

Artificial Intelligence

CS 165A

Apr 21, 2022

Instructor: Prof. Yu-Xiang Wang

Today

→ Search algorithms

Coding Project 1 is due today


- Submit both your code and report
- Declare your collaboration (help you've received)

Recap: Problem Formulation and Search

- Problem formulation
 - State-space description $\langle \{S\}, S_0, \{S_G\}, \{O\}, \{g\} \rangle$
 - **S**: Possible states
 - **S₀**: Initial state of the agent
 - **S_G**: Goal state(s)
 - Or equivalently, a goal test **G(S)**
 - **O**: Operators $O: \{S\} \Rightarrow \{S\}$
 - Describes the possible actions of the agent
 - **g**: Path cost function, assigns a cost to a path/action
- At any given time, which possible action **O_i** is best?
 - Depends on the goal, the path cost function, the future sequence of actions....
- Agent's strategy: Formulate, Search, and Execute
 - This is *offline* problem solving


Recap: PACMAN

- The goal of a simplified PACMAN is to get to the pellet as quick as possible.
 - For a grid of size 30*30. Everything static.
 - What is a reasonable representation of the State, Operators, Goal test and Path cost?



Quiz: PACMAN with static ghosts

- The goal is to eat all pellets as quickly as possible while staying alive. Eating the “Power pellet” will allow the pacman to eat the ghost.



- State (how many?)
- Operators?
- Goal-Test?
- Path-Cost?

Recap: General Tree Search Algorithm

- Uses a queue (a list) and a **queuing function** to implement a *search strategy*
 - **Queuing-Fn**(*queue*, *elements*) inserts a set of elements into the queue and determines the order of node expansion

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

Recap: Breadth-First Search

- All nodes at depth d in the search tree are expanded before any nodes at depth $d+1$
 - First consider all paths of length N , then all paths of length $N+1$, etc.
- Doesn't consider path cost – finds the solution with the shortest path
- Uses FIFO queue

```
function BREADTH-FIRST-SEARCH(problem) returns a solution or failure  
return GENERAL-SEARCH(problem, ENQUEUE-AT-END)
```

Recap: Breadth-First Search

- Complete? **Yes**
- Optimal? **If shallowest goal is optimal**
- Time complexity? **Exponential: $O(b^{d+1})$**
- Space complexity? **Exponential: $O(b^{d+1})$**

In practice, the memory requirements are typically worse than the time requirements

b = branching factor (require finite b)
d = depth of shallowest solution

This lecture: Search algorithms

- Uninformed search
 - DFS
 - Depth-limited search
 - Iterative Deepening search
 - Bidirectional search
 - Uniform cost search
- Tree search vs Graph search
- Informed Search
 - A*-Search


Depth-First Search

- Always expands one of the nodes at the deepest level of the tree
 - Low memory requirements
 - Problem: depth could be infinite
- Uses a stack (LIFO)


```
function DEPTH-FIRST-SEARCH(problem) returns a solution or failure  
return GENERAL-SEARCH(problem, ENQUEUE-AT-FRONT)
```

Example

State space graph



Search tree



Queue

- (A)
- (B C)
- (D C)
- (C)
- (B D E)
- (D D E)
- (D E)
- (E)
- (F)

Depth-First Search

- Complete? No
- Optimal? No
- Time complexity? Exponential: $O(b^m)$
- Space complexity? Polynomial: $O(bm)$

m = maximum depth of the search tree
(may be infinite)

What is the difference between the BFS / DFS that you learned from the algorithm / data structure course?

- Nothing, except:
 - Now you are applying them to solve an AI problem
 - The graph can be infinitely large
 - The graph does not need to be known ahead of time (you only need local information: goal-state checker, successor function)

Space complexity of DFS

- Why is the *space* complexity (memory usage) of depth-first search $O(bm)$?
 - Remove expanded node when all descendents evaluated
 - At each of the m levels, you have to keep b nodes in memory


Example:

$$b = 3$$

$$m = 6$$

$$\text{Nodes in memory: } bm+1 = 19$$

Actually, $(b-1)m + 1 = 13$ nodes, the way we have been keeping our node list



Depth-Limited Search

- Like depth-first search, but uses a depth cutoff to avoid long (possibly infinite), unfruitful paths
 - Do depth-first search up to depth limit l
 - Depth-first is special case with limit = inf
- Problem: How to choose the depth limit l ?
 - Some problem statements make it obvious (e.g., TSP), but others don't (e.g., MU-puzzle, from the supplementary slide last time)

```
function DEPTH-LIMITED-SEARCH(problem, depth-limit) returns a  
  solution or failure  
return GENERAL-SEARCH(problem, ENQUEUE-AT-FRONT-IF-UNDER-  
  DEPTH-LIMIT)
```

Must explicitly represent node depth

Depth-Limited Search

l = depth limit

- Complete? No, unless $d \leq l$
- Optimal? No
- Time complexity? Exponential: $O(b^l)$
- Space complexity? Exponential: $O(bl)$

Iterative-Deepening Search

- Since the depth limit is difficult to choose in depth-limited search, use depth limits of $l = 0, 1, 2, 3, \dots$
 - Do depth-limited search at each level

```
function ITERATIVE-DEEPENING-SEARCH(problem) returns a solution or failure
  for depth  $\leftarrow$  0 to  $\infty$  do
    if DEPTH-LIMITED-SEARCH(problem, depth) succeeds then return result
  end
return failure
```

Iterative-Deepening Search


- IDS has advantages of
 - Breadth-first search – similar optimality and completeness guarantees
 - Depth-first search – Modest memory requirements
- This is the preferred blind search method when the search space is *large* and the solution depth is *unknown*
- Many states are expanded multiple times
 - Is this terribly inefficient?
 - No... and it's great for memory (compared with breadth-first)
 - Why is it not particularly inefficient?

Iterative-Deepening Search: Efficiency

- Complete? **Yes**
- Optimal? **Same as BFS**
- Time complexity? **Exponential: $O(b^d)$**
- Space complexity? **Polynomial: $O(bd)$**


Bidirectional Search

Forward search only:




Bidirectional Search

Simultaneously search forward from the initial state and backward from the goal state



Much more efficient!

Bidirectional Search



Example:
 $4^{10} \approx 1,000,000$
 $2 * 4^5 \approx 2,000$

- $O(b^{d/2})$ rather than $O(b^d)$ – hopefully
- Both actions and predecessors (inverse actions) must be defined
- Must test for intersection between the two searches
 - Constant time for test?
- Really a search strategy, not a specific search method
 - Often not practical....

Bidirectional Search

- Complete? Yes
- Optimal? Same as BFS
- Time complexity? Exponential: $O(b^{d/2})$
- Space complexity? Exponential: $O(b^{d/2})$

* Assuming breadth-first search used from both ends


Uniform Cost Search

- Similar to breadth-first search, but always expands the lowest-cost node, as measured by the path cost function, $g(n)$
 - $g(n)$ is (actual) cost of getting to node n
 - Breadth-first search is actually a special case of uniform cost search, where $g(n) = \text{DEPTH}(n)$
 - If the path cost is **monotonically increasing**, uniform cost search will find the optimal solution

function **UNIFORM-COST-SEARCH**(*problem*) **returns** a solution or failure
return **GENERAL-SEARCH**(*problem*, **ENQUEUE-IN-COST-ORDER**)

(Dijkstra's algorithm of an potentially infinite graph)

Example (3 min work)




Try breadth-first and uniform cost

Example (3 min work): Breath-First Search

Node to expand:


Frontier:



Example (3 min work): Uniform Cost Search

Node to expand:

Frontier:



Uniform-Cost Search

C = optimal cost
 ϵ = minimum step cost

- Complete? Yes, if $\epsilon > 0$
- Optimal? Yes
- Time complexity? Exponential: $O(b^{\lfloor C/\epsilon \rfloor})$
- Space complexity? Exponential: $O(b^{\lfloor C/\epsilon \rfloor})$


Same as breadth-first if all edge costs are equal

Can we do better than Tree Search?

- Sometimes.
- When the number of states are small
 - Dynamic programming (smart way of doing exhaustive search)


State Space vs. Search Tree (cont.)

Search tree (partially expanded)



Search Tree => Search Graph

Dynamic programming (with book keeping)



Graph Search vs Tree Search

- Tree Search
 - We might repeat some states
 - But we do not need to remember states
- Graph Search
 - We remember all the states that have been explored
 - But we do not repeat some states

Summary table of uninformed search

Criteria	BFS	Uniform-cost	DFS	Depth-limited	IDS	Bidirectional
Complete?	Yes [#]	Yes ^{#&}	No	No	Yes [#]	Yes ^{#+}
Time	$O(b^d)$	$O(b^{1+[C^*/e]})$	$O(b^m)$	$O(b^l)$	$O(b^d)$	$O(b^{d/2})$
Space	$O(b^d)$	$O(b^{1+[C^*/e]})$	$O(bm)$	$O(bl)$	$O(bd)$	$O(b^{d/2})$
Optimal?	Yes ^{\$}	Yes	No	No	Yes ^{\$}	Yes ^{\$+}

b : Branching factor

d : Depth of the shallowest goal

l : Depth limit

m : Maximum depth of search tree

e : The lower bound of the step cost

[#]: Complete if b is finite

[&]: Complete if step cost $\geq e$

^{\$}: Optimal if all step costs are identical

⁺: If both direction use BFS

(Section 3.4.6 in the AIMA book.)

Practical note about search algorithms

- The computer can't "see" the search graph like we can
 - No "bird's eye view" – make relevant information explicit!
- What information should you keep for a node in the search tree?
 - State
 - (1 2 0)
 - Parent node (or perhaps complete ancestry)
 - Node #3 (or, nodes 0, 2, 5, 11, 14)
 - Depth of the node
 - $d = 4$
 - Path cost up to (and including) the node
 - $g(\text{node}) = 12$
 - Operator that produced this node
 - Operator #1

Remainder of the lecture

- Informed search
- Some questions / desiderata
 1. Can we do better with some side information?
 2. We do not wish to make strong assumptions on the side information.
 3. If the side information is good, we hope to do better.
 4. If the side information is useless, we perform as well as an uninformed search method.

Best-First Search (with an Eval-Fn)

function **BEST-FIRST-SEARCH**(*problem*, EVAL-FN) **returns** a solution or failure

QUEUING-FN ← a function that orders nodes by EVAL-FN

return **GENERAL-SEARCH**(*problem*, QUEUING-FN)

- Uses a heuristic function, $h(n)$, as the EVAL-FN
- $h(n)$ estimates the cost of the best path from state n to a goal state
 - $h(goal) = 0$

Greedy Best-First Search

- Greedy search – always expand the node that appears to be the closest to the goal (i.e., with the smallest h)
 - Instant gratification, hence “greedy”

```
function GREEDY-SEARCH(problem, h) returns a solution or failure  
return BEST-FIRST-SEARCH(problem, h)
```

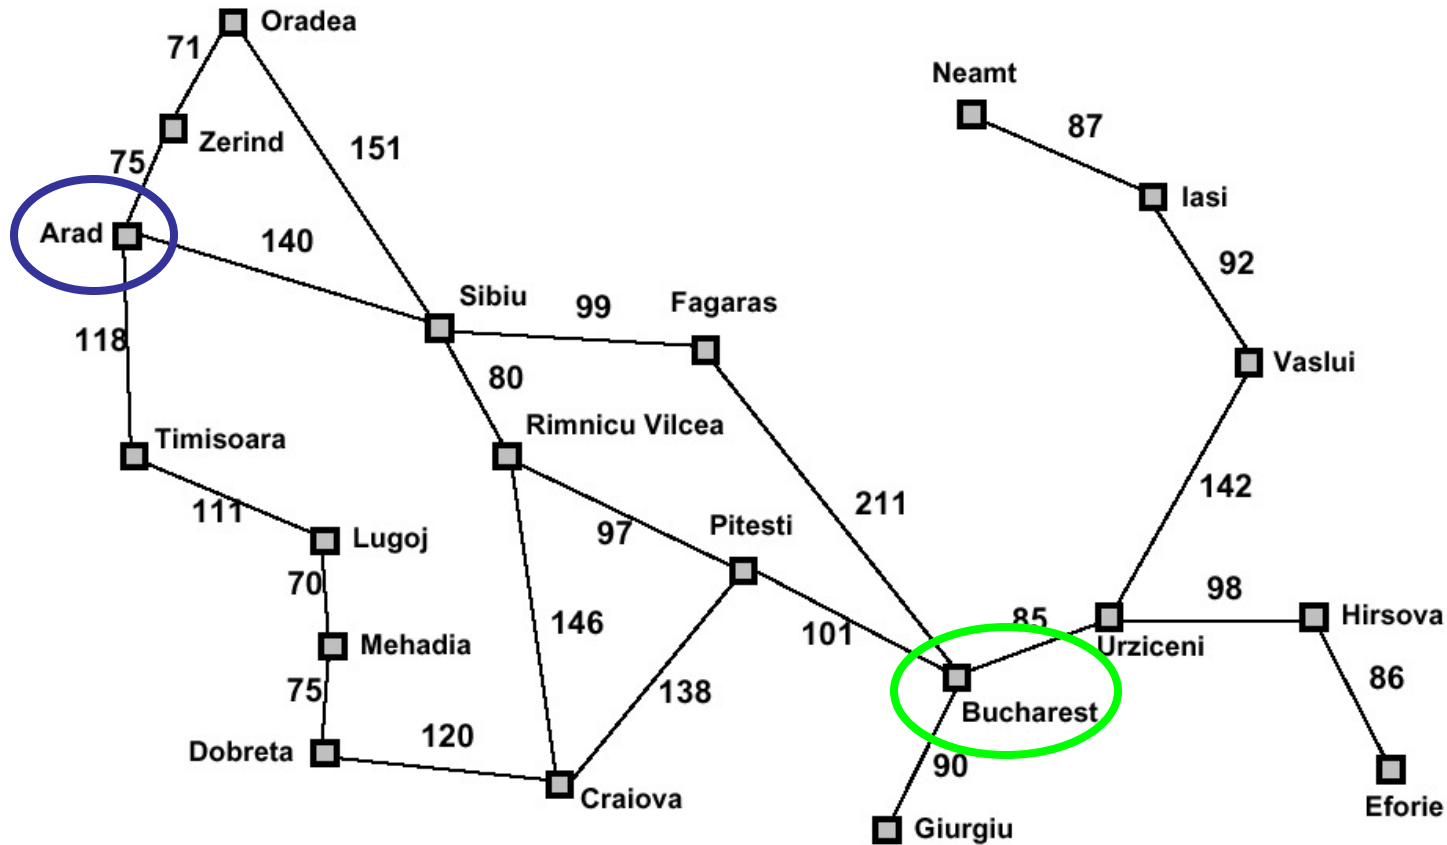

- Greedy search often performs well, but:
 - It doesn't always find the best solution / or any solution
 - It may get stuck
 - Its performance completely depends on the particular h function

A* Search (Pronounced “A-Star”)

- Uniform-cost search minimizes $g(n)$ (“past” cost)
- Greedy search minimizes $h(n)$ (“expected” or “future” cost)
- “A* Search” combines the two:
 - Minimize $f(n) = g(n) + h(n)$
 - Accounts for the “past” and the “future”
 - Estimates the cheapest solution (complete path) through node n

```
function A*-SEARCH(problem, h) returns a solution or failure  
return BEST-FIRST-SEARCH(problem, f)
```


A* Example




Straight-line distance to Bucharest

Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	178
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	98
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

$$f(n) = g(n) + h(n)$$



A* Example




When does A* search “work”?

- Focus on optimality (finding the optimal solution)
- “A* Search” is optimal if h is **admissible**
 - h is optimistic – it never overestimates the cost to the goal
 - $h(n) \leq$ true cost to reach the goal
 - So $f(n)$ never overestimates the actual cost of the best solution passing through node n

Visualizing A* search

- A* expands nodes in order of increasing f value
- Gradually adds " f -contours" of nodes
- Contour i has all nodes with $f=f_i$, where $f_i < f_{i+1}$
-



Optimality of A^* with an Admissible h

- Let OPT be the optimal path cost.
 - All non-goal nodes on this path have $f \leq \text{OPT}$.
 - Positive costs on edges
 - The goal node on this path has $f = \text{OPT}$.
- A^* search does not stop until an f -value of OPT is reached.
 - All other goal nodes have an f cost higher than OPT.
- All non-goal nodes on the optimal path are eventually expanded.
 - The optimal goal node is eventually placed on the priority queue, and reaches the front of the queue.

Optimal Efficiency of A*

A* is **optimally efficient** for any particular $h(n)$

That is, no other optimal algorithm is guaranteed to expand fewer nodes with the same $h(n)$.

- Need to find a good and efficiently evaluable $h(n)$.

A* Search with an Admissible h


- Optimal? Yes
- Complete? Yes
- Time complexity? Exponential; better under some conditions
- Space complexity? Exponential; keeps all nodes in memory

Recall: Graph Search vs Tree Search

- Tree Search
 - We might repeat some states
 - But we do not need to remember states
- Graph Search
 - We remember all the states that have been explored
 - But we do not repeat some states

Avoiding Repeated States using A* Search

- Is GRAPH-SEARCH optimal with A*?



Graph Search


Step 1: Among B, C, E, Choose C

Step 2: Among B, E, D, Choose B

Step 3: Among D, E, Choose E. (you are not going to select C again)

Avoiding Repeated States using A* Search

- Is GRAPH-SEARCH optimal with A*?




Solution 1: Remember all paths: Need extra bookkeeping

Solution 2: Ensure that the first path to a node is the best!

Consistency (Monotonicity) of heuristic h

- A heuristic is consistent (or monotonic) provided
 - for any node n , for any successor n' generated by action a with cost $c(n,a,n')$
 - $h(n) \leq c(n,a,n') + h(n')$
 - akin to triangle inequality.
 - guarantees admissibility (proof?).
 - values of $f(n)$ along any path are non-decreasing (proof?).
 - Contours of constant f in the state space
- GRAPH-SEARCH using consistent $f(n)$ is optimal.
- Note that $h(n) = 0$ is consistent and admissible.



Next lecture

- Examples
- Choosing heuristics
- Games and Minimax Search

Heuristics

- What's a heuristic for
 - Driving distance (or time) from city A to city B ?
 - 8-puzzle problem ?
 - M&C ?
 - Robot navigation ?
 - Reaching the summit ?
- **Admissible** heuristic
 - Does not overestimate the cost to reach the goal
 - “Optimistic”
- Are the above heuristics admissible? Consistent?

Example: 8-Puzzle

5	4	
6	1	8
7	3	2

Start State

1	2	3
8		4
7	6	5

Goal State

Comparing and combining heuristics

- Heuristics generated by considering relaxed versions of a problem.
- Heuristic h_1 for 8-puzzle
 - Number of out-of-order tiles
- Heuristic h_2 for 8-puzzle
 - Sum of Manhattan distances of each tile
- h_2 dominates h_1 provided $h_2(n) \geq h_1(n)$.
 - h_2 will likely prune more than h_1 .
- $\max(h_1, h_2, \dots, h_n)$ is
 - admissible if each h_i is
 - consistent if each h_i is
- Cost of sub-problems and pattern databases
 - Cost for 4 specific tiles
 - Can these be added for disjoint sets of tiles?

Effective Branching Factor

- Though informed search methods may have poor *worst-case* performance, they often do quite well if the heuristic is good
 - Even if there is a huge branching factor
- One way to quantify the effectiveness of the heuristic: the **effective branching factor, b^***
 - N: total number of nodes expanded
 - d: solution depth
 - $N = 1 + b^* + (b^*)^2 + \dots + (b^*)^d$
- For a good heuristic, b^* is close to 1

Example: 8-puzzle problem

Averaged over 100 trials each at different solution lengths

d	Search Cost			Effective Branching Factor		
	IDS	$A^*(h_1)$	$A^*(h_2)$	IDS	$A^*(h_1)$	$A^*(h_2)$
2	10	6	6	2.45	1.79	1.79
4	112	13	12	2.87	1.48	1.45
6	680	20	18	2.73	1.34	1.30
8	6384	39	25	2.80	1.33	1.24
10	47127	93	39	2.79	1.38	1.22
12	364404	227	73	2.78	1.42	1.24
14	3473941	539	113	2.83	1.44	1.23
16	–	1301	211	–	1.45	1.25
18	–	3056	363	–	1.46	1.26
20	–	7276	676	–	1.47	1.27
22	–	18094	1219	–	1.48	1.28
24	–	39135	1641	–	1.48	1.26

Ave. # of nodes expanded

Solution length

Summary of informed search

- How to use a heuristic function to improve search
 - Greedy Best-first search + Uniform-cost search = A* Search
- When is A* search optimal?
 - h is Admissible (optimistic) for Tree Search
 - h is Consistent for Graph Search
- Choosing heuristic functions
 - A good heuristic function can reduce time/space cost of search by orders of magnitude.
 - Good heuristic function may take longer to evaluate.