Today

- Uninformed search: summary
- Informed search,
- Search Heuristics

See Russell and Norvig, Chapters 3,4

Summary of algorithms

Criterion	Breadth-	Uniform-	Depth-	Depth-	Iterative
	First	Cost	First	Limited	Deepening
Complete? Time Space Optimal?	Yes* b^{d+1} b^{d+1} Yes*	Yes* $b^{\lceil C^*/\epsilon \rceil}$ $b^{\lceil C^*/\epsilon \rceil}$ Yes*	No b ^m bm	Yes, if $l \ge d$ b^l bl	Yes b^d bd

Here * indicates conditions stated earlier.



Graph search

The state space with actions leading from state to state corresponds naturally to a **graph** rather than a **tree**; the state appears only once in the graph.

There are data structures corresponding to graphs, and graph search algorithms that avoid repetition of states already seen.

The idea is to keep track of nodes that have already been expanded; if search arrives back at such a node, it is ignored in future search.

See Russell and Norvig for details.

Informed search strategies

Informed strategies use heuristic "rule of thumb" ideas to guide search based on some estimation of where the solution is most likely to be found.

We look at some such strategies:

- Best-first search
- A* search
- Heuristics

informatics



- estimate of "desirability"

 \Rightarrow Expand most desirable unexpanded node

Implementation:

fringe is a queue sorted in decreasing order of desirability

Special cases:

Alan Smaill

greedy search

 A^* search

Alan Smaill	Fundamentals of Artificial Intelligence	Oct 8, 2007

Reminder: Searching a Tree

if GOAL-TEST[*problem*] applied to STATE(*node*) succeeds **return** *node*

function TREE-SEARCH(problem, fringe) returns a solution, or failure

 $fringe \leftarrow \text{INSERTALL}(\text{EXPAND}(node, problem), fringe)$

A strategy is defined by picking the order of node expansion

if *fringe* is empty then return failure

 $node \leftarrow \text{REMOVE-FRONT}(fringe)$

ob gool

 $fringe \leftarrow \text{INSERT}(\text{MAKE-NODE}(\text{INITIAL-STATE}[problem]), fringe)$

informatics

Romania with step costs in km

Use a straight line heuristic: distance to goal in straight line.

Arad	366	Mehadia	241
Bucharest	0	Neamt	234
Craiova	160	Oradea	380
Dobreta	242	Pitesti	100
Eforie	161	Rimnicu Vilcea	193
Fagaras	176	Sibiu	253
Giurgiu	77	Timisoara	329
Hirsova	151	Urziceni	80
Iasi	226	Vaslui	199
Lugoj	244	Zerind	374

informatics

Oct 8, 2007

Greedy search

Fundamentals of Artificial Intelligence

Evaluation function h(n) (heuristic)

= estimate of cost from \boldsymbol{n} to the closest goal

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

Greedy search expands the node that **appears** to be closest to goal prefer the action that takes us to the state with the **minimum** heuristic cost.



Alan Smaill



Greedy search example



Properties of greedy search

 $\label{eq:complete} \underbrace{ \mbox{Complete} ?? \mbox{No-can get stuck in loops, e.g., with Oradea as goal, } \\ \hline \mbox{Iasi} \rightarrow \mbox{Neamt} \rightarrow \mbox{Iasi} \rightarrow \mbox{Neamt} \rightarrow \\ \mbox{Complete in finite space with repeated-state checking} \end{cases}$

Time??



Properties of greedy search

<u>Complete</u>?? No-can get stuck in loops, e.g., <u>lasi</u> \rightarrow Neamt \rightarrow lasi \rightarrow Neamt \rightarrow Complete in finite space with repeated-state checking T: 22 O(lm) by the state sta

<u>Time</u>?? $O(b^m)$, but a good heuristic can give dramatic improvement

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

Optimal?? No

18 informatics

\mathbf{A}^* search

Idea: avoid expanding paths that are already expensive

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

 $g(n) = \operatorname{cost}$ so far to reach n

h(n) = estimated cost to goal from n

 $f(\boldsymbol{n}) = \text{estimated total cost of path through } \boldsymbol{n}$ to goal

A* search uses an **admissible** heuristic, i.e.

 $h(n) \le h^*(n)$

where $h^*(n)$ is the **true** cost of **cheapest** solution from n. (Also require $h(n) \ge 0$, so h(G) = 0 for any goal G.)

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





Alan Smaill









Optimality of A* (standard proof)

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 .



Then:

 $\begin{array}{rcl} f(G_2) &=& g(G_2) & \text{since } h(G_2) = 0 \\ &>& g(G_1) & \text{since } G_2 \text{ is suboptimal} \\ &\geq& f(n) & \text{since } h \text{ is admissible} \end{array}$

Properties of A*

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

Since $f(G_2) > f(n)$, A^{*} will never select G_2 for expansion.

Δ.		<u> </u>		
A	lan	Sn	nai	П

Time??

Fundamentals of Artificial Intelligence

Oct 8, 2007

28 informatics

27 informatics

Optimality of A* (more useful)

Lemma: A^* expands nodes in order of increasing f value

Gradually adds "*f*-contours" of nodes (cf. breadth-first adds layers) Contour *i* has all nodes with $f = f_i$, where $f_i < f_{i+1}$



Alan Smaill

Properties of A*

<u>Complete</u>?? Yes, unless there are infinitely many nodes with $f \le f(G)$ <u>Time</u>?? Exponential in [relative error in $h \times$ length of soln.] Space??

Properties of A*

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

<u>Time</u>?? 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

- A^* expands all nodes with $f(n) < C^*$
- A^* expands some nodes with $f(n) = C^*$
- A^* expands no nodes with $f(n) > C^*$



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 for solution at length d:

$$\begin{array}{ll} d = 14 & {\sf IDS} = 3,473,941 \mbox{ nodes} \\ {\sf A}^*(h_1) = 539 \mbox{ nodes} \\ {\sf A}^*(h_2) = 113 \mbox{ nodes} \\ d = 24 & {\sf IDS} \approx 54,000,000,000 \mbox{ nodes} \\ {\sf A}^*(h_1) = 39,135 \mbox{ nodes} \\ {\sf A}^*(h_2) = 1,641 \mbox{ nodes} \end{array}$$

Relaxed problems

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

Key point: the optimal solution cost of a relaxed problem is no greater than the optimal solution cost of the real problem



- A* is optimal search strategy
- Heuristics for the eight puzzle