

CSEP 573:

Artificial Intelligence

Hanna Hajishirzi

Search

(Un-informed, Informed Search)



slides adapted from
Dan Klein, Pieter Abbeel ai.berkeley.edu
And Dan Weld, Luke Zettlemoyer

Announcements

- PS1
 - Due Jan 19th
- Office hours:
 - Check the website
- Classes:
 - Online in January
- HW1 will be released soon.

Recap: Search

- Search problem:


- States (configurations of the world)
- Actions and costs
- Successor function (world dynamics)
- Start state and goal test

- Search tree:

- Nodes: represent plans for reaching states

- Search algorithm:

- Systematically builds a search tree
- Chooses an ordering of the fringe (unexplored nodes)
- Optimal: finds least-cost plans


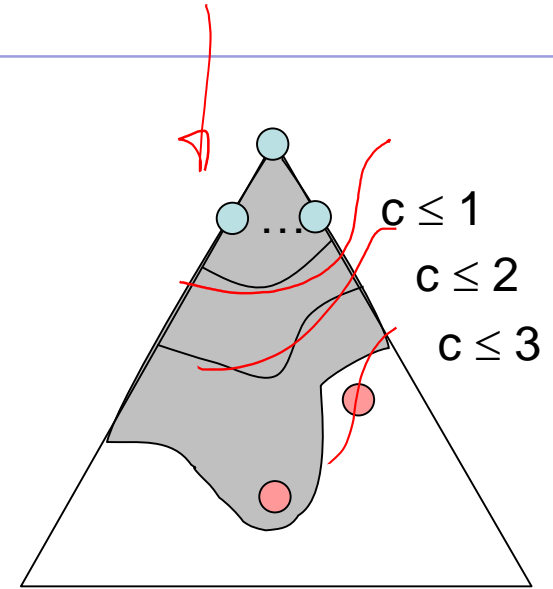


General Tree Search

```
function TREE-SEARCH(problem, strategy) returns a solution, or failure
  initialize the search tree using the initial state of problem
  loop do
    if there are no candidates for expansion then return failure
    choose a leaf node for expansion according to strategy
    if the node contains a goal state then return the corresponding solution
    else expand the node and add the resulting nodes to the search tree
  end
```

Uniform Cost Issues


- Remember: UCS explores increasing cost contours
- The good: UCS is complete and optimal!
- The bad:
 - Explores options in every “direction”
 - No information about goal location
- We'll fix that soon!



Up next: Informed Search

○ Uninformed Search ↓

- DFS
- BFS
- UCS



■ Informed Search

- Heuristics
- Greedy Search
- A* Search
- Graph Search






Search Heuristics

- A heuristic is:

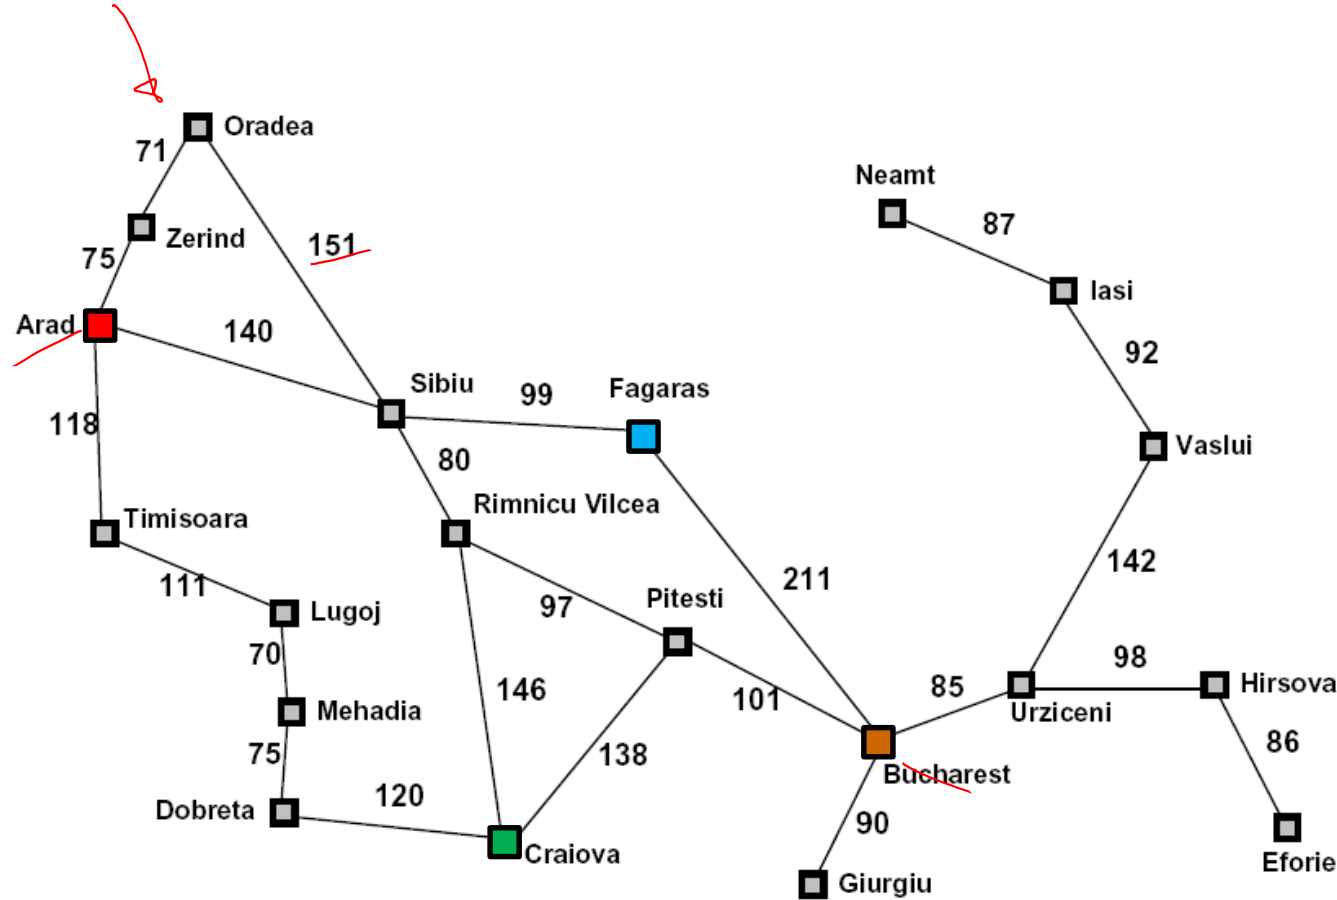
- A function that *estimates* how close a state is to a goal
- Designed for a particular search problem

→ ~~Pathing?~~

- Examples: Manhattan distance, Euclidean distance for pathing



Example: Heuristic Function



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

$h(x)$


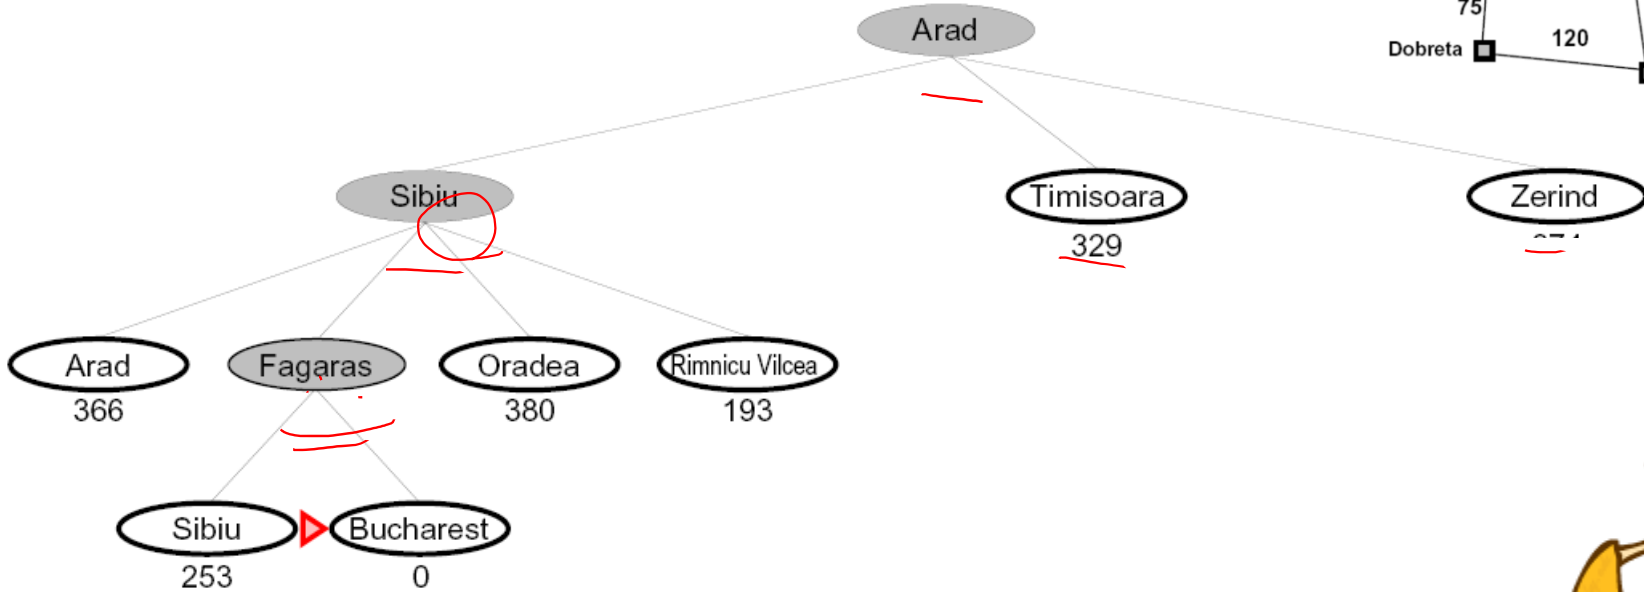
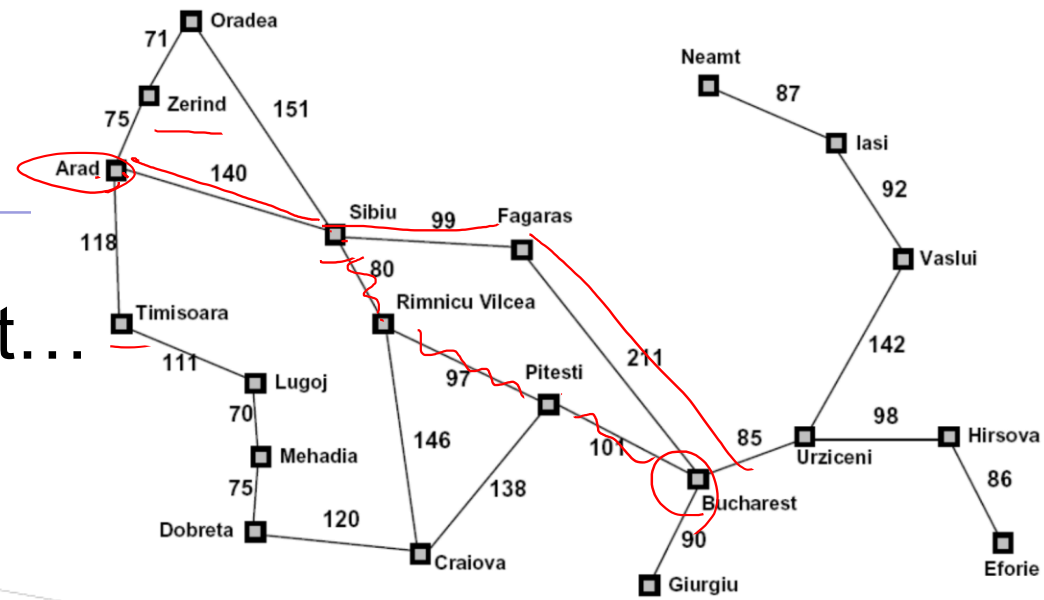
Greedy Search



$h(x)$

Greedy Search


- Expand the node that seems closest...





- Is it optimal?

- No. Resulting path to Bucharest is not the shortest!


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 (wrong) goal
- Worst-case: like a badly-guided DFS



Video of Demo Contours Greedy (Empty)



A* Search



A* Search

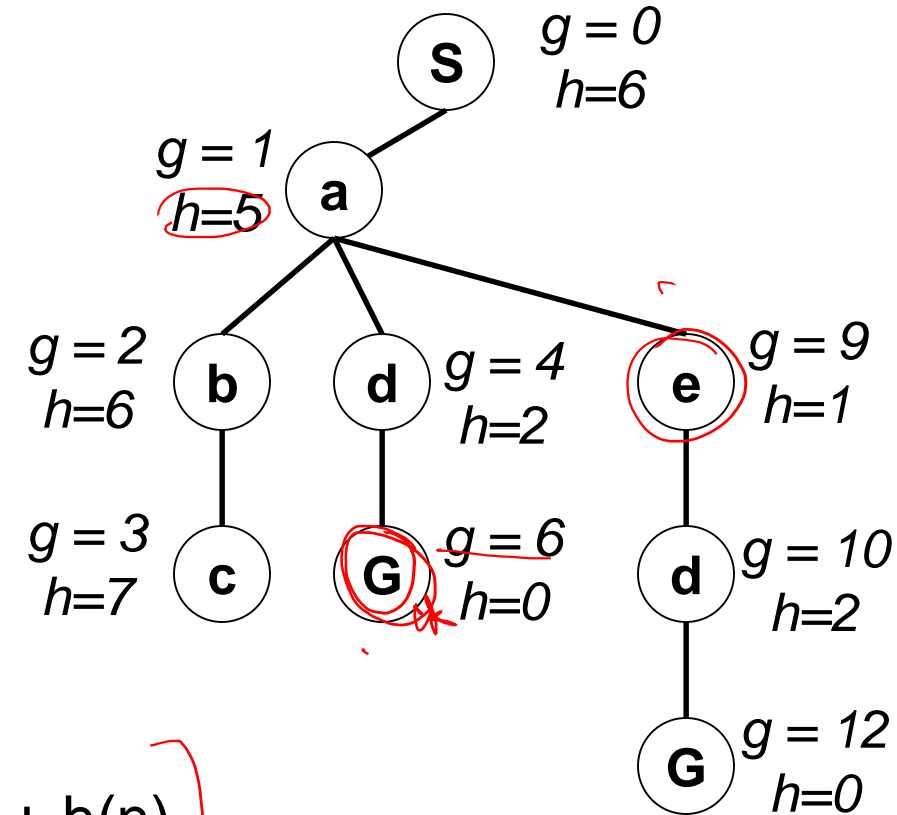
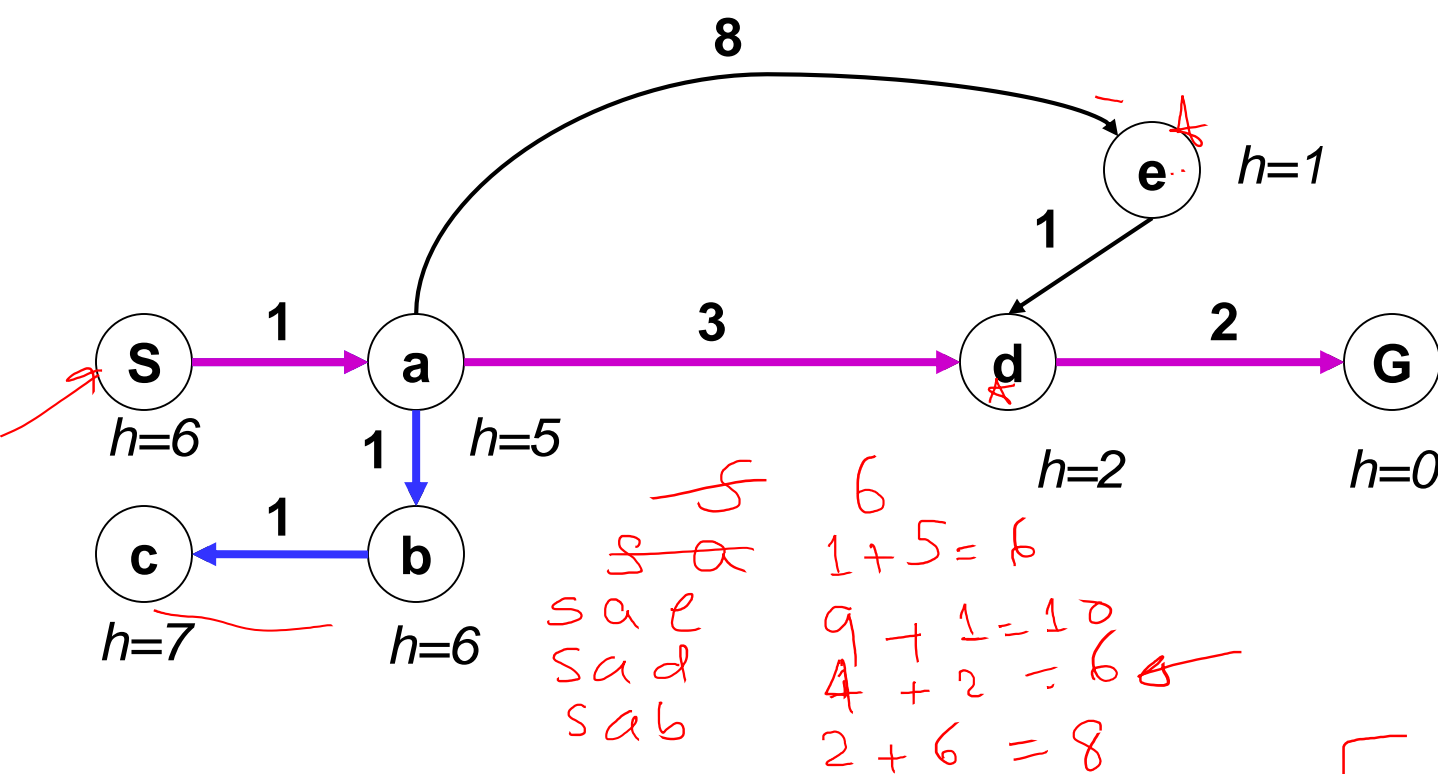
UCS

→ A

UCS, greedy
greedy, greedy

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

When should A* terminate?

- Should we stop when we enqueue a goal?




	g	h	+
S	0	3	3
S->A	2	2	4
S->B	2	1	3
S->B->G	5	0	5
S->A->G	4	0	4

- No: only stop when we dequeue a goal

if admissible \rightarrow optimal

Is A* Optimal?


SAG




	g	h	+
S	0	7	7
S \rightarrow A	1	6	7
S \rightarrow G	5	0	5

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

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:



0.0


$$f = g + \frac{b}{s}$$

↓
UCS


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

Properties of A^*

Uniform-Cost




A^*




UCS vs A* Contours


- Uniform-cost expands equally in all “directions”



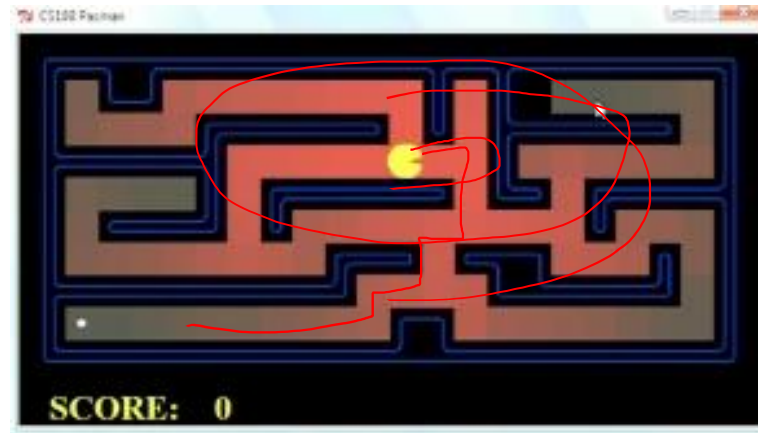
- A* expands mainly toward the goal, but does hedge its bets to ensure optimality



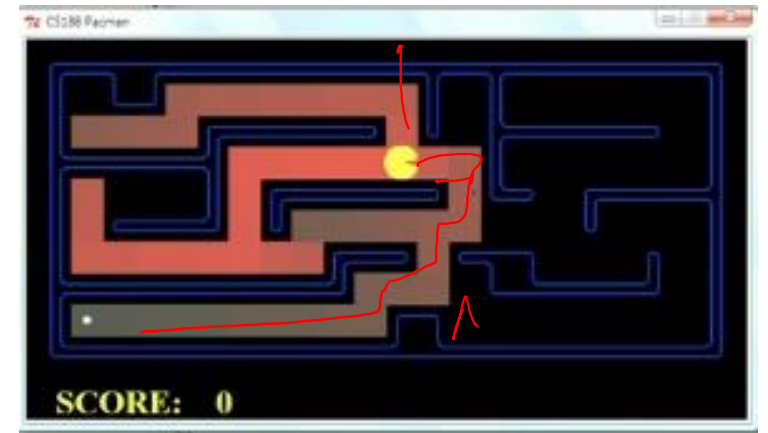
Comparison



Greedy




Uniform Cost




A*


Video of Demo Contours (Empty) -- UCS



Video of Demo Contours (Empty) -- Greedy



Video of Demo Contours (Empty) – A^*




A*: Summary



$UCS + greedy$
 $g(n) + h(n)$

A*: Summary

- A* uses both backward costs and (estimates of) forward costs
- A* is optimal with admissible (optimistic) heuristics
- Heuristic design is key: often use relaxed problems



Video of Demo Empty Water Shallow/Deep – Guess Algorithm

DFS

BFS

UCS

Greedy

A*

The screenshot shows the Pydev Eclipse IDE interface. A search menu is open, listing various search algorithms. The console at the bottom displays the output of a search operation.


Search Menu:

- 1 search -- plan tiny astar
- 2 search -- plan tiny ucs
- 3 search demo empty
- 4 search -- contours greedy vs ucs (greedy)
- 5 search -- contours greedy vs ucs (ucs)
- 6 search -- contours greedy vs ucs (astar)
- 7 search -- greedy bad
- 8 search -- greedy good
- 9 search demo maze
- search demo no costs
- Run As
- Run Configurations...
- Organize Favorites...

Console Output:


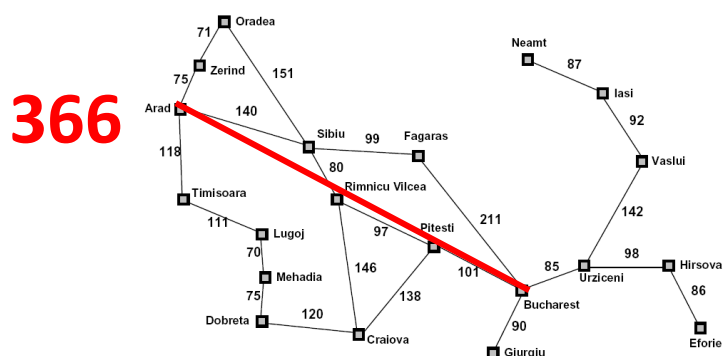
```
<terminated> 1.5  
Total cost: 27  
Number of nodes expanded: 182  
Number of unique nodes expanded: 182  
Pacman emerges victorious! Score: 573  
{'numKills': [0], 'results': ['Win'], 'numMoves': [27], 'scores': [573]}
```

Creating Heuristics



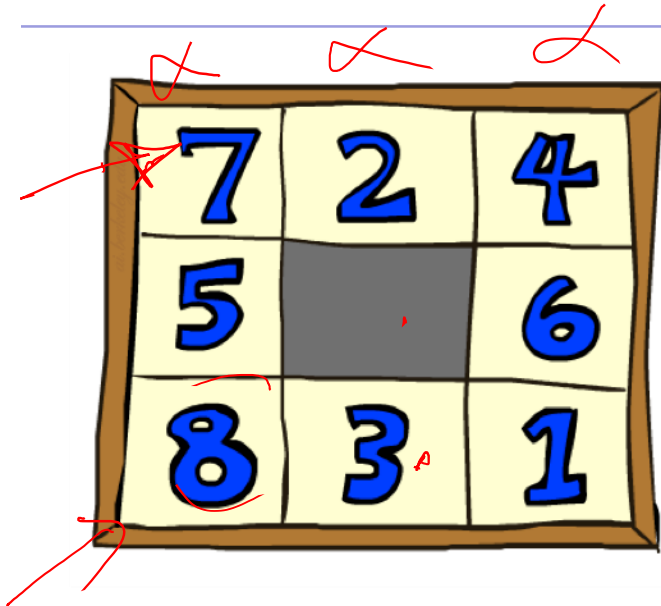
Creating Admissible Heuristics

- Most of the work in solving hard search problems optimally is in coming up with admissible heuristics
- Often, admissible heuristics are solutions to *relaxed problems*, where new actions are available




- Inadmissible heuristics are often useful too


Example: 8 Puzzle



Start State



Actions




Goal State

- What are the states?
- How many states?
- What are the actions?
- How many successors from the start state?
- What should the costs be?

Admissible heuristics?


8 Puzzle I

- Heuristic: Number of tiles misplaced
- Why is it admissible?
- $h(\text{start}) = 8$
- This is a *relaxed-problem* heuristic



Start State

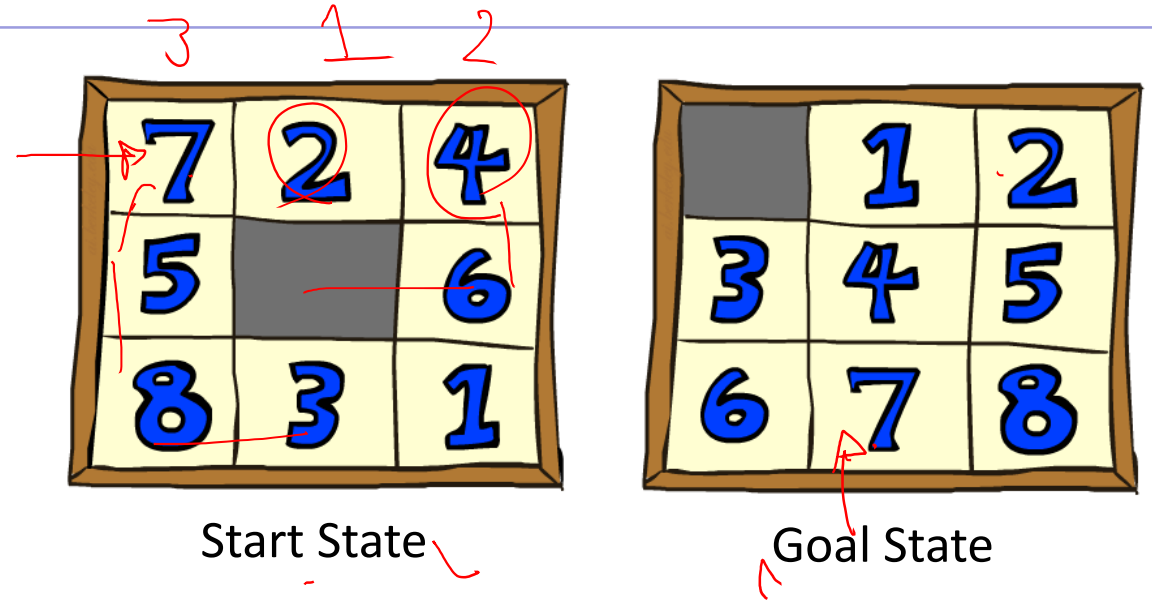
Goal State



	Average nodes expanded when the optimal path has...		
	...4 steps	...8 steps	...12 steps
UCS	112	6,300	3.6×10^6
TILES	13	39	227

8 Puzzle II



- What if we had an easier 8-puzzle where any tile could slide any direction at any time, ignoring other tiles?
- Total *Manhattan* distance
- Why is it admissible?
- $h(\text{start}) = \underline{3 + 1 + 2 + \dots} = 18$



	Average nodes expanded when the optimal path has...		
	<u>...4 steps</u>	...8 steps	...12 steps
TILES	<u>13</u>	39	227
MANHATTAN	<u>12</u>	25	73


8 Puzzle III

- How about using the *actual cost* as a heuristic?
 - Would it be admissible?
 - Would we save on nodes expanded?
 - What's wrong with it?



- With A^* : a trade-off between quality of estimate and work per node
 - As heuristics get closer to the true cost, you will expand fewer nodes but usually do more work per node to compute the heuristic itself

Example: Pancake Problem



Cost: Number of pancakes flipped

Example: Pancake Problem

BOUNDS FOR SORTING BY PREFIX REVERSAL

William H. GATES

Microsoft, Albuquerque, New Mexico

Christos H. PAPANIMITRIOU*†

Department of Electrical Engineering, University of California, Berkeley, CA 94720, U.S.A.


Received 18 January 1978

Revised 28 August 1978

For a permutation σ of the integers from 1 to n , let $f(\sigma)$ be the smallest number of prefix reversals that will transform σ to the identity permutation, and let $f(n)$ be the largest such $f(\sigma)$ for all σ in (the symmetric group) S_n . We show that $f(n) \leq (5n+5)/3$, and that $f(n) \geq 17n/16$ for n a multiple of 16. If, furthermore, each integer is required to participate in an even number of reversed prefixes, the corresponding function $g(n)$ is shown to obey $3n/2 - 1 \leq g(n) \leq 2n + 3$.

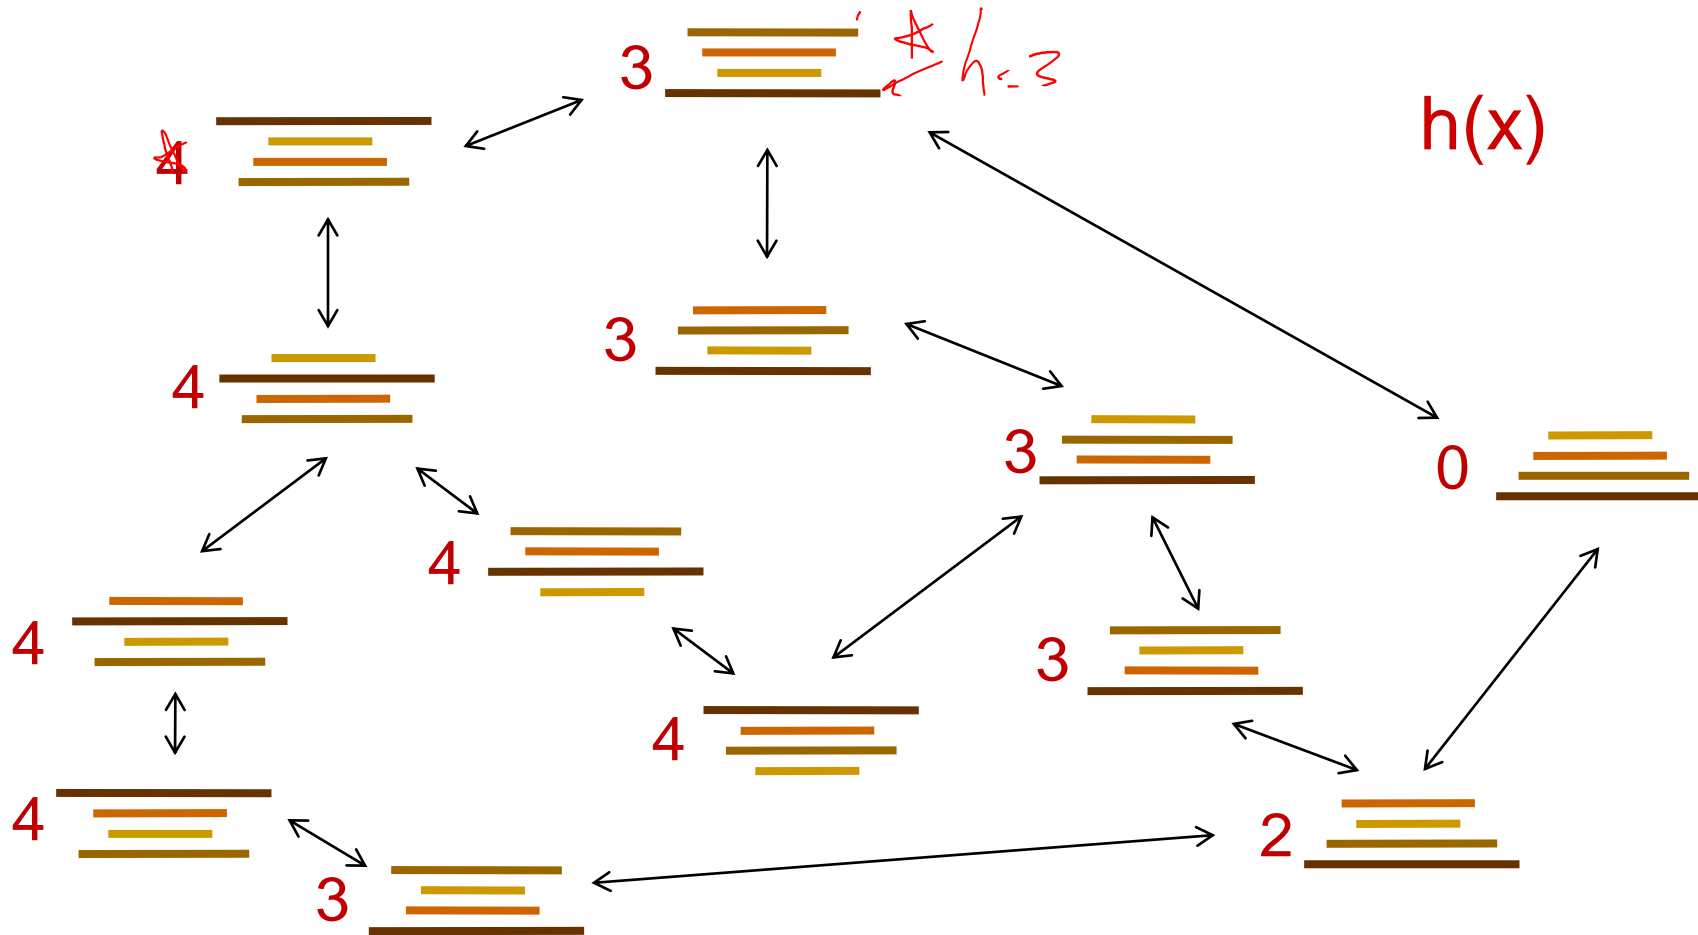
Example: Pancake Problem

State space graph with costs as weights



Example: Heuristic Function

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



Semi-Lattice of Heuristics

Trivial Heuristics, Dominance

- Dominance: $h_a \geq h_c$ if

$$\forall n : h_a(n) \geq h_c(n)$$


- Heuristics form a semi-lattice:

- Max of admissible heuristics is admissible


$$h(n) = \max(h_a(n), h_b(n))$$

- Trivial heuristics

- Bottom of lattice is the zero heuristic (what does this give us?)
- Top of lattice is the exact heuristic



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: Blocking

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 $f(A)$



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

$$\underline{f(n)} \leq \underline{g(A)}$$

$$\underline{g(A)} = \underline{f(A)}$$

Definition of f-cost


Admissibility of h

$h = 0$ at a goal

Optimality of A* Tree Search: Blocking

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 $f(B)$



$$g(A) < g(B)$$

$$f(A) < f(B)$$


B is suboptimal

$h = 0$ at a goal

Optimality of A* Tree Search: Blocking




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



$$\underline{f(n)} \leq \underline{f(A)} < \underline{f(B)}$$


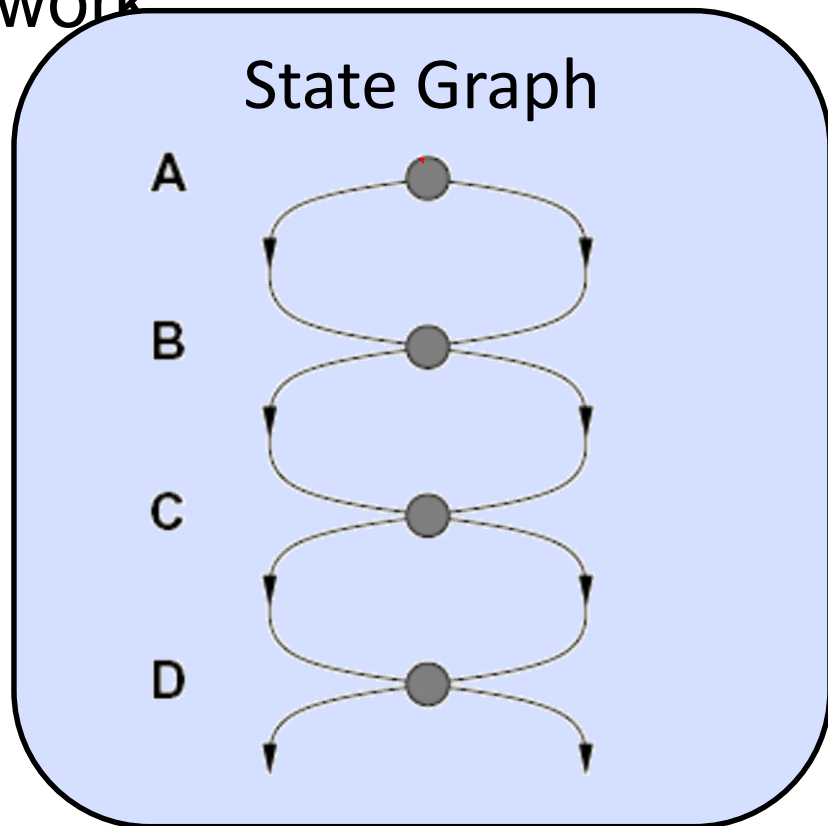
Graph Search



degrees


Tree Search: Extra Work!

- Failure to detect repeated states can cause exponentially more work



Graph Search

- In BFS, for example, we shouldn't bother expanding the circled nodes (why?)



Graph Search


- Idea: never **expand** a state twice
- How to implement:
 - Tree search + set of expanded states (“closed set”)
 - Expand the search tree node-by-node, but...
 - Before expanding a node, check to make sure its state has never been expanded before
 - If not new, skip it, if new add to closed set
- Important: **store the closed set as a set, not a list**
- Can graph search wreck completeness? Why/why not?
- How about optimality?

A* Graph Search Gone Wrong?


admissible:

consistent length:

State space graph




Search tree




Closed Set S B C A

Consistency of Heuristics




- Main idea: estimated heuristic costs \leq actual costs
 - Admissibility: heuristic cost \leq actual cost to goal
 $h(A) \leq$ actual cost from A to G
 - Consistency: heuristic "arc" cost \leq actual cost for each arc
 $h(A) - h(C) \leq \text{cost}(A \text{ to } C)$
- Consequences of consistency:
 - The f value along a path never decreases
 $h(A) \leq \text{cost}(A \text{ to } C) + h(C)$
 - A* graph search is optimal

A* Graph Search



- Sketch: consider what A* does with a consistent heuristic:
 - Fact 1: In tree search, A* expands nodes in increasing total f value (f -contours)
 - Fact 2: For every state s , nodes that reach s optimally are expanded before nodes that reach s suboptimally
- Result: A* graph search is optimal



Optimality of A* Search



- With a ⁿadmissible heuristic, Tree A* is optimal.
- With a consistent heuristic, Graph A* is optimal.
- With h=0, the same proof shows that UCS is optimal.

Pseudo-Code

```
function TREE-SEARCH(problem, fringe) return 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, STATE[node]) then return node
    for child-node in EXPAND(STATE[node], problem) do
      fringe ← INSERT(child-node, fringe)
    end
  end
```

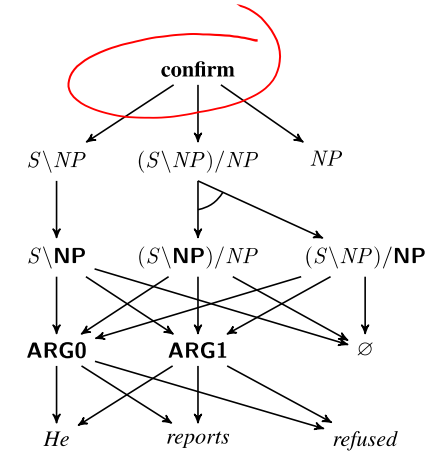
```
function GRAPH-SEARCH(problem, fringe) return a solution, or failure
  closed ← an empty set
  fringe ← INSERT(MAKE-NODE(INITIAL-STATE[problem]), fringe)
  loop do
    if fringe is empty then return failure
    node ← REMOVE-FRONT(fringe)
    if GOAL-TEST(problem, STATE[node]) then return node
    if STATE[node] is not in closed then
      add STATE[node] to closed
      for child-node in EXPAND(STATE[node], problem) do
        fringe ← INSERT(child-node, fringe)
      end
    end
  end
```

A* Applications

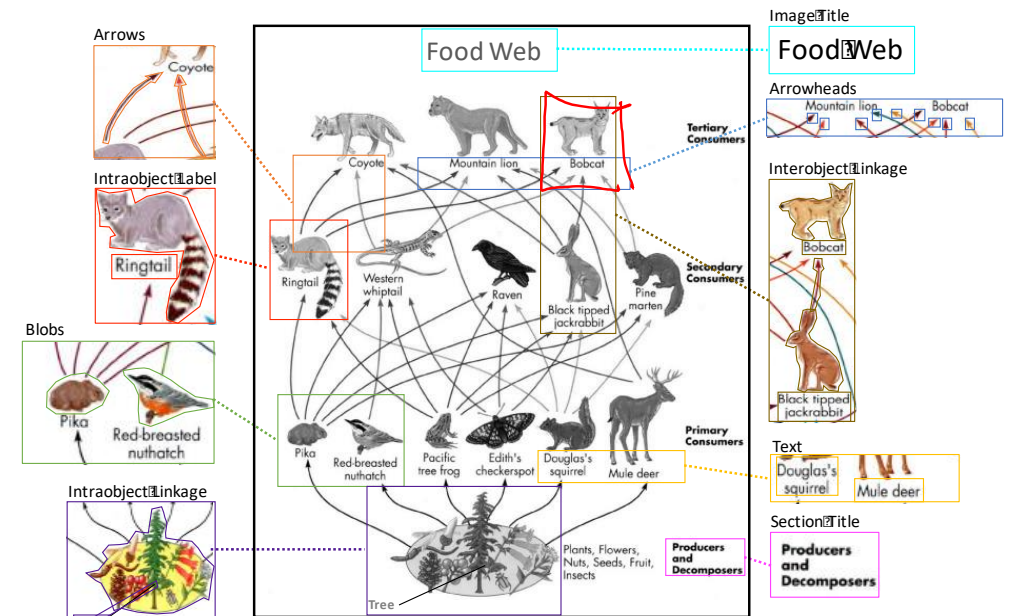
- Video games
- Pathing / routing problems
- Resource planning problems
- Robot motion planning
- Language analysis 
- Machine translation 
- Speech recognition
- ...

A* in Recent Literature

- Joint A* CCG Parsing and Semantic Role Labeling (EMLN'15)




- Diagram Understanding (ECCV'17)




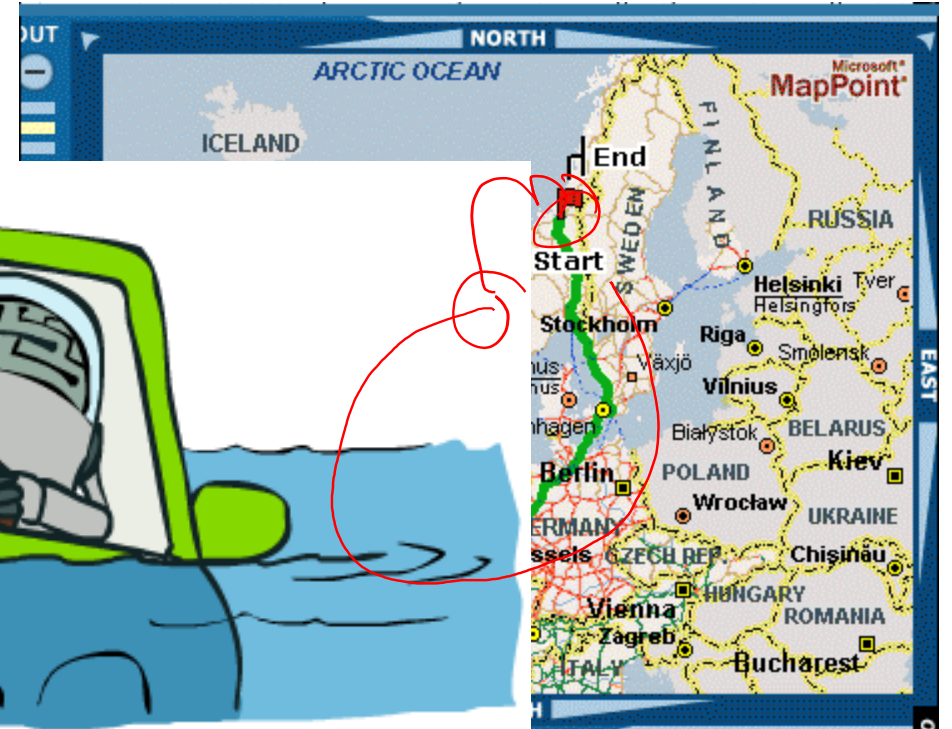


Multiple Choice Question: From the above food web diagram, what will lead to an increase in the population of deer? (A) increase in (B) decrease in (C) decrease in (D) increase in pika

Search and Models

- Search operates over models of the world
 - The agent doesn't actually try all the plans out in the real world!
 - Planning is all “in simulation”
 - Your search is only as good as your models...



Search Gone Wrong?



km 500 1000 Legend Zoom on map click

mi 200 400 600

Start: Haugesund, Rogaland, Norway
End: Trondheim, Sør-Trøndelag, Norway
Total Distance: 2713.2 Kilometers
Estimated Total Time: 47 hours, 31 minutes

nrk.no/alltidmoro