# Heuristic Search: BestFS and A\*

#### Computer Science cpsc322, Lecture 8

#### (Textbook Chpt 3.5)

January, 21, 2009

$$\mathcal{M}$$

### **Course Announcements**

#### Posted on WebCT

- Second Practice Exercise (uninformed Search)
- Assignment 1

#### DEPARTMENT OF COMPUTER SCIENCE Distinguished Lecture Series 2008 - 2009 Speaker: Michael Littman Rutgers University Date: Thursday, January 22, 2009 Time: 3:30 - 4:50pm Venue: Hugh Dempster Pavilion Room 310 Title: Efficiently Learning to Behave Efficiently

### **Lecture Overview**

- Recap Heuristic Function
- Best First Search
- A\*



states? Where it is dirty and robot location actions? Left, Right, Suck Possible goal test? no dirt at all locations

#### **Lecture Overview**

## Recap Heuristic Function

- Best First Search
- A\*

#### **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 (for the end node).
- It treats the frontier as a priority queue ordered by h.
   (similar to ?) ∠ ∠ F ≤ ∠ ∞ ∠ 5 ℃
- This is a greedy approach: it always takes the path which appears locally best

### **Analysis of Best-First Search**

• Complete no: a low heuristic value can mean that a cycle gets followed forever.



- Optimal: no (why not?)
- Time complexity is O(b<sup>m</sup>)
- Space complexity is O(b<sup>m</sup>)

### **Lecture Overview**

- Recap Heuristic Function
- Best First Search
- A\* Search Strategy

## A\* Search Algorithm

- $A^*$  is a mix of:
  - lowest-cost-first and
  - best-first search



- $A^*$  treats the frontier as a priority queue ordered by  $f(p) = \zeta(\rho) + h(\rho)$
- It always selects the node on the frontier with the lowest distance.

## Analysis of A\*

Let's assume that arc costs are strictly positive.

- Time complexity is  $O(b^m)$ 
  - the heuristic could be completely uninformative and the edge costs could all be the same, meaning that A<sup>\*</sup> does the same thing as BFS
- Space complexity is O(b<sup>m</sup>) like <u>BFS</u>, A<sup>\*</sup> maintains a frontier which grows with the size of the tree
- Completeness: yes.
- Optimality: yes.

## Optimality of A<sup>\*</sup>

If A<sup>\*</sup> returns a solution, that solution is guaranteed to be optimal, as long as

When

- the branching factor is finite
- arc costs are strictly positive



Theorem

If A<sup>\*</sup> selects a path p, p is the shortest (i.e., lowest-cost) path.

## Why is $A^*$ optimal? $f = c_+ b_1$

- Assume for contradiction that some other path p' is actually the shortest path to a goal  $cost(p') \leq cost(p)$
- Consider the moment when p is chosen from the frontier. Some part of path p'will also be on the frontier; let's call this partial path p".



## Why is A<sup>\*</sup> optimal? (cont')

• Because p was expanded before p",  $+(p) \leq$ 

p

- Because p is a goal, h(p) = OThus C(p) < C(p') + h(p'')
- Because *h* is admissible, <u>cost(p") + h(p") ≤ cost(p')</u> for any path p'to a goal that extends p"

 $C(p) + h(p) \leq$ 

• Thus  $(p) \leq G(p)$  for any other path p'to a goal.

This contradicts our assumption that p' is the shortest path. that cost(p') < cost(p)

CPSC 322, Lecture 8

Slide 13



- In fact, we can prove something even stronger about A<sup>\*</sup>: in a sense (given the particular heuristic that is available) no search algorithm could do better!
- Optimal Efficiency. Among all optimal algorithms that start from the same start node and use the same heuristic h, A\* expands the minimal number of paths.

## Why is A<sup>\*</sup> optimally efficient?

**Theorem:** *A*<sup>\*</sup> is optimally efficient.

- Let *f*\*be the cost of the shortest path to a goal.
- Consider any algorithm A'
  - the same start node as  $A^*$ ,
  - uses the same heuristic
  - fails to expand some path p'expanded by A<sup>\*</sup>, for which f(p') < f<sup>\*</sup>.
- Assume that *A* is optimal.



## Why is A<sup>\*</sup> optimally efficient? (cont')

- Consider a different search problem
  - identical to the original
  - on which *h* returns the same estimate for each path
  - except that p'has a child path p"which is a goal node, and the true cost of the path to p" is f(p').
  - that is, the edge from p' to p'' has a cost of h(p'): the heuristic is exactly right about the cost of getting from p' to a goal.  $p' \bigcirc p$

## Why is A<sup>\*</sup> optimally efficient? (cont')

- A'would behave identically on this new problem.
  - The only difference between the new problem and the original problem is beyond path *p*', which *A*'does not expand.
- Cost of the path to p'' is lower than cost of the path found by A'.  $p' \_ \bigcirc p$

This violates our assumption that A' is optimal.

р"

## Learning Goals for today's class

- Define/read/write/trace/debug different search algorithms
  With / Without cost they use h
  Informed / Uninformed Beat First h
  A\*mm c+h
- Formally prove A\* optimality.

• Define optimally efficient and formally prove that A\* is optimally efficient to be done

#### **Next class**

Finish Search (finish Chpt 3)

- Branch-and-Bound
- A\* enhancements
- Non-heuristic Pruning
- Backward Search
- Dynamic Programming