

Note

- Homework #1
 - Accessible in *Canvas's* "Assignment" section, click the link
 - On Lecture 1-3
 - Due in one week (2/18)
 - Save your hand-written or typed answers into a single PDF file, combining all pages in order of questions. Use any PDF scan app (e.g., CamScanner) or multifunction printers.
 - Submit the PDF file to "Submission for HW #1" link by 2/18 4pm. **No late submission allowed!** Strictly applied.
 - Must check all pages are present and readable after you submit!

KR-PF-PS

PF/PS: Search Methods

CSC 872

Pattern Analysis and Machine Intelligence

Review

Classic AI!


- Last Lecture: Agent-based AI (KR&PF in AI)
 - Learned how to formulate a problem as an AI agent
 - View as the cycle of Percept-Reason-Action interacting with Environment
 - Environment types: PEAS
 - Agent types:
 - simple and model-based reflex agents
 - goal- and utility-based agents
 - learning agents
- Today
 - We will look at one instance of **actual implementation** of the agent-based program for Goal- and Utility-based ones

CSC872: PAMI – Kazunori Okada (C) 2025

3

3

Measuring with Bucket



Problem: Using these three buckets,
measure 7 liters of water.

CSC872: PAMI – Kazunori Okada (C) 2025

From CSC561@USC


4

4

Measuring with Bucket

A Solution:

a	b	c	
0	0	0	start



CSC872: PAMI – Kazunori Okada (C) 2025


5

5

Measuring with Bucket

Another Solution:

a	b	c	
0	0	0	start



CSC872: PAMI – Kazunori Okada (C) 2025

6

6

Which solution is preferred?

- Solution 1:**

a	b	c	
0	0	0	start
3	0	0	
0	0	3	
3	0	3	
0	0	6	
3	0	6	
0	3	6	
3	3	6	
1	5	6	
0	5	7	goal

- Solution 2:**

a	b	c	
0	0	0	start
0	5	0	
3	2	0	
3	0	2	
3	5	2	
3	0	7	goal

CSC872: PAMI – Kazunori Okada (C) 2025

7

7

Problem-Solving Agent

- Four general steps to design this type of agent:

- Goal formulation**

- What are the successful world states

- Problem formulation**

- What actions and states to consider given the goal

- Search strategy (Find Solution)**

- Determine the possible sequence of actions that lead to the states of known values and then choosing the best sequence.

- Execute**

- Give the solution perform the actions.

Perception

Reasoning

Action

CSC872: PAMI – Kazunori Okada (C) 2025

8

8

Problem-Solving Agent

```
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 do
    goal ← FORMULATE-GOAL(state)
    problem ← FORMULATE-PROBLEM(state, goal)
    seq ← SEARCH(problem)
  action ← FIRST(seq)
  seq ← REST(seq)
  return action
```

CSC872: PAMI – Kazunori Okada (C) 2025

9

9

Example: Measuring with Bucket

- **Formulate goal:**
 - *Have 7 liters of water in 9-liter bucket*
- **Formulate problem:**
 - States: *amount of water in the 3 buckets*
 - Operators: *fill bucket from source, empty bucket to others*
- **Find Solution:**
 - *Sequence of operators that bring you from current state (0,0,0) to the goal state (x,x,7)*


CSC872: PAMI – Kazunori Okada (C) 2025

10

10

Problem Types

- **Deterministic, fully observable**
 - **single-state problem (chess) (buckete)**
 - Agent knows exactly which state it will be in; solution is a sequence
- **Non-observable** *Dead Recovery*
 - **sensorless problem (walking in dark)**
 - Agent may have no idea where it is; solution is a sequence
- **Nondeterministic and/or partially observable**
 - **contingency problem (poker)**
 - percepts provide **new** information about current state
 - often **interleave** search and execution
- **Unknown state space**
 - **exploration problem (maze)**



CSC872: PAMI – Kazunori Okada (C) 2025

11

11

Toy Problem/Model

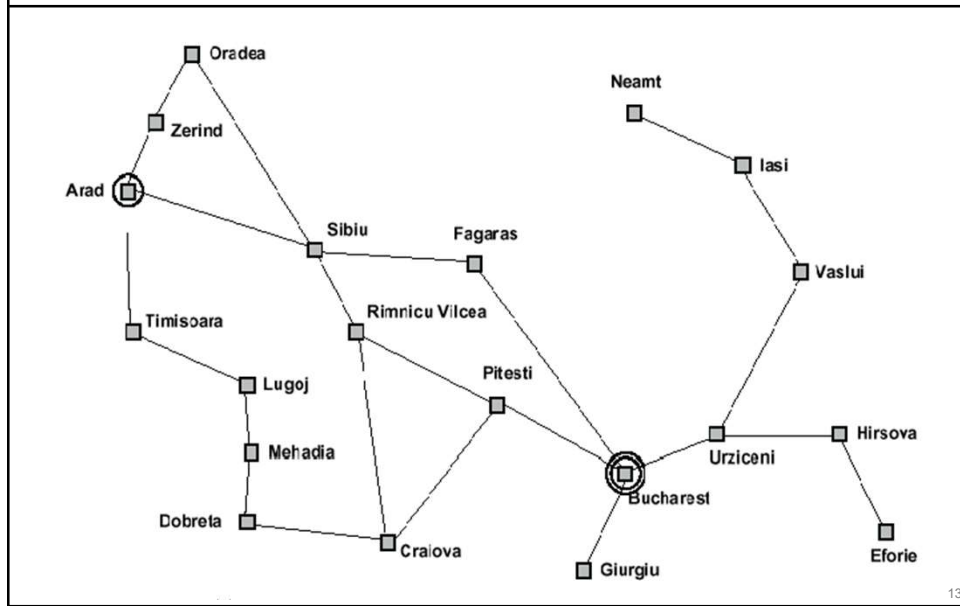
- Intended to illustrate or exercise various methods with concise and exact description
 - Vacuum World
 - Measuring with Buckets
 - ...
- Real-World Problem is the one we want to solve but often hard to describe and solve
 - Robot navigation
 - Playing the game of Go
- Toy problem is used to explore and understand behavior of an algorithm for certain type of problem

CSC872: PAMI – Kazunori Okada (C) 2025

12


12

Toy Problem: Romania



13

Toy Problem: 8-puzzle



14

Selecting a State Space

- Real world is absurdly complex
 - State space must be **abstracted** for problem solving
- Abstracting a set of real **states**
 - "Arad" or "Zerind" represents a complex multi-aspect real city whose boundary may be difficult to define.
- Abstracting a complex combination of real **actions**
 - **Abstraction is to say "going from the city A to B costs L_{AB} "** rather than actually driving from A to B on possible routes, detours, rest stops etc.
- Abstracting a set of real paths that are **solutions** in the real world
 - What is true in the abstracted state space must also be true in the real world (**correctness**).
- Finding the **right level** of abstraction is difficult
- Each abstraction should be "**easier**" than the original problem



CSC872: PAMI – Kazunori Okada (C) 2025

15

15

Problem Formulation


- A **problem** is defined by four items given a state space & a goal:
 - **initial state:**
 - e.g., "at Arad"
 - **operator** (or successor function $S(x)$):
 - e.g., "Arad \rightarrow Zerind" and "Arad \rightarrow Sibiu" etc
 - **goal test:**
 - Explicit: "at Bucharest?"
 - Implicit: Checkmate(x)
 - **path cost** (additive: how long traveled?):
 - e.g., "the sum of distances" and "number of operators applied" etc
- A **solution** is a sequence of operators leading from the initial state to a goal state

Handwritten notes in red:

Arad \rightarrow Sibiu

Sibiu \rightarrow Fagaras

Fagaras \rightarrow Bucharest




CSC872: PAMI – Kazunori Okada (C) 2025


16

16

Example: 8-puzzle



Start State



Goal State


- states?
- actions?
- goal test?
- path cost?
-

CSC872: PAMI – Kazunori Okada (C) 2025

17

17

STOP: Example: Robot Hand




- states?
- actions?
- goal test?
- path cost?
-

CSC872: PAMI – Kazunori Okada (C) 2025

18

18

Search (Finding Solutions)



Basic idea: **offline**, systematic exploration of **simulated** state-space by generating successors of explored states (expanding)



Function General-Search(*problem*, *strategy*) returns a *solution*, or failure
initialize the search tree using the initial state

loop do

if no more candidates for expansion **then return** failure

choose a leaf node for expansion according to the **strategy**

if the node contains a goal state **then return** the corresponding solution

else expand the node and add resulting nodes to the search tree

end

Strategy: the order of node expansion

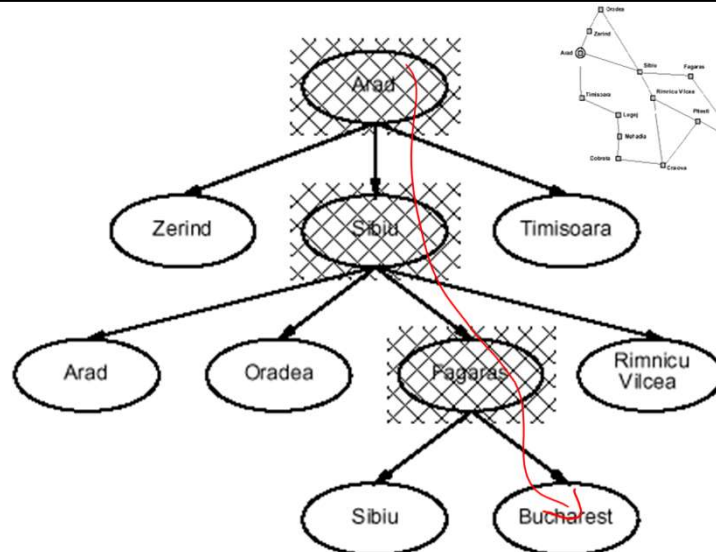
Solution: a sequence from initial to goal states

CSC872: PAMI – Kazunori Okada (C) 2025

19

19

Tree Search Example




CSC872: PAMI – Kazunori Okada (C) 2025

20

20

State Space vs. Search Tree

- Tree node encapsulates state information
 - Node: State, Parent, Action, Depth, Path-Cost
 - **Expand**: create new nodes
 - **Operator**: create corresponding state



CSC872: PAMI – Kazunori Okada (C) 2025

21

21

Search Strategy

- **Order of node expansion** defines a search strategy
- Strategies are evaluated in terms of:
 - **completeness**: does it always find a solution if one exists?
 - **time complexity**: number of nodes generated
 - **space complexity**: maximum number of nodes in memory
 - **optimality**: does it always find a least-cost solution?
- Time and space complexity are measured in:
 - **b**: maximum branching factor of the search tree
 - **d**: depth of the least-cost solution
 - **m**: maximum depth of the state space (may be ∞)

CSC872: PAMI – Kazunori Okada (C) 2025

22

22

Uninformed Search

Search
Strategy

- Use only information available in the problem formulation (Blind Search)
- **Breadth-first**
- **Uniform-cost**
- **Depth-first**
- Depth-limited
- Iterative deepening
- Bidirectional

CSC872: PAMI – Kazunori Okada (C) 2025


23

23

Breadth-First Search

BFS

- Expand *shallowest* unexpanded node
- Implementation: build a FIFO queue



CSC872: PAMI – Kazunori Okada (C) 2025

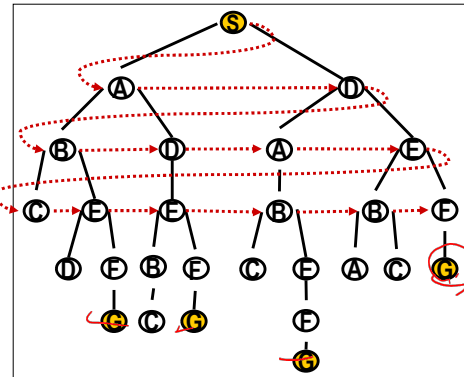
24

24

Breadth-First Search BFS

Completeness: Yes, if b is finite
 Time complexity: $1+b+b^2+\dots+b^d = O(b^d)$, i.e., exponential in d
 Space complexity: $O(b^d)$ all visited must be stored
 Optimality: Yes (assuming cost = 1 per step)

Move downwards level by level, until goal is reached



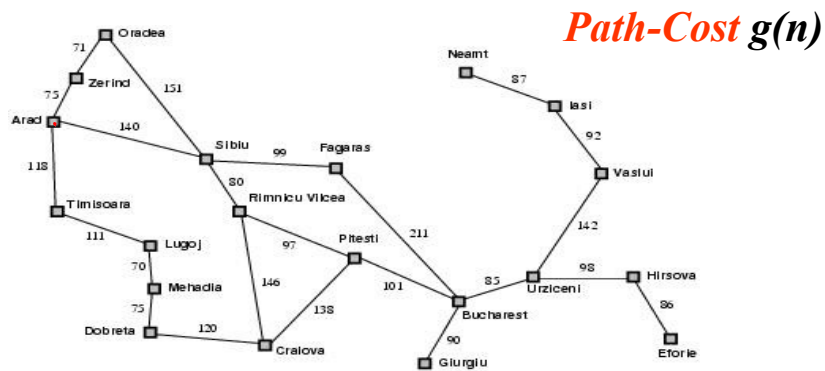
CSC872: PAMI – Kazunori Okada (C) 2025

25

25

Uniform-Cost Search

- Expand node with *lowest path-cost*
- Implementation: a queue sorted by path-cost




CSC872: PAMI – Kazunori Okada (C) 2025

26

26

Uniform-Cost Search

Completeness: Yes, if step cost $\geq \epsilon > 0$
 Time complexity: $\leq O(b^d)$
 Space complexity: $\leq O(b^d)$
 Optimality: Yes, as long as path cost never decreases




CSC872: PAMI – Kazunori Okada (C) 2025

27

27

Depth-First Search DFS

- Expand *deepest* unexpanded node
- Implementation: build a LIFO queue (stack)



I.e., depth-first search can perform infinite cyclic excursions
 Need a finite, non-cyclic search space (or repeated-state checking)

CSC872: PAMI – Kazunori Okada (C) 2025


28

28

Depth-First Search

Completeness: No, fails in infinite or cyclic state-space
Time complexity: $O(b^m)$
Space complexity: $O(bm)$
Optimality: No

Move downwards
as deep as you can,
then back up



CSC872: PAMI – Kazunori Okada (C) 2025

29

29

Informed Search

- Use problem-specific heuristic to guide search
- Utility-based vs Goal-based Agent
- **Best-First Search**
- **Greedy Search**
- **A* search**
- Local Search (Revisited Later)
 - Hill-Climbing
 - Simulated Annealing
 - Local Beam Search

CSC872: PAMI – Kazunori Okada (C) 2025

30

30

Best-First Search

Utility Function


- Idea: Use *evaluation function* $f(n)$ to estimate *desirability* of each node
- Expand node that *appears* best (most desirable)
- Implementation: a queue sorted by *desirability*
- Special Case of $f(n)$
 - Greedy Search
 - A* Search

CSC872: PAMI – Kazunori Okada (C) 2025

31

31

Heuristics



- [dictionary] “A rule of thumb, simplification, or educated guess that reduces or limits the search for solutions in domains that are difficult and poorly understood.”
- $h(n)$ = **estimated** cost of the cheapest path from node n to goal node.
- If n is goal then $h(n)=0$

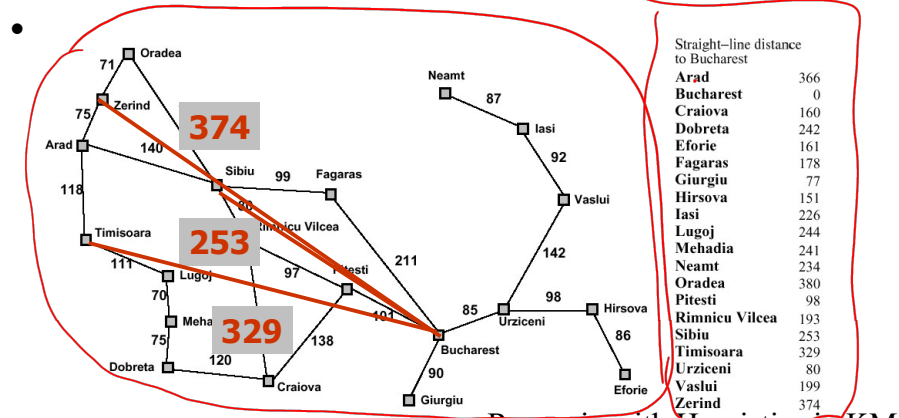
CSC872: PAMI – Kazunori Okada (C) 2025

32

32

Straight-Line Heuristics

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



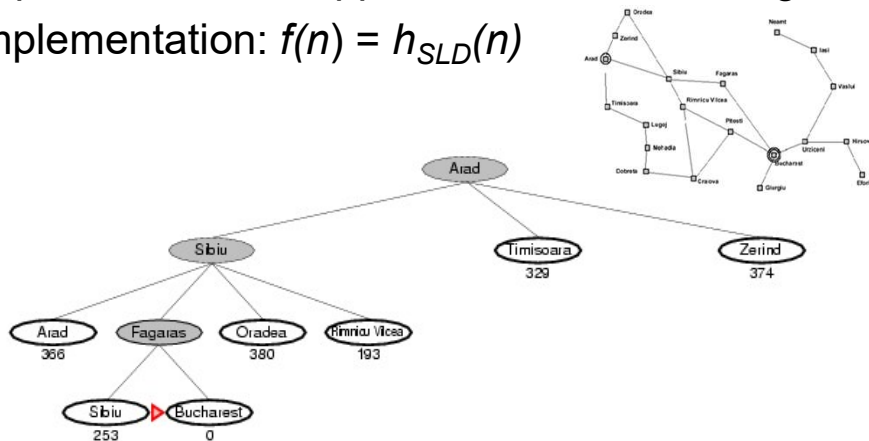
CSC872: PAMI – Kazunori Okada (C) 2025

Romania with Heuristics in KM

33

Greedy Search

- Expand node that appears to be *closest* to goal
- Implementation: $f(n) = h_{SLD}(n)$




CSC872: PAMI – Kazunori Okada (C) 2025

34

34

Greedy Search

Completeness: **No** (cf. DF-search)
 Time complexity: $O(b^m)$ but good heuristic can improve this
 Space complexity: $O(b^m)$ keep all nodes in memory
 Optimality? **No** (cf. DF-search)



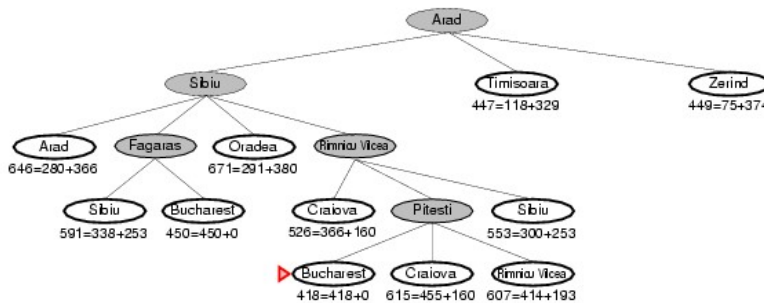
CSC872: PAMI – Kazunori Okada (C) 2025

35

35

A* Search

- Avoid expanding paths that are already expensive
- Implementation: $f(n) = g(n) + h(n)$
past *future*
path cost



CSC872: PAMI – Kazunori Okada (C) 2025

36

36

Admissible Heuristics

- A heuristic is admissible if it *never overestimates* the true cost to reach a goal
- $h(n) \leq h^*(n)$ for all n where $h^*(n)$ is the *true cost* from n .
 - $h_{SLD}(n)$ is admissible because it never overestimates actual road distance.
- Admissible heuristic is *optimistic*



CSC872: PAMI – Kazunori Okada (C) 2025

37

37

Optimality of A*

- **Theorem:** If $h(n)$ is admissible, A* using TREE-SEARCH is *optimal*
- Complete?
 - **Yes** (unless there are infinitely many nodes with $f \leq f(G)$)
- Time?
 - Exponential in length of solution
- Space?
 - Keeps all nodes in memory
- Optimal?
 - **Yes** if $h(n)$ is **admissible**

CSC872: PAMI – Kazunori Okada (C) 2025

38

38

Summary

- Overview
 - PF: Problem-Solving Agent
 - PF: Goal-based Problem Formulation
 - PS: Uninformed Search (Breadth-First, Depth-First)
 - PS: Informed Search (Greedy, A*)
- MATLAB exercise 2 after the break
- Work on HW1!
- Next Lecture
 - PF: Knowledge-based Agent
 - KR: Propositional Logic
 - PF: Logical Inference
 - PS: Resolution, Model Checking, Forward Chaining
 - MATLAB exercise 3

CSC872: PAMI – Kazunori Okada (C) 2025

39