# **Introduction to Artificial Intelligence**

Informed search algorithms

Chapter 4, Sections 1–2, 4

Dieter Fox

# Outline

- $\diamond$  Best-first search
- $\diamond$  A\* search
- $\diamond$  Heuristics
- $\diamond$  Hill-climbing
- $\diamondsuit\,$  Simulated annealing

## **Review: General search**

```
function GENERAL-SEARCH( problem, QUEUING-FN) returns a solution, or failure
nodes ← MAKE-QUEUE(MAKE-NODE(INITIAL-STATE[problem]))
loop do
    if nodes is empty then return failure
    node ← REMOVE-FRONT(nodes)
    if GOAL-TEST[problem] applied to STATE(node) succeeds then return node
    nodes ← QUEUING-FN(nodes, EXPAND(node, OPERATORS[problem]))
end
```

A strategy is defined by picking the order of node expansion

Idea: use an *evaluation function* for each node – estimate of "desirability"

 $\Rightarrow$  Expand most desirable unexpanded node

#### Implementation:

QUEUEINGFN = insert successors in decreasing order of desirability

Special cases:

greedy search A\* search

Based on AIMA Slides ©S. Russell and P. Norvig, 1998

### Romania with step costs in km



Evaluation function h(n) (heuristic) = estimate of cost from n to goal

E.g.,  $h_{SLD}(n)$  = straight-line distance from n to Bucharest

Greedy search expands the node that appears to be closest to goal









### **Properties of greedy search**

**Complete**??

Time??

Space??

**Optimal**??

Based on AIMA Slides ©S. Russell and P. Norvig, 1998

## **Properties of greedy search**

**Complete**: No–can get stuck in loops, e.g., Iasi  $\rightarrow$  Neamt  $\rightarrow$  Iasi  $\rightarrow$  Neamt  $\rightarrow$ Complete in finite space with repeated-state checking

**Time**:  $O(b^m)$ , but a good heuristic can give dramatic improvement

**Space**:  $O(b^m)$ —keeps all nodes in memory

**Optimal**: No

### $\mathbf{A}^*$ search

**Idea**: avoid expanding paths that are already expensive

Evaluation function f(n) = g(n) + h(n)

g(n) = cost so far to reach nh(n) = estimated cost to goal from nf(n) = estimated total cost of path through n to goal

A\* search uses an *admissible* heuristic i.e.,  $h(n) \le h^*(n)$  where  $h^*(n)$  is the *true* cost from n.

E.g.,  $h_{SLD}(n)$  never overestimates the actual road distance

Theorem: A\* search is optimal

Based on AIMA Slides ©S. Russell and P. Norvig, 1998

### $\mathbf{A}^*$ search example



### $A^*$ search example



#### $A^*$ search example



#### $A^*$ search example



#### A<sup>\*</sup> search example



Based on AIMA Slides ©S. Russell and P. Norvig, 1998

#### A<sup>\*</sup> search example



### Optimality of $A^*$

Suppose some suboptimal goal  $G_2$  has been generated and is in the queue. Let *n* be an unexpanded node on a shortest path to an optimal goal  $G_1$ .



$$f(G_2) = g(G_2) \qquad \text{since } h(G_2) = 0$$
  
>  $g(G_1) \qquad \text{since } G_2 \text{ is suboptimal}$   
$$\geq f(n) \qquad \text{since } h \text{ is admissible}$$

Since  $f(G_2) > f(n)$ , A\* will never select  $G_2$  for expansion

**Complete**?? Yes, unless there are infinitely many nodes with  $f \leq f(G)$ 

**Time**?? Exponential in [relative error in  $h \times$  length of soln.]

**Space**?? Keeps all nodes in memory

**Optimal**?? Yes—cannot expand  $f_{i+1}$  until  $f_i$  is finished

### **Admissible heuristics**

E.g., for the 8-puzzle:

 $h_1(n)$  = number of misplaced tiles  $h_2(n)$  = total Manhattan distance

(i.e., no. of squares from desired location of each tile)









 $h_1(S) =??$  $h_2(S) =??$ 

### **Admissible heuristics**

E.g., for the 8-puzzle:

 $h_1(n)$  = number of misplaced tiles  $h_2(n)$  = total Manhattan distance

(i.e., no. of squares from desired location of each tile)









 $h_1(S) =: 7$  $h_2(S) =: 2+3+3+2+4+2+0+2 = 18$ 

## Dominance

If  $h_2(n) \ge h_1(n)$  for all *n* (both admissible) then  $h_2$  dominates  $h_1$  and is better for search

Typical search costs:

 $\begin{array}{ll} d = 14 & {\sf IDS} = 3,473,941 \mbox{ nodes} \\ {\sf A}^*(h_1) = 539 \mbox{ nodes} \\ {\sf A}^*(h_2) = 113 \mbox{ nodes} \\ d = 24 & {\sf IDS} = \mbox{ too many nodes} \\ {\sf A}^*(h_1) = 39,135 \mbox{ nodes} \\ {\sf A}^*(h_2) = 1,641 \mbox{ nodes} \\ \end{array}$ 

Admissible heuristics can be derived from the *exact* solution cost of a *relaxed* version of the problem

If the rules of the 8-puzzle are relaxed so that a tile can move *anywhere*, then  $h_1(n)$  gives the shortest solution

If the rules are relaxed so that a tile can move to *any adjacent square*, then  $h_2(n)$  gives the shortest solution

### **Iterative improvement algorithms**

In many optimization problems, *path* is irrelevant; the goal state itself is the solution

Then state space = set of "complete" configurations; find *optimal* configuration, e.g., TSP or, find configuration satisfying constraints, e.g., n-queens

In such cases, can use *iterative improvement* algorithms; keep a single "current" state, try to improve it

Constant space, suitable for online as well as offline search

# **Example: Travelling Salesperson Problem**

Find the shortest tour that visits each city exactly once



Put n queens on an  $n \times n$  board with no two queens on the same row, column, or diagonal



# Hill-climbing (or gradient ascent/descent)

"Like climbing Everest in thick fog with amnesia"

**Problem**: depending on initial state, can get stuck on local maxima



# Simulated annealing

**Idea**: escape local maxima by allowing some "bad" moves but gradually decrease their size and frequency

```
function SIMULATED-ANNEALING (problem, schedule) returns a solution state
  inputs: problem, a problem
             schedule, a mapping from time to "temperature"
  local variables: current, a node
                       next, a node
                       T, a "temperature" controlling the probability of downward steps
  current \leftarrow Make-Node(INITIAL-STATE[problem])
  for t \leftarrow 1 to \infty do
       T \leftarrow schedule[t]
       if T=0 then return current
       next \leftarrow a randomly selected successor of current
       \Delta E \leftarrow \text{VALUE}[next] - \text{VALUE}[current]
       if \Delta E > 0 then current \leftarrow next
       else current \leftarrow next only with probability e^{\Delta E/T}
```

# **Properties of simulated annealing**

T decreased slowly enough  $\Longrightarrow$  always reach best state

Is this necessarily an interesting guarantee??

Devised by Metropolis et al., 1953, for physical process modelling

Widely used in VLSI layout, airline scheduling, etc.