

CSE 473: Artificial Intelligence

Autumn 2018

Heuristic Search and A* Algorithms

Steve Tanimoto

With slides from :
Dieter Fox, Dan Weld, Dan Klein, Stuart Russell, Andrew Moore, Luke Zettlemoyer

Today

- A* Search
- Heuristic Design
- Graph search

Recap: Search

- **Search problem:**
 - States (configurations of the world)
 - Successor function: a function from states to lists of (state, action, cost) triples; drawn as a graph
 - Start state and goal test
- **Search tree:**
 - Nodes: represent plans for reaching states
 - Plans have costs (sum of action costs)
- **Search Algorithm:**
 - Systematically builds a search tree
 - Chooses an ordering of the fringe (unexplored nodes)

Example: Pancake Problem

Action: Flip over the top n pancakes



Cost: Number of pancakes flipped

Example: Pancake Problem

BOUNDS FOR SORTING BY PREFIX REVERSAL

William H. GATES

Microsoft, Albuquerque, New Mexico

Christos H. PAPADIMITRIOU*†

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

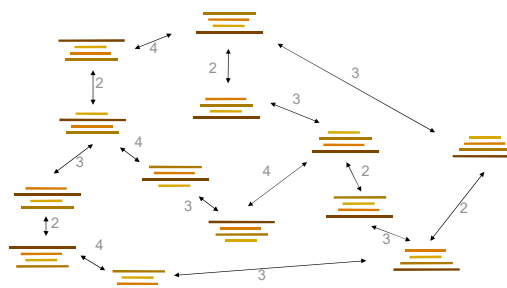
Received 18 January 1978

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

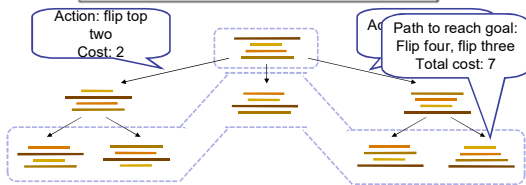


General Tree Search

```

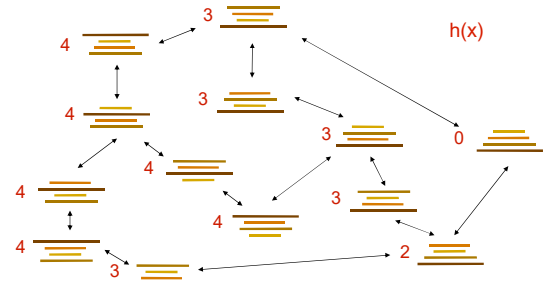
function TREE-SEARCH(problem, strategy) returns a solution, or failure
  initialize the search tree using the initial state of problem
  loop do
    if there are no candidates for expansion then return failure
    choose a leaf node for expansion according to strategy
    if the node contains a goal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

```



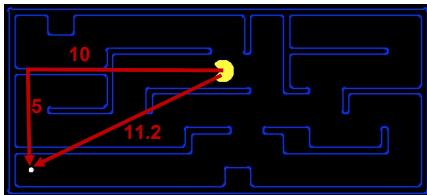
Example: Heuristic Function

Heuristic: the largest pancake that is still out of place



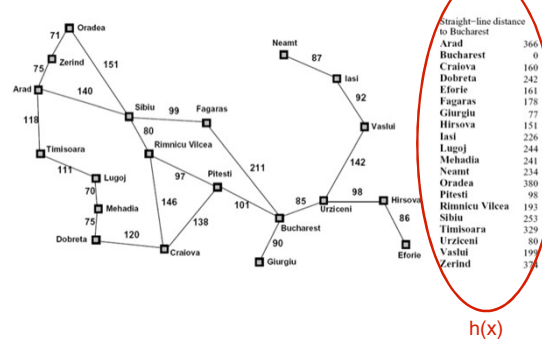
What is a *Heuristic*?

- An *estimate* of how close a state is to a goal
- Designed for a particular search problem



- Examples: Manhattan distance: $10 + 5 = 15$
- Euclidean distance: 11.2

Example: Heuristic Function

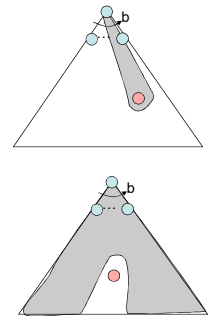


Greedy Search



Best First (Greedy)

- 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



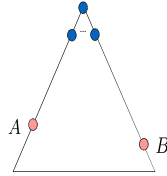
Optimality of A* Tree Search

Assume:

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

Claim:

- A will exit the fringe before B

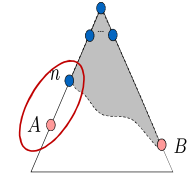


Optimality of A* Tree Search

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 $f(A)$



$$f(n) = g(n) + h(n) \quad \text{Definition of f-cost}$$

$$f(n) \leq g(A) \quad \text{Admissibility of h}$$

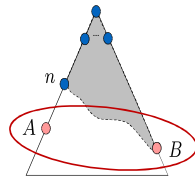
$$g(A) = f(A) \quad h = 0 \text{ at a goal}$$

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 $f(B)$



$$g(A) < g(B) \quad \text{B is suboptimal}$$

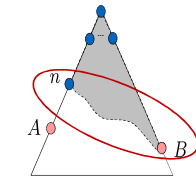
$$f(A) < f(B) \quad h = 0 \text{ at a goal}$$

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

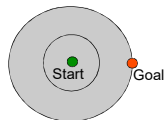


$$f(n) \leq f(A) < f(B)$$

- All ancestors of A expand before B
- A expands before B
- A* search is optimal

UCS vs A* Contours

- Uniform-cost expanded in all directions

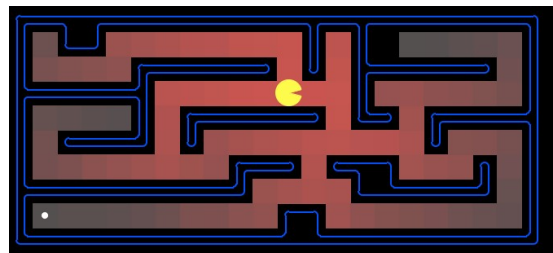


- A* expands mainly toward the goal, but hedges its bets to ensure optimality



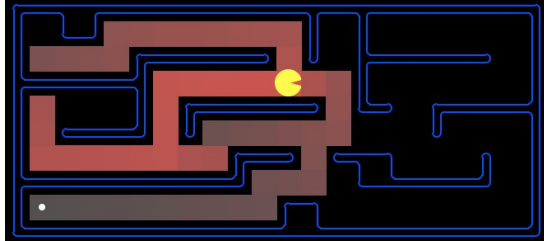
Which Algorithm?

- Uniform cost search (UCS):



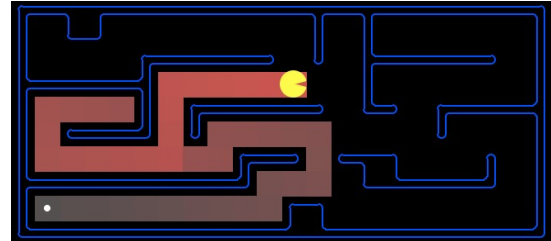
Which Algorithm?

- A*, Manhattan Heuristic:



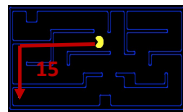
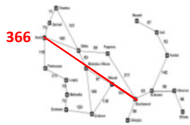
Which Algorithm?

- Best First / Greedy, Manhattan Heuristic:



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

Creating Heuristics

8-puzzle:

7	2	4
5		6
8	3	1

Start State

	1	2
3	4	5
6	7	8

Goal State

- What are the states?
- How many states?
- What are the actions?
- What states can I reach from the start state?
- What should the costs be?

8 Puzzle I

- Heuristic: Number of tiles misplaced

▪ $h(\text{start}) = 8$

- Is it admissible?

7	2	4
5		6
8	3	1

Start State

	1	2
3	4	5
6	7	8

Goal State

	Average nodes expanded when optimal path has length...		
	...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

▪ $h(\text{start}) = 3 + 1 + 2 + \dots = 18$

7	2	4
5		6
8	3	1

Start State

	1	2
3	4	5
6	7	8

Goal State

- Admissible?

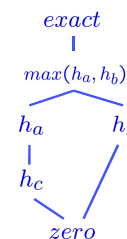
	Average nodes expanded when optimal path has length...		
	...4 steps	...8 steps	...12 steps
TILES	13	39	227
MANHATTAN	12	25	73

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!

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

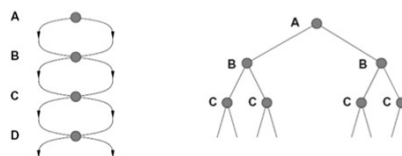


A* Applications

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

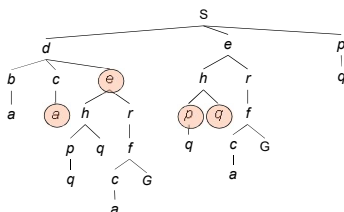
Tree Search: Extra Work!

- Failure to detect repeated states can cause exponentially more work. Why?



Graph Search

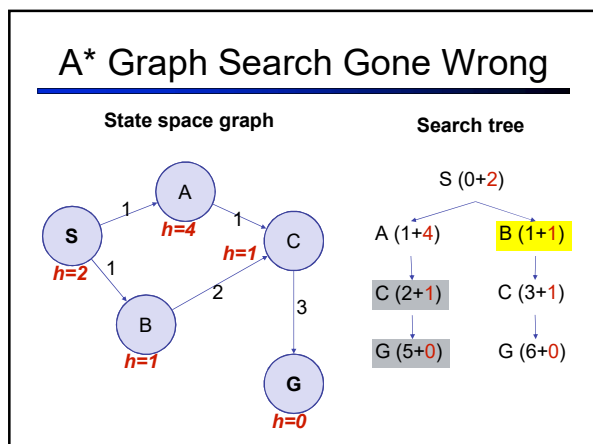
- In BFS, for example, we shouldn't bother expanding some nodes (which, and why?)



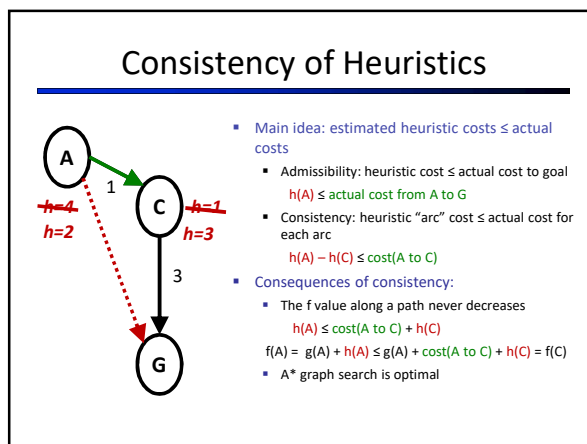
Graph Search

- 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
- Hint: in python, 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



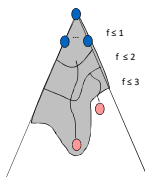
Consistency of Heuristics



Optimality of A* Graph Search

- Sketch: consider what A* does with a consistent heuristic:

- Nodes are popped with non-decreasing f-scores: for all n, n' with n' popped after n : $f(n') \geq f(n)$
 - Proof by induction: (1) always pop the lowest f-score from the fringe, (2) all new nodes have larger (or equal) scores, (3) add them to the fringe, (4) repeat!
- For every state s , nodes that reach s optimally are expanded before nodes that reach s sub-optimally
- Result: A* graph search is optimal



Optimality

- Tree search:**
 - A* optimal if heuristic is admissible (and non-negative)
 - UCS is a special case ($h = 0$)
- Graph search:**
 - A* optimal if heuristic is consistent
 - UCS optimal ($h = 0$ is consistent)
- Consistency implies admissibility
- In general, natural admissible heuristics tend to be consistent, especially if from relaxed problems

Summary: A*

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