

# AI: A Modern Approach, Chpts. 3-4

## Russell and Norvig

Sequential Decision Making in Robotics

CS 599

Geoffrey Hollinger and Gaurav Sukhatme

(Some slide content from Stuart Russell and Hwee Tou Ng)

Spring, 2010

# Decision Making Problem


```
function SIMPLE-PROBLEM-SOLVING-AGENT(percept) returns an action
  static: seq, an action sequence, initially empty
           state, some description of the current world state
           goal, a goal, initially null
           problem, a problem formulation
  state ← UPDATE-STATE(state, percept)
  if seq is empty then
    goal ← FORMULATE-GOAL(state)
    problem ← FORMULATE-PROBLEM(state, goal)
    seq ← SEARCH(problem)
  action ← RECOMMENDATION(seq, state)
  seq ← REMAINDER(seq, state)
  return action
```

# Problem types

- Deterministic, fully observable
  - Agent knows exactly which state it will be in
  - solution is a sequence
- Non-observable
  - Agent may have no idea where it is
  - solution (if any) is a sequence
- Partially observable
  - observations provide new information about current state
  - solution is a contingent plan or a policy
  - often interleave search, execution
- Unknown state space
  - exploration problem


# Vacuum world

- Observable, start in #5.
  - [Right; Suck]
- Non-observable, start in: {1;2;3; 4;5;6;7;8}  
e.g., Right goes to {2;4;6;8}
  - [Right; Suck; Left; Suck]
- Partially observable, start in #5, suck can dirty a clean carpet, local sensing only
  - [Right; if dirt then Suck]




- Possible actions:  
left, right, suck

# Vacuum world state space




- States:
- Successor function:
- Actions:
- Goal test:
- Path cost:

# Vacuum world state space



- States: cross product of dirtiness and robot locations
- Successor function: Left/Right changes location, suck changes dirtiness
- Actions: Left, Right, Suck, NoOp
- Goal test: no dirt
- Path cost: 1 per action (0 for NoOp)

# Non-observable vacuum world



# 8-puzzle state space

7	2	4
5		6
8	3	1

Start State

1	2	3
4	5	6
7	8	

Goal State

- States:
- State transition
- Actions:
- Goal test:
- Path cost:



[Note: optimal solution of n-Puzzle family is NP-hard]

## 8-puzzle state space

7	2	4
5		6
8	3	1

Start State

1	2	3
4	5	6
7	8	

Goal State

- States: integer locations of tiles (ignore intermediate positions)
- State transitions: tile locations change
- Actions: move tile left, right, up, down
- Goal test: goal state (given)
- Path cost: 1 per move

# Basic tree search algorithm

```
function TREE-SEARCH(problem, strategy) returns a solution, or failure
  initialize the search tree using the initial state of problem
  loop do
    if there are no candidates for expansion then return failure
    choose a leaf node for expansion according to strategy
    if the node contains a goal state then return the corresponding solution
    else expand the node and add the resulting nodes to the search tree
  end
```

Basic idea:

offline, simulated exploration of state space by generating successors of already-explored states (a.k.a. expanding states)

# Search strategies

- A strategy is defined by picking the order of node expansion
- Strategies are evaluated along the following dimensions:
  - Completeness: does it always find a solution if one exists?
  - Time complexity: number of nodes generated/expanded
  - Space complexity: maximum number of nodes in memory
  - Optimality: does it always find a least-cost solution?
- Time and space complexity are measured in terms of
  - $b$ : maximum branching factor of the search tree
  - $d$ : depth of the least-cost solution
  - $m$ : maximum depth of the state space (may be 1)
- For instance, a strategy might have time or space complexity  $O(bd)$  or  $O(b^d)$

# Uninformed search strategies

- Breadth-first search
  - Uniform-cost search
  - Depth-first search
  - Depth-limited search
  - Iterative deepening search
- $b$  = branching factor
  - $d$  = solution depth
  - $l$  = depth limit
  - $m$  = max tree depth
  - $C^*$  = cost of optimal
  - $\epsilon$  = min action cost

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

# Avoiding repeated states

```
function GRAPH-SEARCH(problem, fringe) returns a solution, or failure
  closed ← an empty set
  fringe ← INSERT(MAKE-NODE(INITIAL-STATE[problem]), fringe)
  loop do
    if fringe is empty then return failure
    node ← REMOVE-FRONT(fringe)
    if GOAL-TEST(problem, STATE[node]) then return node
    if STATE[node] is not in closed then
      add STATE[node] to closed
      fringe ← INSERTALL(EXPAND(node, problem), fringe)
  end
```

Key idea:


Maintaining a closed list can be exponentially more efficient

# Best-first search

- Idea: use an evaluation function for each node
  - estimate of “desirability“
- Expand most desirable unexpanded node
  
- Implementation:
  - fringe is a queue sorted in decreasing order of desirability
  
- Special cases:
  - Greedy search
  - A\* search

# Map of Romania

- Straight line distances provide estimate of cost to goal

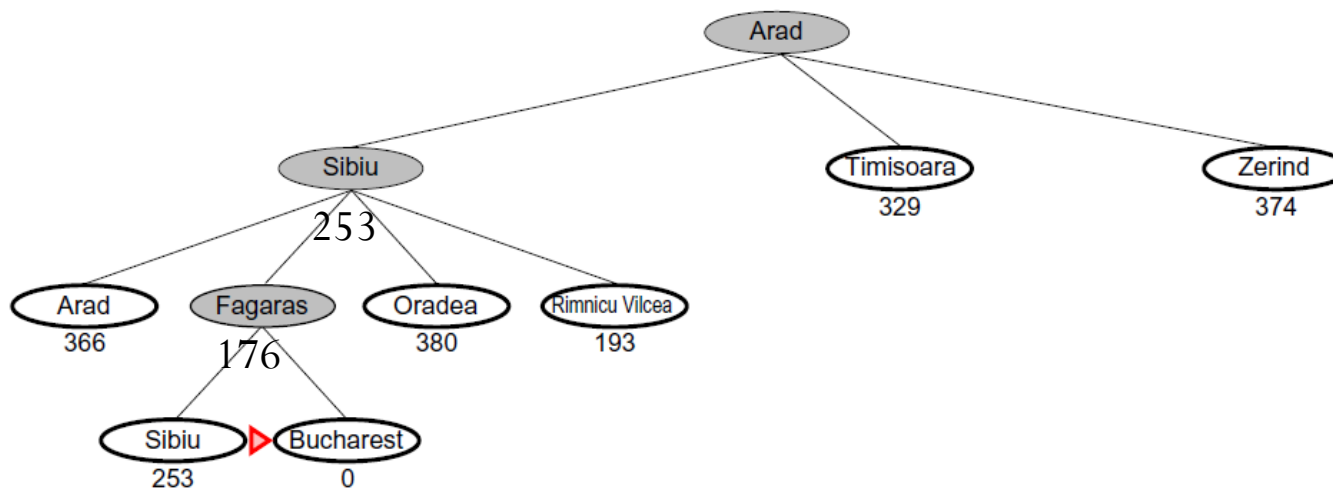


Straight-line distance  
to Bucharest

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

# Greedy search

- Evaluation function  $h(n)$  (heuristic)
  - estimate of cost from  $n$  to the closest goal
- E.g.,  $h_{\text{SLD}}(n) =$  straight-line distance from  $n$  to Bucharest
- Greedy search expands the node that appears to be closest to goal





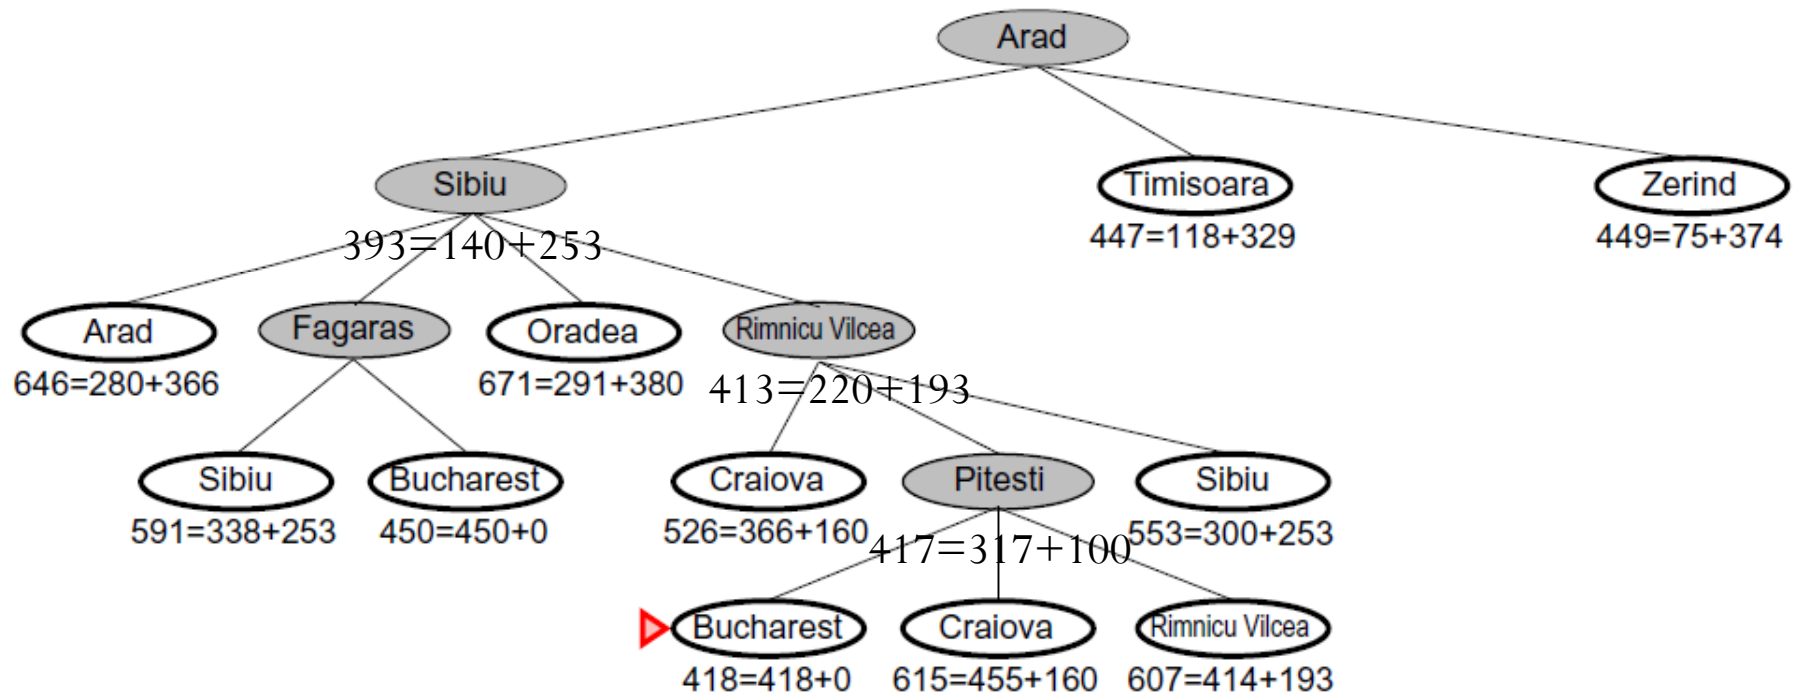
# Greedy search

- Complete: No
  - Can get stuck in loops
  - Complete in finite space with repeated-state checking
- Time:  $O(b^m)$  but better with a good heuristic
- Space:  $O(b^m)$  keeps all nodes in memory
- Optimal: No

# 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  $n$
  - $h(n) =$  estimated cost to goal from  $n$
  - $f(n) =$  estimated total cost of path through  $n$  to goal
- A\* search uses an admissible heuristic
  - i.e.,  $h(n) \leq h^*(n)$  where  $h^*(n)$  is the true cost from  $n$ .  
(Also require  $h(n) \geq 0$ , so  $h(G) = 0$  for any goal  $G$ .)
  - e.g.,  $h_{\text{SLD}}(n)$  never overestimates the actual road distance
- One dirty trick
  - Inflating the heuristic

# A\* example



# A\* Optimality 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  $G$ .

$$f(G_2) = g(G_2) \quad \text{since } h(G_2) = 0$$

$$g(G_2) > g(G) \quad \text{since } G_2 \text{ is suboptimal}$$


$$f(G) = g(G) \quad \text{since } h(G) = 0$$

$$f(G_2) > f(G) \quad \text{from above}$$

$$h(n) \leq h^*(n) \quad \text{since } h \text{ is admissible}$$

$$g(n) + h(n) \leq g(n) + h^*(n)$$

$$f(n) \leq f(G)$$



Hence  $f(G_2) > f(n)$ , and A\* will not have selected  $G_2$  for expansion

# A\* Optimality proof

- Complete: Yes
  - unless there are infinitely many nodes with  $f \leq f(G)$
- Time: Exponential in [relative error in  $h$  \* length of soln.]
- Space: Keeps all nodes in memory
- Optimal: Yes
  
- A\* expands all nodes with  $f(n) < C$
- A\* expands some nodes with  $f(n) = C$
- A\* expands no nodes with  $f(n) > C$

# 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)

7	2	4
5		6
8	3	1

Start State

1	2	3
4	5	6
7	8	

Goal State

- If  $h_2(n) \leq h_1(n)$  for all  $n$  (both admissible) then  $h_2$  *dominates*  $h_1$  and is better for search
- Given any admissible heuristics  $h_a, h_b$ ,  
 $h(n) = \max(h_a(n); h_b(n))$ , is also admissible and dominates  $h_a, h_b$

# Which heuristics are admissible?

- Consider an 8-connected grid
  - Euclidean distance to goal
  - Manhattan distance to goal
  - Euclidean distance divided by 2
  - Euclidean distance times 2
  - Distance to goal along optimal path
  
- Which heuristics dominate?


# 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



# Relaxed version of TSP

- Well-known example: travelling salesperson problem (TSP)
- Find the shortest tour visiting all cities exactly once
- Minimum spanning tree can be computed in  $O(n^2)$  and is a lower bound on the shortest (open) tour



- This also leads to a factor 2 approximation algorithm (much more on this later)

# Fun questions

- Iterative deepening  $A^*$  with a zero heuristic reduces to?
- $A^*$  with a zero heuristic and uniform costs reduces to?

# Important things to understand

- Completeness
- Optimality
- Time/space complexity
- State and action spaces
- State transitions
- Full, partial, and non-observability
- Informed/admissible heuristics
- Relaxed problems
  
- To look up if you don't understand:
  - Complexity classes (NP-hard, etc.)