

















How do we plan our holiday?

- This kind of hypothetical reasoning involves asking
- "what state will I be in after the following sequence of events?"
- From this we can reason about what sequence of events one should try to bring about to achieve a desirable state.
- Search is a computational method for capturing a particular version of this kind of reasoning.

CSE 3401 Fall 2012 Yves Lesperance & Fahiem Bacchu

Search

- There are many difficult questions that are not resolved by search. In particular, the whole question of how does an intelligent system formulate its problem as a search problem is not addressed by search.
- Search only shows how to solve the problem once we have it correctly formulated.

CSE 3401 Fall 2012 Yves Lesperance & Fahiem Bacchus

The formalism.

- To formulate a problem as a search problem we need the following components:
- Formulate a state space over which to search. The state space necessarily involves abstracting the real problem.
- Formulate actions that allow one to move between different states. The actions are abstractions of actions you could actually perform.
- Identify the initial state that best represents your current state and the desired condition one wants to achieve.
- Formulate various heuristics to help guide the search process.

CSE 3401 Fall 2012 Yves Lesperance & Fahiem Bacchu

Example 1: Romania Travel.

13

The formalism.
Once the problem has been formulated as a state space search, various algorithms can be utilized to solve the problem.
A solution to the problem will be a sequence of actions/moves that can transform your current state into state where your desired condition holds.

Example 1.

CSE 3401 Fall 2012 Yves Lesperance & Fahiem Bacchus

• State space.

- States: the various cities you could be located in.
 - Note we are ignoring the low level details of driving, states where you are on the road between cities, etc.
- Actions: drive between neighboring cities.
- Initial state: in Arad
- Desired condition (Goal): be in a state where you are in Bucharest. (How many states satisfy this condition?)

CSE 3401 Fall 2012 Yves Lesperance & Fahiem Bacchu

• Solution will be the route, the sequence of cities to travel through to get to Bucharest.



16





Example 2. The 8-Puzzle

- Although there are 9! different configurations of the tiles (362,880), in fact the state space is divided into two disjoint parts.
- Only when the blank is in the middle are all four actions possible.
- Our goal condition is satisfied by only a single state. But one could easily have a goal condition like
 - The 8 is in the upper left hand corner.
 - How many different states satisfy this goal?

19

Example 3. Vacuum World.

- In the previous two examples, a state in the search space corresponded to a unique state of the world (modulo details we have abstracted away).
- However, states need not map directly to world configurations. Instead, a state could map to the agent's mental conception of how the world is configured: the agent's knowledge state.

CSE 3401 Fall 2012 Yves Lesperance & Fahiem Bacchus



















More complex situations.

- The agent might be able to perform some sensing actions. These actions change the agent's mental state, not the world configuration.
- With sensing can search for a contingent solution: a solution that is contingent on the outcome of the sensing actions
 - $\blacksquare\!<\!\!\text{right},$ if dirt then suck>
- Now the issue of interleaving execution and search comes into play.

CSE 3401 Fall 2012 Yves Lesperance & Fahiem Bacchu

29

More complex situations.

- Instead of complete lack of knowledge, the agent might think that some states of the world are more likely than others.
- This leads to probabilistic models of the search space and different algorithms for solving the problem.
- Later we will see some techniques for reasoning and making decisions under uncertainty.

Algorithms for Search.

- Inputs:
 - a specified initial state (a specific world state or a set of world states representing the agent's knowledge, etc.)
 - **a** successor function $S(x) = \{set of states that can be reached from state x via a single action\}.$
 - a goal test a function that can be applied to a state and returns true if the state is satisfies the goal condition.
 - A step cost function C(x,a,y) which determines the cost of moving from state x to state y using action a. $(C(x,a,y) = \infty \text{ if a does not yield y from } x)$

31

Algorithms for Search.

CSE 3401 Fall 2012 Yves Lesperance & Fahiem Bacchus

- Output:
 - a sequence of states leading from the initial state to a state satisfying the goal test.
 - The sequence might be
 - annotated by the name of the action used.

CSE 3401 Fall 2012 Yves Lesperance & Fahiem Bacchus

• optimal in cost for some algorithms.















 Breadth First Example.

 {0}

 {1,2}

 {2,3,3}

 {2,3,3,4}

 {3,4,3,4}

 {3,4,3,4,5}

 ...

 •[Draw search tree]



Breadth First Properties

- Time Complexity?
- # nodes generated at...
- Level 0 (root): 1
- Level 1: 1* b [each node has at most b successors]
- Level 2: $b^* b = b^2$
- Level 3: $b * b^2 = b^3 \dots$
- Level d: bd
- Level d + 1: $b^{d+1} b = b(b^d 1)$ [when last node is successful]
- Total: $1 + b + b^2 + b^3 + ... + b^{d-1} + b^d + b(b^d 1) = O(b^{d+1})$
- Exponential, so can only solve small instances

47

CSE 3401 Fall 2012 Yves Lesperance & Fahiem Ba





Breadth First Properties

- Space complexity is a real problem.
- E.g., let b = 10, and say 1000 nodes can be expanded per second and each node requires 100 bytes of storage:

Depth	Nodes	Time	Memory
1	1	1 millisec.	100 bytes
6	10 ⁶	18 mins.	111 MB
8	10 ⁸	31 hrs.	11 GB

• Run out of space long before we run out of time in most applications.

50

CSE 3401 Fall 2012 Yves Lesperance & Fahiem

Uniform Cost Search. Keep the frontier sorted in increasing cost of the path to a node; behaves like priority queue. Always expand the least cost node. Identical to Breadth First if each transition has the same cost. Example:

- let the states be the positive integers {0,1,2,...}
- let each state n have as successors n+1 and n+2
- Say that the n+1 action has cost 2, while the n+2 action has cost 3.

51

[Draw search space graph]

CSE 3401 Fall 2012 Yves Lesperance & Fahiem Bacchus

Uniform Cost Search. {0[0]} {1[2],2[3]} {2[3],2[4],3[5]} {2[4],3[5],3[5],4[6]} {3[5],3[5],4[6],3[6],4[7]} ...





Uniform-Cost Search. Proof of Optimality.

- Let c(n) be the cost of the path to node n. If n2 is expanded after n1 then c(n1) ≤ c(n2).
 - If n2 was on the frontier when n1 was expanded, in which case $c(n2) \ge c(n1)$ else n1 would not have been selected for expansion.
 - If n2 was added to the frontier when n1 was expanded, in which case $c(n2) \ge c(n1)$ since the path to n2 extends the path to n1.
 - If n2 is a successor of a node n3 that was on the frontier or added when n1 was expanded, then c(n2) > c(n3) and $c(n3) \ge c(n1)$ by the above arguments.

56

CSE 3401 Fall 2012 Yves Lesperance & Fahiem

Uniform-Cost Search. Proof of Optimality.

2. When n is expanded every path with cost strictly less than c(n) has already been expanded (i.e., every node on it has been expanded).

Proof:

- Let <Start, n0, n1, ..., nk> be a path with cost less than c(n). Let ni be the last node on this path that has been expanded. <Start, n0, n1, ni-1, ni, ni+1, ..., nk>.
- ni+1 must be on the frontier, also c(ni+1) < c(n) since the cost of the entire path to nk is < c(n).
- But then uniform-cost would have expanded ni+1 not n!
- So every node on this path must already be expanded, i.e. this path has already been expanded. QED

57

CSE 3401 Fall 2012 Yves Lesperance & Fahiem Bacchu

Uniform-Cost Search. Proof of Optimality.

3. The first time uniform-cost expands a state, it has found the minimal cost path to it (it might later find other paths to the same state).

Proof:

- No cheaper path exists, else that path would have been expanded before.
- No cheaper path will be discovered later, as all those paths must be at least as expensive.

58

CSE 3401 Fall 2012 Yves Lesperance & Fahiem

So, when a goal state is expanded, the path to it must be optimal.

Depth First Search Depth First Search Example (applied to the example of Breadth First • Place the successors of the current state at search) the front of the frontier. {0} • Frontier behaves like a stack. {1,2} {2,3,2} {3,4,3,2} {4,5,4,3,2} {5,6,5,4,3,2} . . . [draw search tree] CSE 3401 Fall 2012 Yves Lesperance & Fahiem I CSE 3401 Fall 2012 Yves Lesperance & Fahien 59 60







Depth Limited Search



Iterative Deepening Search. Take the idea of depth limited search one step further. Starting at depth limit L = 0, we iteratively increase the depth limit, performing a depth limited search for each depth limit. Stop if no solution is found, or if the depth limited search failed without cutting off any nodes because of the depth limit.

67

Iterative Deepening Search Example

$\{0\} [DL = 0]$	$\{0\} [DL = 3]$ $\{1,2\}$
{0} [DL = 1] {1,2} {2}	$\{2,3,2\}$ $\{3,4,3,2\}, \{4,3,2\}, \{3,2\}$ $\{4,5,2\}, \{5, 2\}$
{0} [DL = 2] {1,2} {2,3,2}, {3,2}, {2} {3, 4}, {4}	Success!
	CSE 3401 Fall 2012 Yves Lesperance & Fahiem Bacchus

















