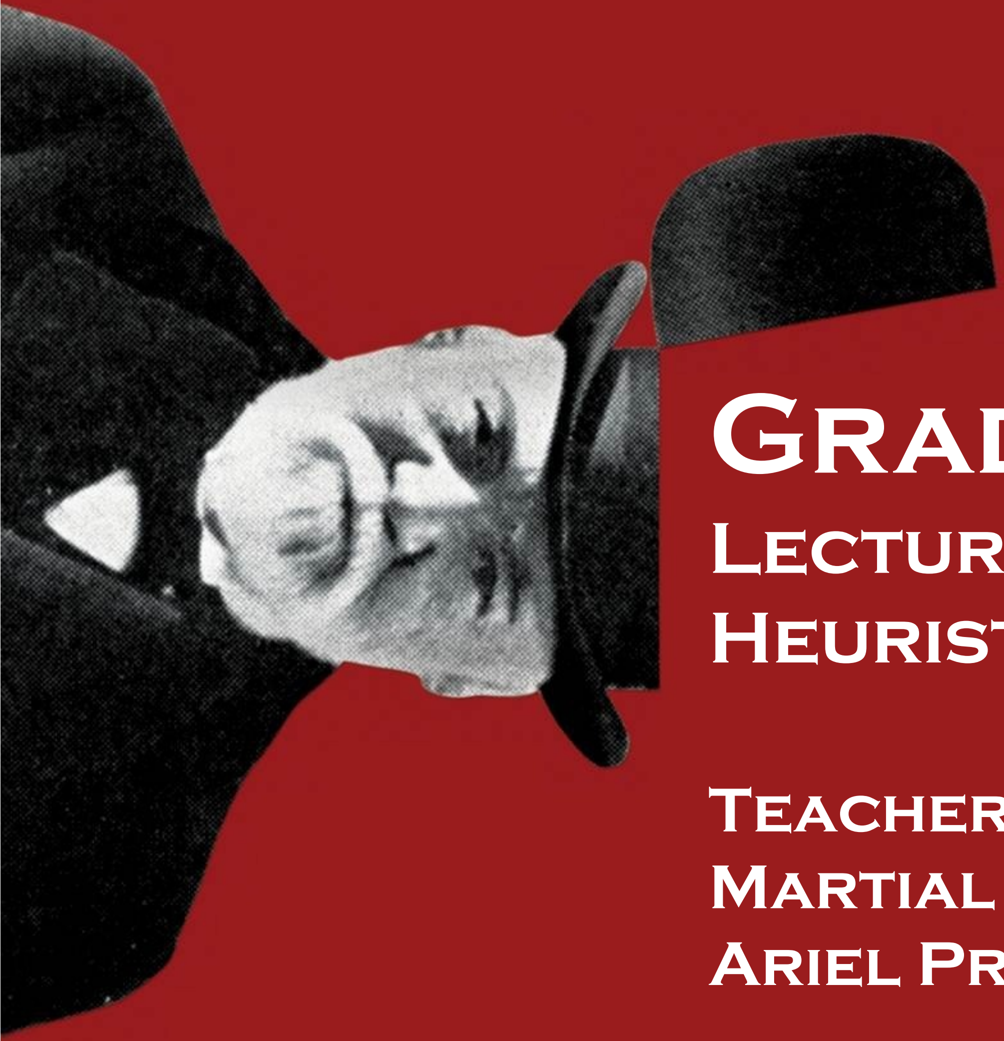

Graduate AI

Lecture 9: Heuristic Search

Teachers:
Martial Hebert
Ariel Procaccia (this time)

Heuristics

- On Monday we saw that heuristics matter!
- Heuristics are usually taken to mean "rules of thumb"
- Practical techniques that work well despite lack of theoretical guarantees
- In this lecture: a bit more formal than that

Best-first search

- Find a path from initial state to goal state
- Relies on an evaluation function
- Nodes with best evaluation value are explored first
- Different evaluation functions induce different algorithms

Greedy search

- Best-first search with evaluation function h(n)
- h(n) = estimated cost from node n to a goal state

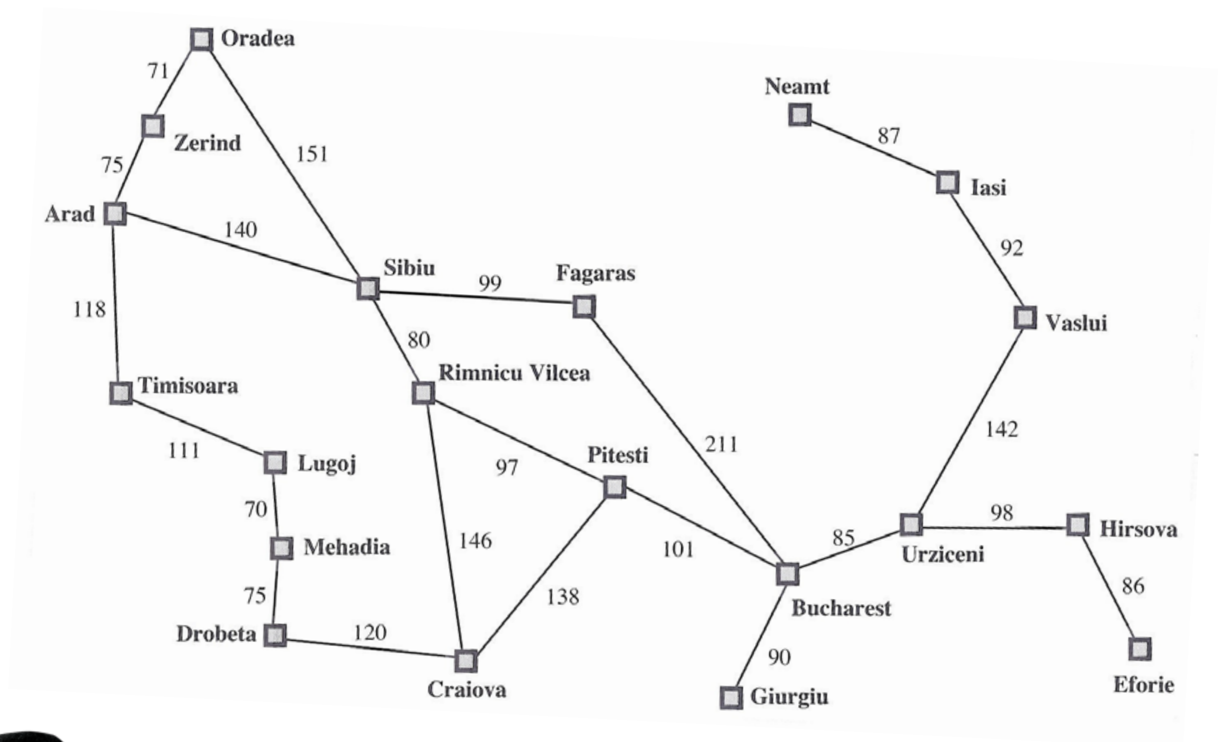

Greedy search: example

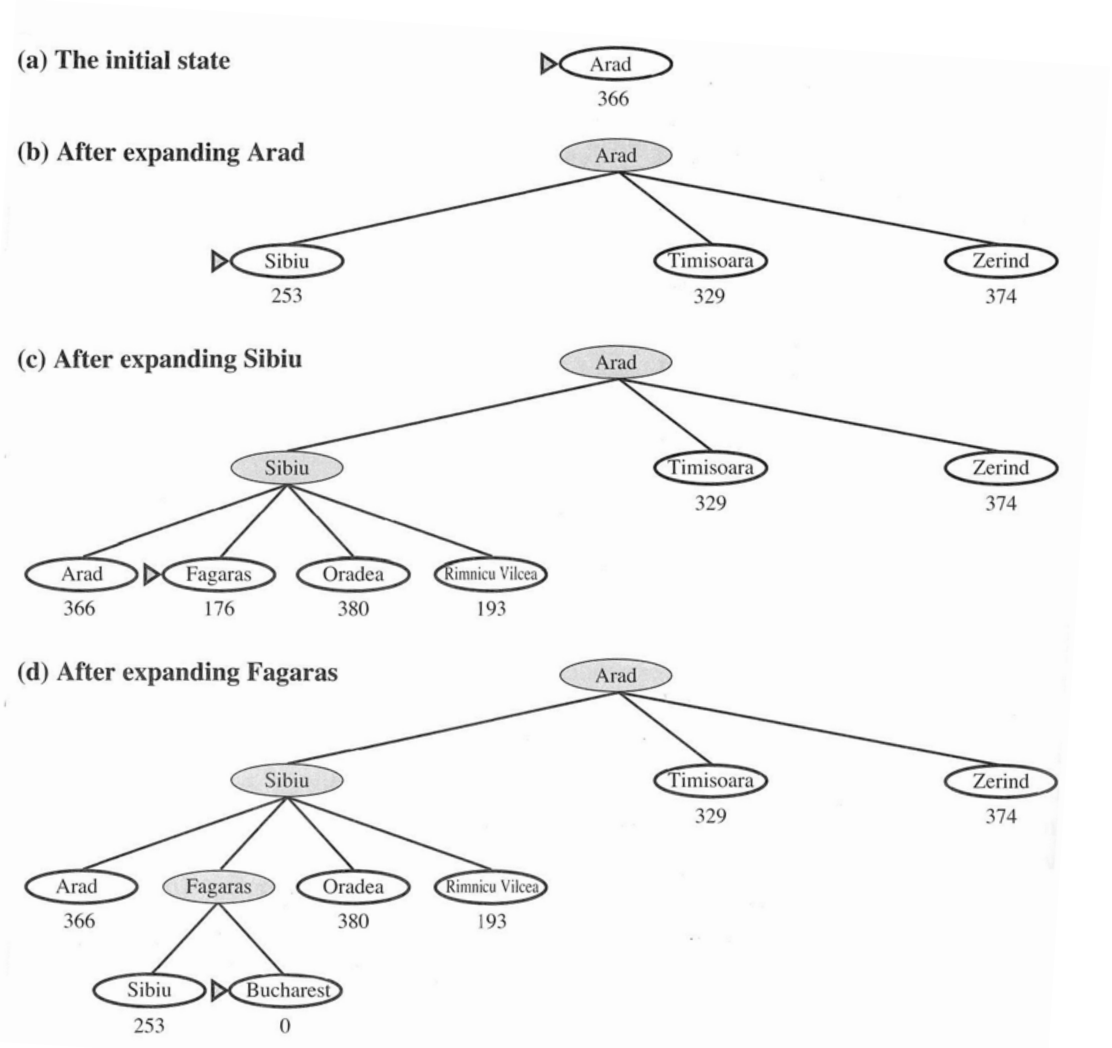

A* Search

- Best-first search with f(n) = g(n) + h(n)
- g(n) = work done so far, h(n) = estimate of remaining work

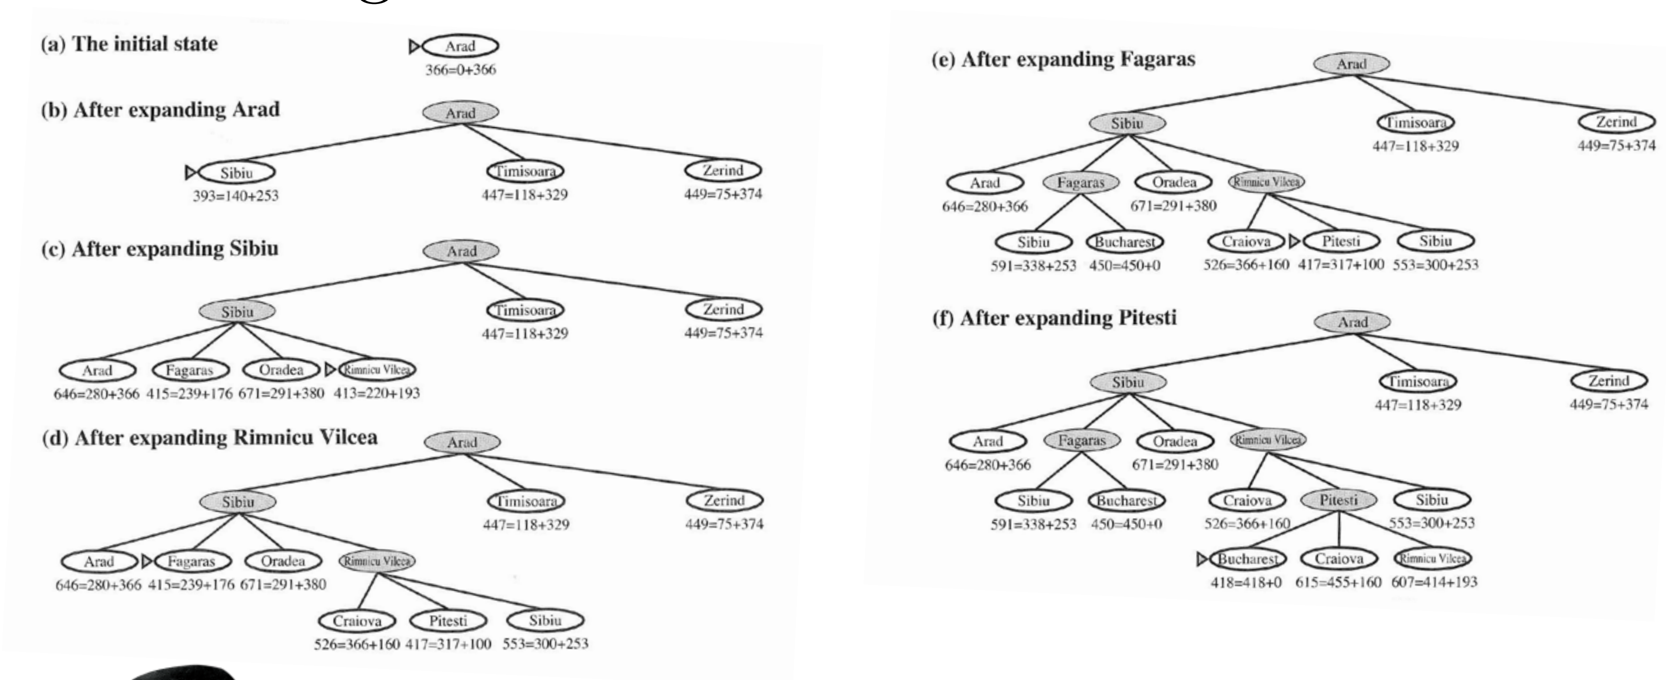

GOOD HEURISTICS

- k(n,n') = cost of cheapest path between n and n'
- h is **consistent** if for every n,n',
  $h(n) \leq k(n,n') + h(n')$
- Line distance heuristic is consistent by the triangle inequality

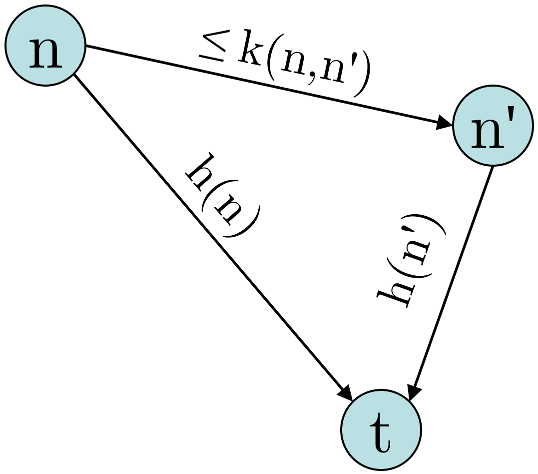

Optimality of A*

- **Theorem:** If h is consistent, A* returns the min cost solution
- Proof:
  - Assume $h(n) \leq k(n,n') + h(n')$
  - Values of $f(n)$ on a path are nondecreasing: if n' is the successor of n then $f(n') = g(n') + h(n') \geq g(n) + h(n) = f(n)$
  - When A* selects n for expansion, the optimal path to n has been found: otherwise there is a frontier node n' on optimal path to n that should be expanded first
  - $\Rightarrow$ Nodes expanded in nondecreasing $f(n)$
  - First goal state that is expanded must be optimal QED

More on consistency

- With a consistent heuristic A* f-costs are nondecreasing
- We can draw contours in the state space

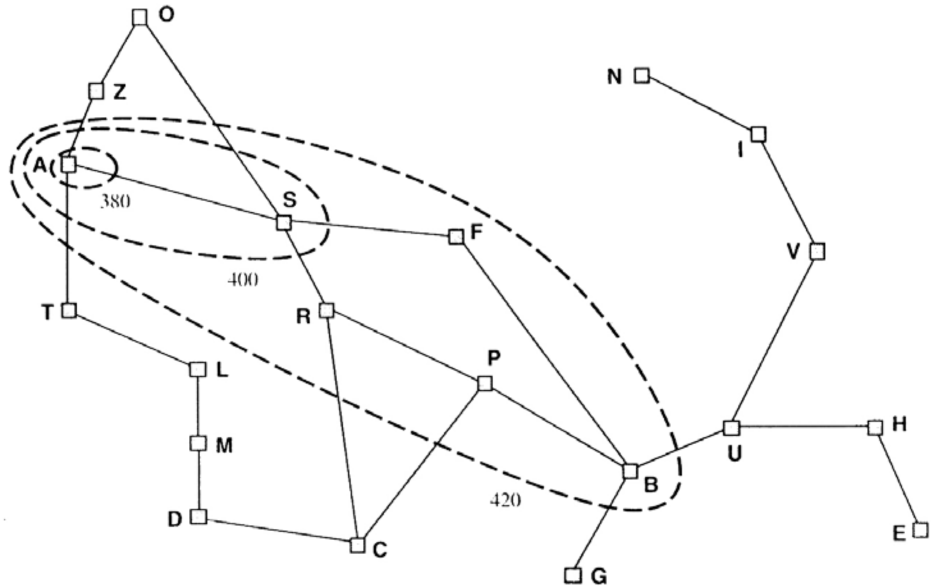

ADMISSIBILITY

- $h^*(n)$ = cost of cheapest path from n to a goal
- h is **admissible** if for all nodes n, $h(n) \leq h^*(n)$
- Consistency implies admissibility
  - For goal t, $h(n) \leq k(n,t) + h(t) = k(n,t) = h^*(n)$
- A* with admissible h is optimal under additional assumptions

8-puzzle heruistics

- h_1: #tiles in wrong position
- h_2: sum of Manhattan distances of tiles from goal
- Both are admissible
- h_2 dominates h_1, i.e., $h_1(n) \leq h_2(n)$ for all n

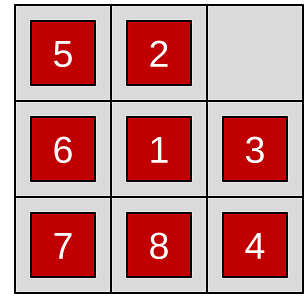

Example state

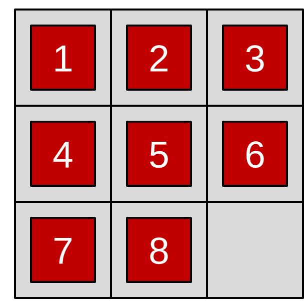

Goal state

THE IMPORTANCE OF A GOOD HEURISTIC

- The following table gives the search cost of A* with the two heuristics, averaged over random puzzles, for various solution lengths

Length	A*(h_1)	A*(h_2)
16	1301	211
18	3056	363
20	7276	676
22	18094	1219
24	39135	1641

Optimality over other algs

- We prove the following statements on the board
- They also appear in:
  R. Dechter and J. Pearl. Best-first search and the optimality of A*. Journal of the ACM 32:506-536, 1985 (link on course website)
- Any alg that is admissible given consistent heuristics will expand all nodes surely expanded by A* [Dechter and Pearl, Thm 8 on page 522]
- This is not true if the heurisitic is merely admissible [Dechter an Pearl, pages 524-525]